

Lecture 2: August 2, 2017

*Lecturer: Swaprava Nath**Scribe(s): Dhawal Upadhyay*

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

2.1 Introduction

Game theory is the formal study of strategic interaction among multiple agents/players, that have individual utilities/payoffs to maximize. In analyzing a game, we may sometimes encounter counter-intuitive results. First, the stage is set with the discussion of one such problem – the Neighboring Kingdoms’ dilemma. Then we discuss preference relations, and the requirements for the preference relation to have a utility representation.

2.2 The Neighboring Kingdoms’ Dilemma

Suppose there are two kingdoms, each having limited resources. The kingdoms have been hostile for some time, and thus need to boost their defense. Meanwhile, they are also starving, and need to improve their agriculture. Due to limited resources, they can only do one of the two things – boost defense or agriculture.

Let’s call the kingdoms **A** and **B**. Kings of A and B need to make a call between the two options. They don’t know what the other king will choose, and they cannot coordinate. We will call each kingdom *players* in this game. Consider the following payoff table. The rows represent the actions of player A and columns represent the actions of player B. The entries in the matrix are tuples that represent the utilities of the two players.

A\B	Agri	Def
Agri	5,5	0,6
Def	6,0	1,1

Let us say the numbers represent the happiness quotient of the two kingdoms. If both A and B choose agriculture, they both are happy. Thus they get high utilities! If A chooses agriculture and B defense, then B will attack A and loot their resources (A cannot attack because of B’s defense). So B gets 6, and A gets 0. Same happens the other way around. When both choose defense, then although they may not have much to eat, but are still content that nobody can loot them! Thus the utility profile is (1, 1). What should the players choose?

Consider this problem from A’s perspective. If B chooses Agri, then A is better off choosing Def, since the quotient he gets is 6 (compared to 5 if he chooses Agri).

If B chooses Def, then A is better off choosing Def, since the quotient in this case is 1 (compared to 0 if he chooses Agri).

So no matter whether B chooses Agri or Def, A is better off choosing Def. Same happens other way around, since the problem is symmetric from B’s perspective.

Thus both end up choosing Def, and get a quotient of 1 each. If they were allowed to cooperate, they could have chosen Agri each, leading to a better quotient of 5 each.

A similar version of this problem is the Prisoner's Dilemma

2.3 Notation

A set of players $N = \{1, \dots, n\}$.

The set of actions of player i is denoted by A_i , a specific action is denoted by a_i , $a_i \in A_i$.

Utility/payoff of agent i : $u_i = A_1 \times A_2 \times \dots \times A_n \mapsto \mathbb{R}$.

The cartesian product of the action sets will also be represented by $A := \times_{i \in N} A_i$. An element of A is called an *action profile* and is represented by the notations $a \in A$ or $(a_i, a_{-i}) \in A$. The notation a_{-i} represents the action profile of all the players except player i .

When all players choose their respective actions, a_1, \dots, a_n , we say that an *outcome* has realized. Hence, an action profile can be considered as an outcome, say o_ℓ . In the example above, there are four outcomes of the game: (Agri, Agri), (Agri, Def), (Def, Agri), (Def, Def). A *lottery* over the outcomes is a probability distribution over these outcomes, and will be denoted by $[p_1 : o_1, p_2 : o_2, \dots, p_k : o_k]$, where p_i is the probability of occurrence of outcome o_i . Denote the set of outcomes by O .

A relation that signifies the preferences of a player/agent over different outcomes is called a *preference relation*. A preference relation \succeq is a subset of $O \times O$. We will denote o_1 to be at least as preferred as o_2 by $o_1 \succeq o_2$ – which also includes the case that the agent is indifferent between o_1 and o_2 . To make the distinction between strict preference and indifference, we will use the notation $o_1 \succ o_2$ which is same as “ $o_1 \succeq o_2$ but *not* $o_2 \succeq o_1$ ” and $o_1 \sim o_2$ which is same as “ $o_1 \succeq o_2$ and $o_2 \succeq o_1$ ”. In the example above, we have encoded the preference using a utility function, for example, $u_1(o_1) = u_1(\text{Agri, Agri}) = 5$. But it is important to remember that not every preference can be encoded into an utility representation. Consider the following example.

2.3.1 Preference relation without utility representation

Consider a family of three people – child, father, and mother. Three food items a, b, c .

Preference ordering:

Child : $a \succ b \succ c$

Father : $b \succ c \succ a$

Mother : $c \succ a \succ b$

Let's we consider the majority preference order, i.e., the order of items by the majority of the family. If we compare a and b , child and father prefer a , while only mother prefers b . So a is preferred over b by the majority. Similarly, if we do this for pairs (b, c) , we will see b is preferred over c by majority. But while comparing (c, a) , we will arrive at c being more preferred to a . Hence, we get a cycle $(a \succ_{\text{maj}} b \succ_{\text{maj}} c \succ_{\text{maj}} a)$. Clearly, this preference relation does not have a utility representation, as one cannot assign real values that satisfy these inequalities.

Hence, we discuss the conditions that are sufficient for a preference relation to have a utility representation.

2.4 von-Neumann-Morgenstern utility theorem

Axioms for utility representation:

1. **Completeness:** $\forall o_1, o_2 \in O$, exactly one of the following holds: $o_1 \succ o_2$, or $o_2 \succ o_1$ or $o_1 \sim o_2$.
2. **Transitivity:** If $o_1 \succ o_2$ and $o_2 \succ o_3$ then $o_1 \succ o_3$.
3. **Substitutability:** If $o_1 \sim o_2$, then for every sequence o_3, \dots, o_k , and probability masses $p, p_3, \dots, p_k \in [0, 1]$, s.t. $p + \sum_{i=2}^n p_i = 1$, the following holds: $[p : o_1, p_3 : o_3, \dots, p_k : o_k] \sim [p : o_2, p_3 : o_3, \dots, p_k : o_k]$.
4. **Decomposability:** For lotteries l_1, l_2 , if $P_{l_1}(o_i) = P_{l_2}(o_i), \forall o_i \in O$, then $l_1 \sim l_2$.
 Example:
 $l_1 = [0.5 : [0.4 : o_1, 0.6 : o_2], 0.5 : o_3]$
 $l_2 = [0.2 : o_1, 0.3 : o_2, 0.5 : o_3]$
 Here $l_1 \sim l_2$.
5. **Monotonicity:** If $o_1 \succ o_2$, and $1 > p > q > 0$, then
 $[p : o_1, 1 - p : o_2] \succ [q : o_1, 1 - q : o_2]$.
6. **Continuity:** If $o_1 \succ o_2$ and $o_2 \succ o_3$, then $\exists p \in [0, 1]$ s.t. $o_2 \sim l_2 = [p : o_1, 1 - p : o_3]$.

Theorem 2.1 (von-Neumann, Morgenstern (1944)) *If a preference relation \succeq satisfies Axioms 1 to 6, then $\exists u : O \mapsto [0, 1]$ s.t.*

1. $u(o_1) \geq u(o_2) \iff o_1 \succeq o_2$,
2. $u([p_1 : o_1, p_2 : o_2, \dots, p_k : o_k]) = \sum_{i=1}^k (p_i u(o_i))$.

Quite naturally, these utilities are called von-Neumann-Morgenstern (vNM) utilities.

2.5 Summary

We studied six axioms that ensure the existence of a utility representation of a preference relation. These axioms are quite intuitive, and the proof can be found in any standard text (e.g., see [SLB08]). In many real life situations, these axioms are always satisfied, and henceforth, we shall assume the existence of a utility function for any preference that we discuss.

References

- [SLB08] Shoham, Yoav, and Kevin Leyton-Brown. "Multiagent systems: Algorithmic, game-theoretic, and logical foundations". Cambridge University Press, 2008.

Lecture 3: August 4, 2017

Lecturer: Swaprava Nath

Scribe(s): Gaurav Kumar

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

3.1 Normal Form Representation

We revisit the game of neighbouring kingdoms' dilemma. This is a one-shot non-cooperative game. The outcomes (o_i) in this game are as follows: (A,A)= o_1 , (A,D)= o_2 , (D,A)= o_3 , (D,D)= o_4 .

1\2	A	D
A	5,5	0,6
D	6,0	1,1

A *normal form* or *strategic form* representation of a game is given by the tuple $\langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$, where $N = \{1, 2, \dots, n\}$ is the set of players, S_i is the set of strategies of player i , and $u_i : S_1 \times \dots \times S_n \mapsto \mathbb{R}$ is the vNM utility function of agent i . We will denote a specific strategy of player i by $s_i \in S_i$. The strategy profile of all agents except player i is denoted by s_{-i} . A strategy profile is the tuple $s := \{s_1, s_2, \dots, s_n\}$ which is also represented by (s_i, s_{-i}) . Note that $s \in S_1 \times \dots \times S_n = S$. Hence, utility of player i is $u_i(s_1, s_2, \dots, s_n)$, where s_j represents strategy picked by player j , $\forall j \in N$.

Difference between Strategy and Action We distinguish action from strategy as follows. An *action* is the smallest available choice a player has in a game. For instance, it is either to pick action Agri or Def in the game above. But a *strategy* is a more complicated object than an action, since it may involve a combination of actions, or a mixture of actions. For instance, in a multi-round game, a strategy will be a *complete contingent plan* of which action to play at which stage and state of the game. However, in the one-shot non-cooperative game example above, the strategy is to pick an action, and therefore the set of strategies is same as the set of actions. But this distinction is worth remembering.

3.2 Behaviour of players

To predict an outcome of a game, we need some behavioral assumptions on the players. We assume following behaviours of the players :

- **Rationality:** Every player picks strategy to maximize her utility.
- **Intelligence:** Every player possesses enough information about the game and is able to find the best strategy for her.

Assumption 3.1 *Every player is rational and intelligent.*

Common Knowledge: A fact is known as a common knowledge if

1. All players know the fact, and
2. All players know that all other players know the fact, and
3. All players know that all other players know that all other players know the fact, and ... ad infinitum.

Example of common knowledge: Consider an isolated island where lives three blue-eyed individuals (eyes can be either blue or black). The individuals do not talk to each other but can listen and the island has no reflecting media, e.g., mirrors etc., where they can see their own eye color. One day a sage comes to the island and says “Blue-eyed people are bad for the island and must leave. There is at least one blue-eyed person in this island”. Assume that the sage’s statements cannot be disputed. Also, if a person realizes that his eye color is blue, he leaves at the end of the day. Let us see the implication of this statement and how common knowledge percolates to the outcome.

If there were only one blue-eyed person, he would have seen that the other two had black eyes, realized that his eye color is blue (since sage is always correct), leaves at the end of day one. Every other player understands this.

If there were two blue-eyed persons, then both of them will see one blue and one black eyed person, hope that he is not the blue eyed one and wait until the second day if the other blue-eyed person leaves on day one. When it does not happen, he realizes that both of them had blue eyes, so they both leave at the end of day two. Every player understands this.

Since there are three blue-eyed persons, then extending the same argument, we see that every player will wait till day three if anyone leaves. When nobody left on day two, it becomes clear that all of them had blue eyes, and they all leave at the end of day three.

Assumption 3.2 *The fact that all players are rational and intelligent is a Common Knowledge.*

3.3 Some important definitions

- **Strictly dominated strategy:** A strategy s'_i is strictly dominated by s_i , if $\forall s_{-i} \in S_{-i}$

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i}).$$

- **Weakly dominated strategy:** A strategy s'_i is weakly dominated by s_i , if $\forall s_{-i} \in S_{-i}$

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}),$$

and $\exists s_{-i} \in S_{-i}$ such that

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i}).$$

- **Strictly/Weakly dominant strategy:** A strategy s_i is strictly/weakly dominant strategy of player i if s_i strictly/weakly dominates all other $s'_i \in S_i \setminus \{s_i\}$
- **Strictly/Weakly dominant strategy equilibrium:** A strategy profile (s_i^*, s_{-i}^*) is an SDSE/WDSE if s_i^* is a SDS/WDS for every $i, i \in N$.

Does every Game have a SDSE/WDSE? Can there be more than one Nash equilibrium? We will answer these questions through the following example.

1\2	C	F
C	(2,1)	(0,0)
F	(0,0)	(1,2)

This game does not have any SDSE or WDSE. Hence, it is clear that every game is not guaranteed to have a SDSE/WDSE. Hence we come to a weaker notion of equilibrium which is called the Nash equilibrium.

Definition 3.3 (Pure strategy Nash equilibrium) A strategy profile (s_i^*, s_{-i}^*) is a pure strategy Nash equilibrium if $\forall i \in N$ and $\forall s_i \in S_i$

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*).$$

We see that there exists two pure strategy Nash equilibria in the game above: (C,C) and (F,F). However when there exists an SDSE, there is exactly one equilibrium.

Lecture 4: August 8, 2017

Lecturer: Swaprava Nath

Scribe(s): Sachin K Salim

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

Definition 4.1 (Best response set) A best response of agent i against the strategy profile s_{-i} of the other players is a strategy that gives the maximum utility against the s_{-i} chosen by other players, i.e.,

$$B_i(s_{-i}) = \{s_i \in S_i : u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}), \forall s'_i \in S_i\}.$$

Observe: If (s_i^*, s_{-i}^*) is a Pure Strategy Nash Equilibrium, then $s_i^* \in B_i(s_{-i}^*) \forall i \in N$.

We know that an SDSE is a WDSE. To observe the relation between WDSE and PSNE, we recap the definition of WDSE. To define a WDSE, we need the definition of *Weakly Dominant Strategy* (WDS).

Definition 4.2 (Weakly Dominant Strategy) s_i^* is WDS if

1. $u_i(s_i^*, s_{-i}) \geq u_i(s'_i, s_{-i}), \forall s'_i \in S_i, \forall s_{-i} \in S_{-i}$
2. $u_i(s_i^*, \bar{s}_{-i}) > u_i(s'_i, \bar{s}_{-i}), \forall s'_i \in S_i, \text{ for some } \bar{s}_{-i} \in S_{-i}$

It is important to note that when s_i dominates all s'_i in the definition above, the profile of other players on which the strict inequality holds, can be different for different s'_i . Here is an example illustrating this fact. Note here $u_1(D, D) > u_1(A, D)$ and $u_1(D, A) > u_1(S, A)$. D is WDS for P_1 .

		P_2	
		A	D
P_1	A	5,5	0,5
	D	5,0	1,1
	S	4,0	1,1

Table 4.1: Example game to illustrate WDS.

A strategy profile (s_i^*, s_{-i}^*) is a WDSE if s_i^* is a WDS for every $i \in N$. Clearly, a WDSE is a PSNE.

However, in a finite game, even PSNE is not guaranteed to exist. Table 4.2 gives an example. Hence, we arrive at a further weak equilibrium concept named *Mixed Strategy Nash Equilibrium* (MSNE).

		P_2	
		H	T
P_1	H	+1, -1	-1, +1
	T	-1, +1	+1, -1

Table 4.2: Matching Coins Game

4.1 Mixed Strategy Nash Equilibrium

For a finite set A , $\Delta(A)$ is defined as the set of all probability distributions over A , $\Delta(A) = \{p \in [0, 1]^{|A|} : \sum_{a \in A} p(a) = 1\}$. Then $\sigma_i \in \Delta(S_i)$ is a **mixed strategy** of player i , where S_i is their finite strategy set. Mixed strategy is a distribution σ_i over the strategies in S_i , i.e., $\sigma_i : S_i \mapsto [0, 1]$ s.t. $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.

Utility of player i at a mixed strategy profile (σ_i, σ_{-i}) is

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s \in S} \left(\prod_{i \in N} \sigma_i(s_i) \right) u_i(s_i, s_{-i}),$$

where $s = (s_1, \dots, s_n)$ and $S = S_1 \times \dots \times S_n$.

Consider the game as given in Table 4.2. Now suppose Player 1 plays the mixed strategy H with probability p and Player 2 plays H with probability q .

Then the utility u_1 of the player 1 is $u_1((p, 1-p), (q, 1-q))$
 $= pq u_1(H, H) + p(1-q) u_1(H, T) + (1-p)q u_1(T, H) + (1-p)(1-q) u_1(T, T)$

For a mixed strategy profile $\sigma' = ((1, 0), (\frac{1}{2}, \frac{1}{2}))$, $u_1(\sigma') = 1 \cdot \frac{1}{2}(+1) + 1 \cdot \frac{1}{2}(-1) = 0$

When player i plays pure strategy while all others play mixed strategy, we denote the utility of the player by

$$u_i(s_i, \sigma_{-i}) := \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) u_i(s_i, s_{-i}).$$

Definition 4.3 (Mixed Strategy Nash Equilibrium) *MSNE is a mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ s.t.*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma'_i, \sigma_{-i}^*) \quad \forall \sigma'_i \in \Delta(S_i), \quad \forall i \in N.$$

One can define a best response set in terms of mixed strategies in a similar spirit and observe that

$$B_i(\sigma_{-i}) = \{\sigma_i \in \Delta(S_i) : u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}), \quad \forall \sigma'_i \in \Delta(S_i)\},$$

and if $(\sigma_i^*, \sigma_{-i}^*)$ is a MSNE, then $\sigma_i^* \in B_i(\sigma_{-i}^*), \forall i \in N$.

Now since we have seen all the important equilibrium concepts, fig. 4.1 shows how one equilibrium implies another and thereby the Venn-diagram of the different equilibria. Each of the subset implication in this figure is strict. It is easy to construct examples to show the strictness.




Figure 4.1: Types of Equilibrium

4.2 Computation of MSNE

To compute an MSNE, we first state a result that helps in formulating the problem of finding the equilibrium. To do this, we define the support of a mixed strategy as follows.

Definition 4.4 (Support of a Mixed Strategy) *The support of a mixed strategy σ_i is the subset of the strategy space of i on which the mixed strategy σ_i has positive mass, and is denoted by*

$$\delta(\sigma_i) = \{s_i \in S_i : \sigma_i(s_i) > 0\}.$$

Theorem 4.5 (Characterization of a MSNE) *A mixed strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is a MSNE iff $\forall i \in N$*

1. $u_i(s_i, \sigma_{-i}^*)$ is the same for all $s_i \in \delta(\sigma_i^*)$, and
2. $u_i(s_i, \sigma_{-i}^*) \geq u_i(s'_i, \sigma_{-i}^*)$, $\forall s_i \in \delta(\sigma_i^*), s'_i \notin \delta(\sigma_i^*)$.

Lecture 5: August 9, 2017

Lecturer: Swaprava Nath

Scribe(s): Gundeep Arora

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

5.1 Recap

In the last lecture, we saw an example where an MSNE existed but PSNE did not. We also looked at the characterization theorem of MSNE. In this lecture, we shall prove the theorem and compute MSNE using this theorem.

5.2 Characterization Theorem of MSNE

For completeness, we restate the theorem from the last lecture. This theorem lists down the essential characteristics of a mixed strategy Nash equilibrium.

Theorem 5.1 *A strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is a MSNE iff $\forall i \in N$*

1. $u_i(s_i, \sigma_{-i}^*)$ is same $\forall s_i \in \delta(\sigma_i^*)$
2. $u_i(s_i, \sigma_{-i}^*) \geq u_i(s'_i, \sigma_{-i}^*) \quad \forall s_i \in \delta(\sigma_i^*), \quad s'_i \notin \delta(\sigma_i^*)$

Before we prove this theorem, we shall state an observation,

Remark 1 *Clearly, while maximizing the expectation of a random variable with finite support w.r.t. the distribution is achieved when the distribution places the whole probability mass at the maximum value (or splits arbitrarily over the maximum values). Hence*

$$\max_{\sigma_i \in \Delta(S_i)} u_i(\sigma_i, \sigma_{-i}^*) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*).$$

Moreover, if σ_i^* is Nash equilibrium strategy for player i , then the strategy that maximizes the expected utility $u_i(s_i, \sigma_{-i}^*)$ will always be a part of the support of σ_i^* , $\delta(\sigma_i^*)$. Otherwise, we can create another strategy with all the probability mass on this maximum one that falls outside the support and this strategy would have a strictly better utility for player i , contradicting σ_i^* being MSNE. Therefore we have

$$\max_{\sigma_i \in \Delta(S_i)} u_i(\sigma_i, \sigma_{-i}^*) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) = \max_{s_i \in \delta(\sigma_i^*)} u_i(s_i, \sigma_{-i}^*).$$

Proof: (\Rightarrow) We first prove the necessity part, i.e., that given a strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE, the two conditions hold true. Given $(\sigma_i^*, \sigma_{-i}^*)$ is a MSNE

$$\begin{aligned} u_i(\sigma_i^*, \sigma_{-i}^*) &= \max_{\sigma_i \in \Delta(S_i)} u_i(\sigma_i, \sigma_{-i}^*) \\ &= \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) \\ &= \max_{s_i \in \delta(\sigma_i^*)} u_i(s_i, \sigma_{-i}^*) \end{aligned} \quad (5.1)$$

Here the first and second equality follows due to remark 1. Also, by definition of expected utility for the given strategy profile we have

$$\begin{aligned} u_i(\sigma_i^*, \sigma_{-i}^*) &= \sum_{s_i \in S_i} \sigma_i^*(s_i) \cdot u_i(s_i, \sigma_{-i}^*) \\ &= \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i) \cdot u_i(s_i, \sigma_{-i}^*) \end{aligned} \quad (5.2)$$

Equating the 5.1 and 5.2, we see that the expectation and the maximum value of a set are equal. This can happen only when either the set is singleton or all the elements take the same value. This proves the first condition mentioned.

We prove the second condition using the idea of contradiction. Suppose the condition does not hold, i.e.

$$\exists s_i \in \delta(\sigma_i^*), s'_i \notin \delta(\sigma_i^*) \text{ s.t. } u_i(s_i, \sigma_{-i}^*) < u_i(s'_i, \sigma_{-i}^*)$$

According to our previous argument, $u_i(s_i, \sigma_{-i}^*)$ is same for all $s_i \in \delta(\sigma_i^*)$. Hence the LHS of the above inequality is equal to $u_i(\sigma_i^*, \sigma_{-i}^*) =: u_i^*$. Choose a strategy σ'_i for player i , such that

$$\begin{aligned} \sigma'_i(s'_i) &= 1 \\ \sigma'_i(s_i) &= 0, \forall s_i \in S_i \setminus \{s'_i\} \end{aligned} \quad (5.3)$$

Using this mixed strategy, we compute the expected utility for player i at the strategy profile $(\sigma'_i, \sigma_{-i}^*)$ and find

$$u_i(\sigma_i^*, \sigma_{-i}^*) < u_i(\sigma'_i, \sigma_{-i}^*).$$

The above inequality contradicts the fact that $(\sigma_i^*, \sigma_{-i}^*)$ is an MSNE. This proves our second condition as well.

(\Leftarrow) To prove the sufficiency, we assume that given the two conditions of the characterization theorem hold. We define $u_i(s_i, \sigma_{-i}^*) =: m_i(\sigma_{-i}^*)$, for all $s_i \in \delta(\sigma_i^*)$. This is possible to define due to condition (1). Similarly, using (2), we conclude $m_i(\sigma_{-i}^*) = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*)$.

$$\begin{aligned} u_i(\sigma_i^*, \sigma_{-i}^*) &= \sum_{s_i \in \delta(\sigma_i^*)} \sigma_i^*(s_i) \cdot u_i(s_i, \sigma_{-i}^*) \\ &= m_i(\sigma_{-i}^*) \\ &= \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}^*) \\ &= \max_{\sigma_i \in \Delta(S_i)} u_i(\sigma_i, \sigma_{-i}^*) \\ &\geq u_i(\sigma_i, \sigma_{-i}^*) \quad \forall \sigma_i \in \Delta(S_i) \end{aligned} \quad (5.4)$$

The first equality holds by definition of $\delta(\sigma_i^*)$. The next two equalities hold due to conditions (1) and (2) as explained above. The last equality is by remark 1 Hence the strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is a MSNE. \blacksquare

5.3 Some Examples for MSNE

With the proof of the characterization theorem done, we shall now look at some examples with the lens of that theorem and argue whether possible strategy profiles are MSNE or not based on their satisfiability of the two conditions of the theorem.

5.3.1 The Matching Coins Game / Penalty Shoot-out Game

In the penalty shoot-out game, that we looked at last time, we had concluded that no pure strategy Nash equilibrium existed for that game. The game poses the utility functions of either player hitting or missing the target.

1 \ 2	H	T
H	(1,-1)	(-1,1)
T	(-1,1)	(1,-1)

Let the probability of player 1 choosing H be p . Similarly, for player 2, let that probability be q . We shall now look at some of the strategy profiles and argue if it can be an MSNE or not.

- $k = (\{\mathbf{H}\}, \{\mathbf{H}\})$: In this case, $u_1(k) = 1$ and $u_2(k) = -1$. The for the given choice of strategy for player 1, the player 2, can choose $\{\mathbf{H}\}$ and have a utility $u_2\{\mathbf{H}\}, \{\mathbf{H}\} = 1$, which is more than $u_2(k)$. We can establish the same result for each of the other pure strategy profiles will not be a MSNE. This is just what we had seen in the previous lecture, where we had concluded that no pure strategy Nash equilibrium existed for this game. These pure strategy profiles also do not satisfy the second condition of the characterization theorem.

- $k = (\{\mathbf{H}\}, \{\mathbf{H}, \mathbf{T}\})$: In this case, $u_1(k) = 1$ and

$$\begin{aligned} u_2(\{\mathbf{H}\}, \{\mathbf{H}\}) &= -1 = 0 \\ u_2(\{\mathbf{H}\}, \{\mathbf{T}\}) &= 1 \\ &= 2p - 1 \end{aligned} \tag{5.5}$$

For this to be a MSNE, by the first condition of the characterization equation, we have

$$\begin{aligned} u_2(\{\mathbf{H}\}, \{\mathbf{H}\}) &= u_2(\{\mathbf{H}\}, \{\mathbf{T}\}) \\ -1 &\neq 1 \end{aligned} \tag{5.6}$$

This hence cannot be a MSNE. By using the argument of symmetry for this game, we can argue that any strategy profile that has a pure strategy for any user will not be a Nash equilibrium.

- $k = (\{\mathbf{H}, \mathbf{T}\}, \{\mathbf{H}, \mathbf{T}\})$: In this case,

$$\begin{aligned} u_2(\{\mathbf{H}, \mathbf{T}\}, \{\mathbf{H}\}) &= (-1) \cdot p + 1 \cdot (1 - p) \\ &= 1 - 2p \\ u_2(\{\mathbf{H}, \mathbf{T}\}, \{\mathbf{T}\}) &= 1 \cdot p + (-1) \cdot (1 - p) \\ &= 2p - 1 \\ u_1(\{\mathbf{H}\}, \{\mathbf{H}, \mathbf{T}\}) &= 1 \cdot q + (-1) \cdot (1 - q) \\ &= 2q - 1 \\ u_1(\{\mathbf{T}\}, \{\mathbf{H}, \mathbf{T}\}) &= (-1) \cdot q + 1 \cdot (1 - q) \\ &= 1 - 2q \end{aligned} \tag{5.7}$$

For this strategy for this to be a valid MSNE, from the characterization theorem that we just proved, we have the following,

$$\begin{aligned}
 u_2(\{H,T\}, \{H\}) &= u_2(\{H,T\}, \{T\}) \\
 1 - 2q &= 2q - 1 \\
 q &= \frac{1}{2} \\
 u_1(\{H\}, \{H,T\}) &= u_1(\{T\}, \{H,T\}) \\
 2p - 1 &= 1 - 2p \\
 p &= \frac{1}{2}
 \end{aligned} \tag{5.8}$$

Since, the system of equations have a solution, the the strategy profile $(\{H : \frac{1}{2}, T : \frac{1}{2}\}, \{H : \frac{1}{2}, T : \frac{1}{2}\})$ is a MSNE. The first condition is vacuously satisfied.

5.3.2 Game-Selection Problem

We also looked at the problem, where the two friends who had different preferences towards going for a game while still going together to have a good time. More formally, the game with the utility function is represented as :

1 \ 2	F	C
F	(2,1)	(0,0)
C	(0,0)	(1,2)

Like the previous example, let the probability of player 1 choosing H be p . Similarly, for player 2, let that probability be q . We shall now look at some of the strategy profiles and argue if it can be an MSNE or not.

- **Pure strategy** : Consider a pure strategy, $k = (\{F\}, \{F\})$. For this strategy profile, we have $u_1(k) = 2$ and $u_2(k) = 1$. This is a PSNE, hence definitely a MSNE. As we had seen in the previous lecture, the profile $k = (\{C\}, \{C\})$ is also a MSNE. Both these profiles trivially satisfy the two conditions of the characterization theorem.
- $(\{F\}, \{F,C\})$: For this strategy profile,

$$\begin{aligned}
 u_2(\{F\}, \{F\}) &= 1 \\
 u_2(\{F\}, \{C\}) &= 0
 \end{aligned} \tag{5.9}$$

Clearly the above equations state that the first condition to be is violated and hence this profile is not a MSNE.

- $(\{C\}, \{F,C\})$: For this strategy profile,

$$\begin{aligned}
 u_2(\{C\}, \{F\}) &= 0 \\
 u_2(\{C\}, \{C\}) &= 2
 \end{aligned} \tag{5.10}$$

Clearly the above equations state that the first condition to be is violated and hence this profile is not a MSNE.

- $(\{F,C\},\{F\})$: For this strategy profile,

$$\begin{aligned} u_1(\{F\},\{F\}) &= 2 \\ u_1(\{C\},\{F\}) &= 0 \end{aligned} \tag{5.11}$$

Clearly the above equations state that the first condition to be is violated and hence this profile is not a MSNE.

- $(\{F,C\},\{C\})$: For this strategy profile,

$$\begin{aligned} u_1(\{C\},\{C\}) &= 1 \\ u_1(\{F\},\{C\}) &= 0 \end{aligned} \tag{5.12}$$

Clearly the above equations state that the first condition to be is violated and hence this profile is not a MSNE.

- $(\{F,C\},\{F,C\})$: In this case,

$$\begin{aligned} u_2(\{F,C\},\{F\}) &= 1 \cdot p + 0 \cdot (1 - p) \\ &= p \\ u_2(\{F,C\},\{C\}) &= 0 \cdot p + 2 \cdot (1 - p) \\ &= 2 - 2p \\ u_1(\{F\},\{F,C\}) &= 2 \cdot q + 0 \cdot (1 - q) \\ &= 2q \\ u_1(\{C\},\{F,C\}) &= 0 \cdot q + 1 \cdot (1 - q) \\ &= 1 - q \end{aligned} \tag{5.13}$$

For this strategy for this to be a valid MSNE, from the characterization theorem that we just proved, we have the following,

$$\begin{aligned} u_2(\{F,C\},\{F\}) &= u_2(\{F,C\},\{C\}) \\ p &= 2 - 2p \\ p &= \frac{1}{3} \\ u_1(\{F\},\{F,C\}) &= u_1(\{C\},\{F,C\}) \\ 2q &= 1 - q \\ p &= \frac{1}{3} \end{aligned} \tag{5.14}$$

Since, the system of equations have a solution, the the strategy profile $(\{F : \frac{1}{3}, T : \frac{2}{3}\}, \{F : \frac{2}{3}, C : \frac{1}{3}\})$ is a MSNE. The first condition is vacuously satisfied.

This problem can be extended to have the following form,

1 \ 2	F	C	D
F	(2,1)	(0,0)	(1,1)
C	(0,0)	(1,2)	(2,0)

For such a case, a mixed strategy for player two, will generally look like, $\{F : q_1; C : q_2; D : 1 - q_1 - q_2\}$. Proceeding similar to the original example, we have,

$$\begin{aligned} u_1(\{F\}, \{F, C, D\}) &= u_1(\{C\}, \{F, C, D\}) \\ 2 \cdot q_1 + 0 \cdot q_2 + 1 \cdot (1 - q_1 - q_2) &= 0 \cdot q_1 + 1 \cdot q_2 + 2 \cdot (1 - q_1 - q_2) \\ 1 + q_1 - q_2 &= 2 - 2q_1 - q_2 \\ q_1 &= \frac{1}{3} \end{aligned} \tag{5.15}$$

However, we do not have any specific value for q_2 , so any value of q_2 between $[0, \frac{2}{3}]$ will yield an MSNE. The solution for the player 1 remains the same.

5.4 General Principle for finding MSNE

As we saw in the two examples above, in order to evaluate the Nash equilibrium. we enumerated all the possible supports of the Cartesian product, $S_1 \times S_2 \times \dots \times S_n$, where S_i is the set of all possible options for player i and then used the two conditions to check if it was a MSNE or not. This means that the number of supports, K , that need to be enumerated is,

$$K = (2^{|S_1|} - 1) \times (2^{|S_2|} - 1) \times \dots \times (2^{|S_n|} - 1)$$

More formally put, the problem that we are trying to solve here is, given a support $X_i \subseteq S_i$ and a support profile $X_1 \times X_2 \times \dots \times X_n$, we have from the first condition of the characterization theorem, i.e., $u_i(s_i, \sigma_{-i})$ should be same for all $s_i \in \delta(\sigma_i)$

$$w_i = \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) u_i(s_i, s_{-i}), \forall s_i \in X_i, \forall i \in N.$$

From the first and second conditions together, we get

$$w_i \geq \sum_{s_{-i} \in S_{-i}} \left(\prod_{j \neq i} \sigma_j(s_j) \right) u_i(s_i, s_{-i}), \forall s_i \in S_i \setminus X_i, \forall i \in N.$$

These equalities and inequalities constitute a feasibility program with variables $w_i, i \in N, \sigma_j(s_j), s_j \in S_j, j \in N$. The $\sigma_j(s_j)$'s must satisfy $\sigma_j(s_j) \geq 0, s_j \in S_j, j \in N$ and $\sum_{s_j \in S_j} \sigma_j(s_j) = 1, \forall j \in N$. This is a linear programming problem iff $n = 2$. For $n > 2$, the first set of equality and inequalities are non-linear. To find an MSNE, a brute-force algorithm is to list these set of equalities and inequalities for every support profile in K and solve them. Note that K is exponential in the size of the strategy spaces. Unfortunately, the problem of finding Nash equilibrium is PPAD complete [DGP08], which implies that it is unlikely to have a better algorithm to find MSNE in general games.

Finding Nash equilibrium and its complexity is an active area of research. For two player zero-sum game, we can have a linear program, the equalities to be solved turn out to be simple and hence we have efficient algorithms for it. There are algorithms like Lemke-Howson Algorithms that are proven to be optimal for two player non-zero sum games. We shall look at the existence of the Nash equilibrium in the next lecture.

References

- [MAS] YOAV SHOHAM and KEVIN LEYTON-BROWN, "Multi Agent Systems Algorithmic, Game-Theoretic, and Logical Foundations,"

- [Narahari] Y. NARAHARI, “Mixed Strategies and Mixed Strategy Nash Equilibrium,” *Game Theory, Lecture Notes*, 2012
- [DGP08] Daskalakis, Constantinos, Paul W. Goldberg, and Christos H. Papadimitriou. “The complexity of computing a Nash equilibrium.” *SIAM Journal on Computing* 39.1 (2009): 195-259.

Lecture 6: August 11, 2017

Lecturer: Swaprava Nath

Scribe(s): Anil Kumar Gupta

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.uitk.ac.in.*

6.1 Introduction

We have seen discussions on computing Nash equilibrium. In this lecture, we will address a more fundamental question: the existence of a mixed Nash equilibrium. Nash showed that MSNE exists in any finite game. To prove this result, we will use a result from real analysis. First, we discuss some basic definitions of sets that will be used in presenting the result.

6.2 Definitions and Standard Results

- A set $S \subseteq \mathbb{R}^n$ is **convex** if $\forall x, y \in S$ and $\forall \lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in S$.
- A set $S \subseteq \mathbb{R}^n$ is **closed** if it contains all its limit points (points whose every neighborhood contains a point in S – e.g., for the point 1 in the interval $[0, 1)$, consider a ball of radius $\epsilon > 0$, arbitrary, clearly, each such ball will contain a point in $[0, 1)$).
- A set $S \subseteq \mathbb{R}^n$ is **bounded** if $\exists x_0 \in \mathbb{R}^n$ and $R \in (0, \infty)$ such that $\forall x \in S$, $\|x - x_0\|_2 < R$.
- A set $S \subseteq \mathbb{R}^n$ is **compact** if it is *closed* and *bounded*.

Now we state the result from real analysis without proof.

Theorem 6.1 (Brouwer’s Fixed Point Theorem) *If $S \subseteq \mathbb{R}^n$ is convex and compact and $T : S \mapsto S$ is continuous, then T has a fixed point, i.e., \exists a point $x^* \in S$ s.t. $T(x^*) = x^*$.*

6.3 Existence of MSNE

Finite game: A game in which the number of players and the strategies are finite.

Theorem 6.2 (Nash (1951)) *Every finite game has a (mixed) Nash equilibrium.*

Proof: Define simplex to be

$$\Delta_k = \{x \in \mathbb{R}_{\geq 0}^{k+1} : \sum_{i=1}^{k+1} x_i = 1\}.$$

Clearly, this is a convex and compact set in \mathbb{R}^{k+1} . Consider two players (the case with n players is an extension of this idea). Say, player 1 has m strategies labeled $1, \dots, m$ and player 2 has n strategies labeled

$1, \dots, n$. So, player 1's mixed strategy is a point in Δ_{m-1} and player 2's mixed strategy is a point in Δ_{n-1} . The set of mixed strategy profiles is a point in $\Delta_{m-1} \times \Delta_{n-1}$. Since we are in a two player game, the utilities can be expressed in terms of two matrices A and B , both in $\mathbb{R}^{m \times n}$, denoting the utilities of players 1 and 2 respectively at the pure strategy profiles given by the rows and columns of the matrices. For mixed strategies $p \in \Delta_{m-1}$ and $q \in \Delta_{n-1}$ for players 1 and 2 respectively

$$u_1(p, q) = p^\top Aq, u_2(p, q) = p^\top Bq.$$

Define the following quantities,

$$c_i(p, q) = \max\{A_iq - p^\top Aq, 0\}, \text{ where } A_i \text{ is the } i^{\text{th}} \text{ row of } A, i \in \{1, \dots, m\}.$$

$$d_j(p, q) = \max\{p^\top B_j - p^\top Bq, 0\}, \text{ where } B_j \text{ is the } j^{\text{th}} \text{ column of } B, j \in \{1, \dots, n\}.$$

Clearly, both quantities are non-negative for all i, j .

Now, we define two functions P and Q as follows

$$P_i(p, q) = \frac{p_i + c_i(p, q)}{1 + \sum_{k=1}^m c_k(p, q)}, i \in \{1, \dots, m\}; \quad Q_j(p, q) = \frac{q_j + d_j(p, q)}{1 + \sum_{k=1}^n d_k(p, q)}, j \in \{1, \dots, n\}.$$

Clearly, $P_i(p, q) \geq 0, \forall i$ and $\sum_{i=1}^m P_i(p, q) = 1$. Hence $P(p, q) \in \Delta_{m-1}$ and similarly we see that $Q(p, q) \in \Delta_{n-1}$. Define the transformation function

$$T(p, q) = (P(p, q), Q(p, q)).$$

We see that, $T : \Delta_{m-1} \times \Delta_{n-1} \mapsto \Delta_{m-1} \times \Delta_{n-1}$, and maps a convex and compact set onto itself. From the definitions it is clear that c_i and d_j 's are continuous in (p, q) , therefore, P_i 's and Q_j 's are also continuous which implies that T is continuous. Hence, by Brouwer's fixed point theorem,

$$\exists (p^*, q^*) \text{ s.t. } T(p^*, q^*) = (p^*, q^*).$$

Claim 6.3

$$\sum_{k=1}^m c_k(p^*, q^*) = 0; \quad \sum_{k=1}^n d_k(p^*, q^*) = 0.$$

Proof:[of Claim] Suppose the claim is false, i.e., $\sum_{k=1}^m c_k(p^*, q^*) > 0$. Since (p^*, q^*) is a fixed point of T

$$p_i^* = \frac{p_i^* + c_i(p^*, q^*)}{1 + \sum_{k=1}^m c_k(p^*, q^*)} \Rightarrow p_i^* \left(\sum_{k=1}^m c_k(p^*, q^*) \right) = c_i(p^*, q^*). \quad (6.1)$$

Define a subset of indices as $I = \{i : p_i^* > 0\}$. We see that

$$I = \{i : p_i^* > 0\} = \{i : c_i(p^*, q^*) > 0\} = \{i : A_iq^* > p^{*\top} Aq^*\}. \quad (6.2)$$

The first equality follows from eq. (6.1) and our assumption that $\sum_{k=1}^m c_k(p, q) > 0$. The second equality come from the definition of c_i . Define $u_i^* := p^{*\top} Aq^*$.

Now we see

$$u_1^* = \sum_{i=1}^m p_i^* A_iq^* = \sum_{i \in I} p_i^* (A_iq^*) > \left(\sum_{i \in I} p_i^* \right) u_1^* = u_1^*.$$

The first equality is by definition, the second inequality holds since p_i^* is positive only on I (by definition), the inequality holds from eq. (6.2), and the last equality holds since u_i^* is a scalar and comes out of the summation. The inequality above is a contradiction. Similarly we can prove the claim for $\sum_k d_k$. Hence our claim is proved. ■

By this claim, $\sum_{k=1}^m c_k(p^*, q^*) = 0$. Since $c_k(p^*, q^*) \geq 0, \forall k = 1, \dots, m$, it implies that $c_k(p^*, q^*) = 0 \forall k = 1, \dots, m$. By definition of c_i 's, we then have

$$\begin{aligned} A_i q^* &\leq p^{*\top} A q^* \\ \Rightarrow \sum_{i=1}^m p'_i A_i q^* &\leq p^{*\top} A q^*. \end{aligned}$$

The implication holds for any arbitrary mixed strategy p' of player 1. Similarly we can show that q^* is a best response for player 2 against the mixed strategy p^* played by player 1. Therefore (p^*, q^*) is a MSNE. ■

Lecture 7: August 16, 2017

Lecturer: Swaprava Nath

Scribe(s): Prakhar Ji Gupta

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

7.1 Correlated Equilibrium

Motivation In the previous lectures, we have seen games with different types of equilibria and finally arrived at the mixed strategy Nash equilibrium (MSNE) which was the weakest and most general. The most special property of MSNE is that it always exists for any finite game and can be found by solving a finite number of (potentially non-linear) equations. However, calculating an MSNE is computationally difficult. In this lecture, we look at another equilibrium notion called *correlated equilibrium* (CE) which is weaker than MSNE.

In a Nash equilibrium, each player chooses his strategy independent of the other player, which may not always lead to the best outcome. However if the players trust a third-party agent, who randomizes over the strategy profiles and suggests the individual strategies to the concerned players, the outcomes can be significantly better. Such a strategy is called correlated strategy. Note that a correlated strategy is *not* a strategy of the players, rather it is a strategy of the third-party agent.

Example 7.1.1 *Consider a busy crossing of two roads. If the traffic of both roads move at the same time, it is a chaos, leading to potential accidents. If both stops, it is useless. The only good outcome is when traffic on one road stops and the other moves. The traffic light accomplishes this objective by periodically asking one road to stop and the other road to move. The traffic light serves the purpose of the (automated) third party agent and the players are the traffic of each road.*

7.2 Definition and Examples

Definition 7.1 *A correlated strategy is a mapping $\pi : S \mapsto [0, 1]$ such that $\sum_{s \in S} \pi(s) = 1$ where $S = S_1 \times S_2 \times \dots \times S_n$ and S_i represents the strategy set of player i .*

Hence, a correlated strategy π is a joint probability distribution over the strategy profiles.

A correlated strategy is a correlated equilibrium if it becomes self-enforcing, i.e., no player ‘gains’ by deviating from the suggested strategy.

Note: Here the suggested strategy π is a common knowledge.

Definition 7.2 *A correlated equilibrium (CE) is a correlated strategy π such that $\forall s_i \in S_i$ and $\forall i \in N$,*

$$\sum_{s_{-i} \in S_{-i}} \pi(s_i, s_{-i}) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} \pi(s_i, s_{-i}) u_i(s'_i, s_{-i}) \quad \forall s'_i \in S_i. \quad (7.1)$$

This means that the player i does not gain any advantage in the expected utility if he deviates from the suggested correlated strategy π , assuming all other players follow the suggested strategy.

Another way to interpret a correlated equilibrium is that π is a single randomization device (a dice e.g.) which gives a random outcome which is a strategy profile, and a specific player only observes the strategy corresponding to her. Given that observation, she computes her expected utility, and if that does not improve if she picks another strategy (and it happens for every player) that randomization device is a correlated equilibrium.

The following examples help us understand the definitions better.

7.2.1 Game Selection Problem

In the problem, two friends want to go to watch a game together, however Player 1 like Cricket more and Player 2 likes Football. The Utility Function is represented in the form:

1 \ 2	C	F
C	2,1	0,0
F	0,0	1,2

In MSNE, we saw that the expected utility of each player was $\frac{2}{3}$. However if the correlated strategy is such that $\pi(C, C) = \frac{1}{2} = \pi(F, F)$. If we assume that Player 1 is suggested to choose F , then the expected utility from following the suggestion is given as

$$\sum_{s_{-1} \in S_{-1}} \pi(s_{-1}|F) u_1(F, s_{-1}) = \frac{1}{\pi(F)} [\pi(F, C) u_1(F, C) + \pi(F, F) u_1(F, F)] = \frac{1}{\frac{1}{2}} \left[0 + \frac{1}{2} 1 \right] = 1 \quad (7.2)$$

where $\pi(F)$ is the probability that F is suggested to Player 1, $\pi(s_{-1}|F)$ is the probability that s_{-1} is strategy of other players when 1 is suggested F and $\pi(F, F)$ is probability that (F, F) is the strategy profile. If Player 1 deviates from the strategy, then his expected utility is

$$\sum_{s_{-1} \in S_{-1}} \pi(s_{-1}|F) u_1(C, s_{-1}) = \frac{1}{\pi(F)} [\pi(C, F) u_1(C, C) + \pi(F, F) u_1(C, F)] = \frac{1}{\frac{1}{2}} [0 + 0] = 0. \quad (7.3)$$

Similarly, if C is suggested to Player 1, his expected utility is 2 when he follows the suggestion and 0 when he does not follow. Similar conclusions hold when we consider Player 2. This proves that the correlated strategy here is a correlated equilibrium.

Note that utility at the expected equilibrium is $\frac{1}{2}(1 + 2) = \frac{3}{2}$ as compared to $\frac{2}{3}$ in MSNE.

7.2.2 Traffic Accident Problem

In the problem, two cars are at a crossroad and wish to cross it. Their utilities are positive if they cross the road, by they cross together, they will collide. The utilities is represented in the form:

1 \ 2	Stop	Go
Stop	0,0	1,2
Go	2,1	-10,-10

Consider a correlated strategy π such that $\pi(S, G) = \pi(S, S) = \pi(G, S) = \frac{1}{3}$. If we assume that Player 1 is suggested to choose Stop, then the expected utility from following the suggestion is given as

$$\sum_{s_{-1} \in S_{-1}} \pi(s_{-1}|S) u_1(S, s_{-1}) = \frac{1}{\pi(S)} [\pi(S, S) u_1(S, S) + \pi(S, G) u_1(S, G)] = \frac{1}{\frac{2}{3}} \left[0 + \frac{1}{3} 1 \right] = \frac{1}{2}. \quad (7.4)$$

where $\pi(S)$ is the probability that S is suggested to Player 1, $\pi(s_{-1}|S)$ is the probability that s_{-1} is strategy of other players when 1 is suggested S and $\pi(S, S)$ is probability that (S, S) is the suggested strategy. If Player 1 deviates from the strategy, then his expected utility is

$$\sum_{s_{-1} \in S_{-1}} \pi(s_{-1}|S) u_1(G, s_{-1}) = \frac{1}{\pi(S)} [\pi(S, S) u_1(G, S) + \pi(S, G) u_1(G, G)] = \frac{1}{\frac{2}{3}} \left[\frac{1}{3} 2 + \frac{1}{3} (-10) \right] = -4. \quad (7.5)$$

If G is suggested to Player 1, the expected utility is 2 on following the suggestion and 0 on deviating from the suggestion. Similarly, we can find that same conclusions hold when we consider Player 2. This proves that the correlated strategy here is a correlated equilibrium.

This game is known as Chicken Game as well.

Note: CE is generally not unique for any game and depends upon the randomization process. In our example itself, $\pi(S, G) = \pi(G, S) = \frac{1}{2}$ is also a CE with expected utility of $\frac{3}{2}$ for each player.

7.2.3 Interpretation

Another way to interpret a CE is that it is a distribution over the strategy profiles such that if a strategy which has a non-zero probability of occurrence for Player i is suggested to i , the player can compute the posterior distribution of the strategies suggested to other players. Player i 's expected payoff according to that distribution will be maximized by following the suggestion if other players follow their respective suggestions as well. More formally, let \bar{s}_i be the strategy suggested to Player i , then it is a CE if $\forall i \in N$:

$$\sum_{s_{-i} \in S_{-i}} \pi(s_{-i}|\bar{s}_i) u_i(\bar{s}_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} \pi(s_{-i}|\bar{s}_i) u_i(s'_i, s_{-i}), \quad \forall s'_i \in S_i \quad (7.6)$$

$$\Rightarrow \sum_{s_{-i} \in S_{-i}} \pi(\bar{s}_i, s_{-i}) u_i(\bar{s}_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} \pi(\bar{s}_i, s_{-i}) u_i(s'_i, s_{-i}), \quad \forall s'_i \in S_i. \quad (7.7)$$

7.2.4 Computing the Correlated Equilibrium:

To find a CE, we need to solve a set of linear equations with the variables as $\pi(s), s \in S$. By 7.1, we know $\pi(s)$ is a CE if $\forall s_i \in S_i$, and $\forall i \in N$

$$\sum_{s_{-i} \in S_{-i}} \pi(s) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} \pi(s) u_i(s'_i, s_{-i}) \quad \forall s'_i \in S_i. \quad (7.8)$$

The total number of inequalities here are $O(nm^2)$, assuming $|S_i| = m, \forall i \in N$. We also need to ensure that $\pi(s)$ is a valid probability distribution. Therefore

$$\pi(s) \geq 0, \quad \forall s \in S \quad m^n \text{ inequalities} \quad (7.9)$$

$$\sum_{s \in S} \pi(s) = 1 \quad (7.10)$$

The inequalities together represent a feasibility LP which is poly-time solvable. For computing MSNE, the number of support profiles are $O(2^{mn})$, which is exponentially larger than the number of inequalities to find a CE ($O(m^n)$). Therefore computing a CE is a much simpler problem than a MSNE. It can also be shown that every MSNE is a CE. The whole space of equilibrium guarantees can be represented as:




Figure 7.1: Venn diagram of the equilibria of a game

7.3 Extensive Form Games

We now discuss a different form of representing games, that is more appropriate for multistage games. These game representation is called *extensive form games* (EFG). We start with the perfect information EFGs.

7.3.1 Perfect Information Extensive Form Games

Normal form games (NFG) are appropriate when players take their actions simultaneously – their action profile decides the outcome. On the other hand, in **Extensive form games**, players take actions depending upon the sequence of actions, which we will call *history*, taken in the game, and the outcome happens at the end of the sequential actions. While NFGs are easily represented by payoff matrices, EFGs are best represented by a tree-like structure. In a **Perfect Information EFG**, every player knows the history till that time *perfectly*. The game below is an example of PIEFG.




Figure 7.2: Brother-sister chocolate sharing game in extensive form

7.3.2 Chocolate Division Game

Suppose a mother gives his elder son two (indivisible) chocolates to share between him and his younger sister. She also warns that if there is any dispute in the sharing, she will take the chocolates back and nobody will get anything. The brother can propose the following sharing options: (2-0): brother gets two, sister gets nothing, or (1-1): both gets one each, or (0-2): both chocolates to the sister. After the brother proposes the sharing, his sister may “Accept” the division or “Reject” it. This can be represented as shown in Figure 7.2.

7.3.3 Notation

We formally denote a PIEFG by the tuple

$$\langle N, A, \mathcal{H}, \mathcal{X}, P, (u_i)_{i \in N} \rangle$$

where,

N	set of players
A	set of all possible actions (of all players)
\mathcal{H}	set of all <i>sequences of actions</i> (histories) satisfying empty sequence $\emptyset \in \mathcal{H}$ if $h \in \mathcal{H}$, any initial continuous sub-sequence h' of h belongs to \mathcal{H} A history $h = (a^{(0)}, a^{(1)}, \dots, a^{(T-1)})$ is <i>terminal</i> if $\nexists a^{(T)} \in A$ s.t. $(a^{(0)}, a^{(1)}, \dots, a^{(T-1)}, a^{(T)}) \in \mathcal{H}$
$Z \subseteq \mathcal{H}$	set of all <i>terminal</i> histories
$\mathcal{X} : H \setminus Z \mapsto 2^A$	action set selection function
$P : H \setminus Z \mapsto N$	player function
$u_i : Z \mapsto \mathbb{R}$	utility function of player i

The *strategy* of a player in an EFG is a sequence of actions at every history where the player plays. Formally

$$S_i = \prod_{\{h \in H : P(h)=i\}} \mathcal{X}(h).$$

In other words, it is a *complete contingency plan* of the player. It enumerates potential actions a player can take at every node where he can play, even though some sequence of actions may never be executed together.

7.3.4 Representing the Chocolate Division Game

With the notation above, we represent the game as follows.

$$N = \{B, S\}, A = \{2-0, 1-1, 0-2, A, R\}$$

$$\mathcal{H} = \{\emptyset, (2-0), (1-1), (0-2), (2-0, A), (2-0, R), (1-1, A), (1-1, R), (0-2, A), (0-2, R)\}$$

$$Z = \{(2-0, A), (2-0, R), (1-1, A), (1-1, R), (0-2, A), (0-2, R)\}$$

$$\mathcal{X}(\emptyset) = \{(2-0), (1-1), (0-2)\}$$

$$\mathcal{X}(2-0) = \mathcal{X}(1-1) = \mathcal{X}(0-2) = \{A, R\}$$

$$P(\emptyset) = B, P(2-0) = P(1-1) = P(0-2) = S$$

$$u_B(2-0, A) = 2, u_B(1-1, A) = 1, u_S(1-1, A) = 1, u_S(0-2, A) = 2$$

$$u_B(0-2, A) = u_B(0-2, R) = u_B(1-1, R) = u_B(2-0, R) = 0$$

$$u_S(0-2, R) = u_S(1-1, R) = u_S(2-0, R) = u_S(2-0, A) = 0$$

$$S_1 = \{2-0, 1-1, 0-2\}$$

$$S_2 = \{Y, N\} \times \{Y, N\} \times \{Y, N\} = \{YYY, YYN, YNY, YNN, NYY, NYN, NNY, NNN\}$$

7.3.5 Representing PIEFG as NFG

Given S_1 and S_2 , we can represent the game as an NFG, which can be written in the form of matrix. This can be generalized for all PIEFG, i.e., each PIEFG can be represented as a NFG. For the given example, we can express the utility function as in the following table:

B \ S	YYY	YYN	YNY	YNN	NYY	NYN	NNY	NNN
2-0	(2,0)	(2,0)	(2,0)	(2,0)	(0,0)	(0,0)	(0,0)	(0,0)
1-1	(1,1)	(1,1)	(0,0)	(0,0)	(1,1)	(1,1)	(0,0)	(0,0)
0-2	(0,2)	(0,0)	(0,2)	(0,0)	(0,2)	(0,0)	(0,2)	(0,0)

Observe that there are many PSNEs in the given game, some of which leads to quite nonintuitive solutions. The PSNEs are marked in **Bold**.

B \ S	YYY	YYN	YNY	YNN	NYY	NYN	NNY	NNN
2-0	(2,0)	(2,0)	(2,0)	(2,0)	(0,0)	(0,0)	(0,0)	(0,0)
1-1	(1,1)	(1,1)	(0,0)	(0,0)	(1,1)	(1,1)	(0,0)	(0,0)
0-2	(0,2)	(0,0)	(0,2)	(0,0)	(0,2)	(0,0)	(0,2)	(0,0)

Some of the results like $\{2-0, NNY\}$, $\{2-0, NNN\}$ and $\{0-2\}$ are not practically useful. As the representation is very clumsy and does not provide us with any advantage, the NFG representation is wasteful and the EFG representation is succinct for such cases. The example also forces us to look for some other equilibrium ideas as the Nash Equilibrium does not serve much purpose in this case.

Lecture 8: August 18, 2017

Lecturer: Swaprava Nath

Scribe(s): Neeraj Yadav

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

8.1 Recap

Extensive form games are used to represent the sequential games. We will only consider perfect information extensive form games (PIEFG) with a *finite* number of stages. It can always be transformed into normal form game and find the Nash equilibrium. However, as we have seen, this representation is often redundant and the NE prediction may lead to non-credible threats. Therefore, we need a new equilibrium notion, which is a *refinement* of Nash equilibrium. A refinement of a Nash equilibrium refers to an equilibrium which is NE but the converse may not hold. We have the following intuitive result, which will be obvious when we define the refined equilibrium.

Theorem 8.1 *Every finite PIEFG has a PSNE.*

Here, by finite PIEFG, we mean that the number of players, actions, and stages of the PIEFG are finite.

Intuition: At every stage of the game, a player has *perfect information* about the action taken by the former player. Hence there is no reason for *randomizing* over the actions. There is always a pure action that is *weakly* superior than the other actions. This will be obvious after the discussion on subgame perfection.

Examples of PIEFG: Chess, Tic-Tac-Toe, Bargaining.

Zermelo (1913) showed using an argument similar to PIEFG that if both the players in chess are infinitely rational and intelligent, chess must be a very boring game.

In the given perfect information extensive normal game $((BH), (CE))$ is indeed a PSNE but it is not very intuitive. Player 1 at his second choice node plays H instead of G whereas he will get more utility had he played G. This suspicious behavior of Player 1 is termed as a *threat* to Player 2 forcing him to play E (as the unilateral deviation of Player 2 to F will yield an utility of 0 to him). Hence it continues to be a Nash equilibrium. But this threat of Player 1 is really not *credible*. If Player 2 ever plays F instead of E and the game reaches the stage where Player 1 has options between G and H, will Player 1 still choose H over G (at his own loss)?

Therefore, the Nash equilibrium $((BH), (CE))$ seems ambiguous in the context of PIEFGs. Hence, there is a need to refine the concept of equilibrium through defining the notion of a subgame to avoid such equilibrium having non-credible threats.

Definition 8.2 (Subgame) *The extensive form game represented by the subtree at a node is called the subgame at that node.*




Fig:8.1 PIEFG Representation of a Sequential Game

Figure 8.1: PIEFG representation of a sequential game

Definition 8.3 (Subgame Perfect Nash Equilibrium) *The subgame perfect Nash equilibrium (SPNE) of a game G are all strategy profiles $s \in S := \times_{i \in N} S_i$ such that for any subgame G' of G the restriction of s to G' is a Nash equilibrium of G' .*

We define the restriction of a strategy profile s to a subgame G (denoted by $s|G'$) as the truncation of the actions of s to the actions that are relevant in G' . For example, in Figure 8.1, let $s = ((BH), (CE))$ and G' be the subtree rooted at the second decision node of Player 1 (at history BF), then $s|G' = (H, \emptyset)$. Clearly, $s|G'$ is not a NE at G' , hence $((BH), (CE))$ is not an SPNE.

8.2 Computing SPNE: Backward Induction Algorithm

1. Start at the leaf having maximum depth.
2. For the player in the parent node find the action which maximizes the utility for that player.
3. Retain that (player,action) and delete all the edges at that level, translate the utilities to the parent node.
4. Go up one level and repeat (1).
5. Stop if root is reached.

The algorithm is implemented as a single depth first traversal of the game tree. It identifies the strategy which has higher pay-off for the player in the bottom-most subgame tree and removes the other counter actions and updates the pay-off of the parent node and repeats the complete process again. The algorithm terminates when the action with higher pay-off among all the possible actions for the player at root is found out. In this manner, the algorithm gives SPNE.

The algorithm yields much better prediction than the Nash equilibria since the treat in SPNE is credible. Also the algorithm always gives an SPNE. Therefore, SPNE is guaranteed to exist in a PIEFG. Though, there could be multiple SPNEs.

However, the complexity of BI algorithm sometimes prohibitive – since it needs to parse all possible paths in the game tree. For games like chess, the extensive form representation has $\sim 10^{150}$ nodes and it is not

feasible to apply backward induction. Game softwares uses heuristic pruning for computer players and does not consider the parts of the game tree which can never be the candidates for equilibrium.

Exercise: check if $((AG), (CF))$ is an SPNE of the previous game.

Note: SPNE is always a PSNE but the reverse is not true.

8.3 Limitations of SPNE

Centipede Game: In this game two players makes alternate decisions, at each turn choosing between going “down” and ending the game or going “across” and continuing it except at the last node where going “across” also ends the game. The payoffs are constructed in such a way that the player achieves higher payoffs by choosing “down”. Consider the last choice, at that point the best choice for the player is to go down. Going down is also the best choice for the other player in the previous choice point. By induction the same argument holds for all choice points.




Fig:8.2 Centipede Game in PIEFG Representation

Figure 8.2: Centipede game

When this game was experimented on people it was found that people played “across” moves also until close to end of the game. This is a practical criticism of the concept of SPNE.

On the theory side, suppose that you are the second player in the game, and in the first step of the game the first player plays across. The SPNE tells that you should go down, but the same analysis suggests that you would not have reached this choice point in the first place. Hence the idea of SPNE is self-contradictory at certain decision nodes, even though it guarantees NE at every subgame.

8.4 Imperfect Information Extensive Form Games

The PIEFG is not able to represent the simultaneous move games like neighboring kingdom’s dilemma. Hence there is a need to move for a more general representation.

Definition 8.4 (Imperfect Information Extensive Form Game) An imperfect-information extensive form game is a tuple

$$\langle N, A, \mathcal{H}, \mathcal{X}, P, (u_i)_{i \in N}, (I_i)_{i \in N} \rangle$$

Where $\langle N, A, \mathcal{H}, \mathcal{X}, P, (u_i)_{i \in N} \rangle$ is a PIEFG and for every $i \in N$, $I_i := (I_i^1, I_i^2, \dots, I_i^{k(i)})$ is a partition of $\{h \in \mathcal{H} \setminus Z : P(h) = i\}$ with the property that if $h, h' \in I_i^j$, then $\mathcal{X}(h) = \mathcal{X}(h')$. The sets in the partition I_i are called information sets of player i , and in a specific information set, the actions available to player i are same.

In the game shown in Figure 8.3, both players 1 and 2 have one information set each, and $I_1^1 = \{\emptyset\}$, $I_2^1 = \{(A), (D)\}$.

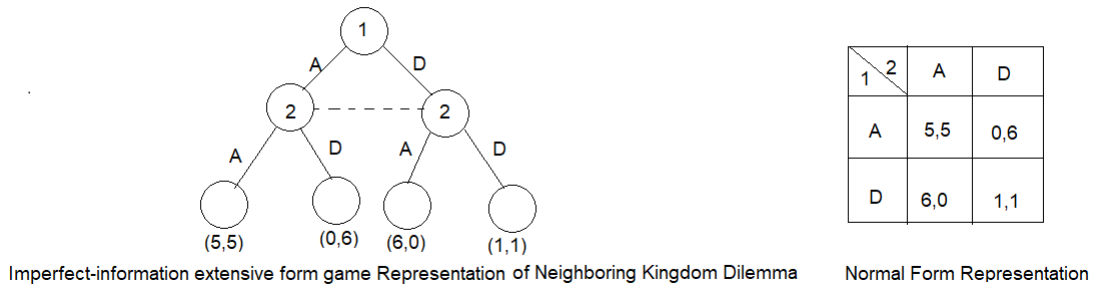


Figure 8.3: An example EFG with imperfect information

An information set is always non-empty but it can be a singleton as PIEFG is also a IIEFG with singleton information sets. IIEFG is a richer representation than PIEFG as an NFG can also be converted into IIEFG.

Lecture 9: August 22, 2017

Lecturer: Swaprava Nath

Scribe(s): Pranav Sao

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

9.1 Imperfect-information extensive-form games

A PIEFG is not able to represent the simultaneous move games like neighboring kingdoms dilemma. Hence we needed to move to a more general representation, namely imperfect information EFG.

Definition 9.1 (Imperfect Information Extensive Form Game) An imperfect-information extensive form game is a tuple

$$\langle N, A, \mathcal{H}, \mathcal{X}, P, (u_i)_{i \in N}, (I_i)_{i \in N} \rangle$$

Where $\langle N, A, \mathcal{H}, \mathcal{X}, P, (u_i)_{i \in N} \rangle$ is a PIEFG and for every $i \in N$, $I_i := (I_i^1, I_i^2, \dots, I_i^{k(i)})$ is a partition of $\{h \in \mathcal{H} \setminus Z : P(h) = i\}$ with the property that if $h, h' \in I_i^j$, then $\mathcal{X}(h) = \mathcal{X}(h')$. The sets in the partition I_i are called information sets of player i , and in a specific information set, the actions available to player i are same.

Set I_i for every player i , is a collection of Information sets $I_i^j, j = 1, \dots, k(i)$. Information sets are collection of histories where the player at that history is uncertain about which history has been reached.

Since by Definition 9.1, the actions at an information set are identical, we can define \mathcal{X} over Information sets I_i^j s, rather than defining them over histories h, h' . Therefore

$$\mathcal{X}(h) = \mathcal{X}(h') = \mathcal{X}(I_i^j).$$

Definition 9.2 (Strategy Set) Strategy set of player $i, i \in N$ is defined as the Cartesian product of the actions available to player i at his information sets, i.e.,

$$S_i = \prod_{\tilde{I} \in I_i} \mathcal{X}(\tilde{I}).$$

The intuition of the information set comes from the fact that a player can be uncertain about the true state of a game when he is deciding his action. Consider the game of neighboring kingdoms dilemma (Fig. 9.1). In IIEFG representation, Player 2 does not know at which node (a or b), he presently is. Thus his action here is independent of it.

9.1.1 Representational equivalence

To draw a connection between different representations of games, we need the equilibrium concepts to follow in newer representations too. However, in IIEFG there could be strategies which are different than a mixed strategy.




Figure 9.1: Neighboring Kingdom's Dilemma




Figure 9.2: IIEFG representation of a game with a forgetful player

Definition 9.3 (Behavioral Strategy) *In a IIEFG, a behavioral strategy assigns a probability distribution over the set of possible actions at each information set.*

This is quite different from the mixed strategies. In a mixed strategy, a player assigns probabilities over the pure strategies (which is a complete contingency plan). But in behavioral strategy, the player independently randomizes over his actions at his every information set.

At this point, it may appear that a transformation between these two strategies is always possible. But we show that these two notions are actually incomparable.

Example 9.1.1 *Consider the IIEFG shown in Figure 9.2. The discontinuous line between the nodes where player 1 plays shows that these two nodes are in the same information set of player 1. Thus his strategy set is $\{L, R\}$. Player 2 has two strategies U and D .*

Mixed strategy *From the NFG representation of this game, we can conclude that for player 1 R strictly dominates L , and for player 2 D weakly dominates U . Thus game have **weakly dominant strategy equilibrium** (R, D) , which is also the unique MSNE given by $((0, 1), (0, 1))$.*

Behavioral strategy In behavioral strategy, players randomize over actions at each information set. From IIEFG representation we can see that player 2's rationality must dictate him to play D, as it is always gives him more payoff. So, with the probability distribution for player 2 being (0 : U, 1 : D), let the probability distribution for player 1 be ($p : L, (1 - p) : R$). The expected payoff for player 1 is

$$p^2 \times 1 + p(1 - p) \times 100 + (1 - p) \times 2,$$

which is maximum at $p = \frac{98}{198}$. Hence, a behavioral strategy equilibrium for this game is $((\frac{98}{198}, \frac{100}{198}), (0, 1))$.

MSNE and BSE are different in this game because player 1 forgets what he played at level 1. This kind of games is known as IIEFGs with imperfect recall. Next, we consider games with perfect recall, where these two strategies are equivalent.

9.2 Games with Perfect Recall

Let us consider type of games where at every opportunity of a player to act, he remembers exactly what he did at his every previous turn to play. Such games are called *games with perfect recall*.

Definition 9.4 Player i has a perfect recall in an IIEFG, if for any two histories $h, h' \in I_i^j$ where $h = (v_0, a_0, v_1, a_1, \dots, v_{m-1}, a_{m-1}, v_m)$ and $h' = (v'_0, a'_0, v'_1, a'_1, \dots, v'_{n-1}, a'_{n-1}, v'_n)$ where v_i, v'_i s are the nodes of the IIEFG and a_i, a'_i s are the actions at corresponding nodes, the following holds.

1. $m = n$
2. for all j s.t. $0 \leq j \leq (m - 1)$, v_j, v'_j must be in same information set of player i .
3. for all j s.t. $0 \leq j \leq (m - 1)$, **if** $P(h_j) = i$ **then** $a_j = a'_j$, where h_j is the truncated history h at level j from the root.

A game has **perfect recall** if every player has perfect recall.

References

- [CW87] M. JACKSON, K. LEYTON- BROWN and Y. SHOHAM, Game Theory Lecture 4-8 "Imperfect Information Extensive Form: Definition, Strategies"

Lecture 10: August 23, 2017

Lecturer: Swaprava Nath

Scribe(s): Asim Unmesh


Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

10.1 Outcome Equivalence of Behavioral and Mixed Strategy

In the context of IIEFGs *behavioral strategy* of a player is a probability distribution over the actions at every information set of the player. Thus the player is randomizing at each node rather than randomizing over the pure strategies of an IIEFG (the complete contingency plan). The latter is a mixed strategy of the player.

Definition 10.1 (Outcome Equivalence) *A behavioral strategy b_i and a mixed strategy σ_i are outcome equivalent if for all σ_{-i} , the probability distribution induced over the terminal vertices are the same for (b_i, σ_{-i}) and (σ_i, σ_{-i}) .*

Figure 10.1: Example of Outcome Equivalence



In the above game, consider the following mixed strategy of player 2:

$$\sigma_2(Ll) = \sigma_2(Lr) = \frac{1}{3}, \sigma_2(Rl) = \frac{1}{12}, \sigma_2(Rr) = \frac{1}{4}.$$

Also, consider the following behavioral strategy of player 2:

$$b_2(L) = \frac{2}{3}, b_2(R) = \frac{1}{3}, b_2(l) = \frac{1}{2} = b_2(r).$$

Suppose player 1 plays U with probability P_U and D with probability P_D . Now for the probability of reaching terminal history LU, according to the mixed strategy is equal to the probability of reaching the same terminal history w.r.t. the behavioral strategy, i.e.,

$$(\sigma_2(Ll) + \sigma_2(Lr))P_U = \frac{2}{3}P_U = b_2(L) \cdot P_U.$$

Similar conclusions hold for the terminal histories LDl, LDr, and R.

For LDl: $\sigma_2(Ll)P_D = \frac{1}{3}P_D = b_2(L) \cdot P_D \cdot b_2(l)$

For LDr: $\sigma_2(Lr)P_D = \frac{1}{3}P_D = b_2(L) \cdot P_D \cdot b_2(r)$

For R: $(\sigma_2(Rl) + \sigma_2(Rr))P_U = \frac{1}{3} \cdot P_U = b_2(R) \cdot P_U$

Thus we see that all the terminal nodes have same probability of occurrence under the mixed strategies and behavioral strategies. Thus, the abovementioned mixed strategy and behavioral strategy for player 2 are outcome equivalent. The following result shows that this equivalence always holds for games of perfect recall.

Theorem 10.2 (Kuhn 1953) *In IIEFGs with perfect recall every mixed strategy is outcome equivalent to behavioral strategies.*

The proof of the theorem is constructive and follows the principle used in the example above. From now on, we will consider games only with perfect recall, and all strategies will be referred to as behavioral strategies.

10.2 Equilibrium Notion and Importance of Belief




Figure 10.2: Football-Cricket Game for 2 players

In the example of Football - Cricket game (Figure 10.2)

- if player 2 believes that, $p_a > \frac{2}{3}$, then it is better for him to play F.
- if player 2 believes that, $p_a < \frac{2}{3}$, then it is better for him to play C.
- if player 2 believes that, $p_a = \frac{2}{3}$, then he can mix F and C in any way.

Thus, the players' belief is important in deciding his/her strategy. In the equilibrium notion of IIEFGs, this idea is made explicit.

10.2.1 Belief

Let the *information sets* of player i be $I_i = \{I_i^1, I_i^2, \dots, I_i^{k(i)}\}$. In an IIEFG, the belief of player i is a map $\mu_i^j : I_i^j \rightarrow [0, 1]$, such that,

$$\sum_{x \in I_i^j} \mu_i^j(x) = 1.$$

10.2.2 Bayesian Belief of Player i

A belief $\mu_i := (\mu_i^j, j = 1, \dots, k(i))$ of player i is Bayesian with respect to the behavioral strategy σ , if it is derived from the strategy profile σ using Bayes rule, i.e.,

$$\mu_i^j(x) = \frac{P_\sigma(x)}{\sum_{y \in I_i^j} P_\sigma(y)}, \forall x \in I_i^j, \forall j = 1, \dots, k(i).$$

10.3 Sequential Rationality

A strategy σ_i of player i at an information set I_i^j is sequentially rational given σ_{-i} and beliefs μ_i if $\forall \sigma'_i$

$$\sum_{x \in I_i^j} \mu_i(x) u_i(\sigma_i, \sigma_{-i} | x) \geq \sum_{x \in I_i^j} \mu_i(x) u_i(\sigma'_i, \sigma_{-i} | x).$$

10.4 Football Cricket Game example for Sequential Rationality

In figure 10.2 we see that the game is played between two players 1 and 2. Players 1 and 2 picks action F with probabilities p and q respectively. Consider a belief (we suppress the superscript since each player has only one information set)

$$\begin{aligned} \mu_2(a) &= p = 0.5 \\ \mu_2(b) &= 1 - p = 0.5 \end{aligned}$$

Then, player 2 can compute,

$$\sum_{x \in I_2^j} \mu_2(x) U_2(\sigma_1, \sigma_2 | x) = 0.5[q \cdot 1 + (1 - q) \cdot 0] + 0.5[2 \times (1 - q)] = 0.5[2 - q]$$

Thus, to maximize his utility given his belief about the moves of player 1, it will be sequentially rational for player 2 to pick $q = 0$, i.e., play football with zero probability.

So, $\sigma = ((0.5, 0.5), (0, 1))$ is sequentially rational for player 2. But is it sequentially rational for player 1?

10.5 Perfect Bayesian Equilibrium

An assessment (σ, μ) is a *perfect Bayesian equilibrium* (PBE) if for every player i

1. μ_i is Bayesian with respect to σ ,
2. σ_i is sequentially rational given σ_{-i} and μ_i at every information set of i .

From the last example, we can find out that $\sigma = ((0.5, 0.5), (0, 1))$ is not sequentially rational for player 1. Hence it is not a PBE. However, $((2/3, 1/3), (1/3, 2/3))$ is.

We provide a result that shows that PBE is a *refinement* of MSNE.

Theorem 10.3 *Every Perfect Bayesian Equilibrium (PBE) is a Mixed Strategy Nash Equilibrium (MSNE).*

Proof: Exercise. ■

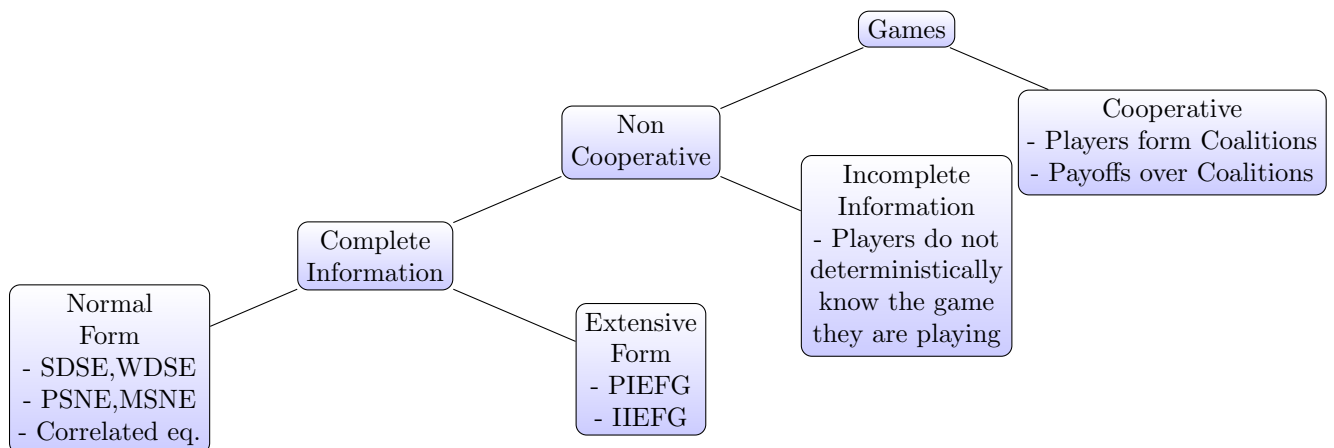
Lecture 11: August 25, 2017

Lecturer: Swaprava Nath

Scribe(s): Ameya Loya

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

11.1 Game Theory Ecosystem



Things not covered: Repeated Games, Stochastic Games,

11.2 Game Theory in Practice: Peer to Peer File Sharing

Conventional server-client models use single server to which every client is connected. However, this is inefficient since the download speed is bottlenecked by the server's bandwidth – while there is enough download/upload bandwidth with the clients that is unused. Peer-to-Peer (denoted P2P) is a technology to mitigate this inefficiency. It is a completely decentralized network in which each client is capable of acting as a server for another, and helps keep the download speed unaffected by the number of users (Figure 11.1). The advantages of P2P is that it is (a) highly scalable, and (b) resilient to server failures.

11.2.1 Basic Terminology

1. **Protocol:** Messages that can be sent, actions that can be taken over the network.
2. **Client:** A particular process for sending messages, taking actions.
3. **Reference Client:** A particular implementation of P2P.

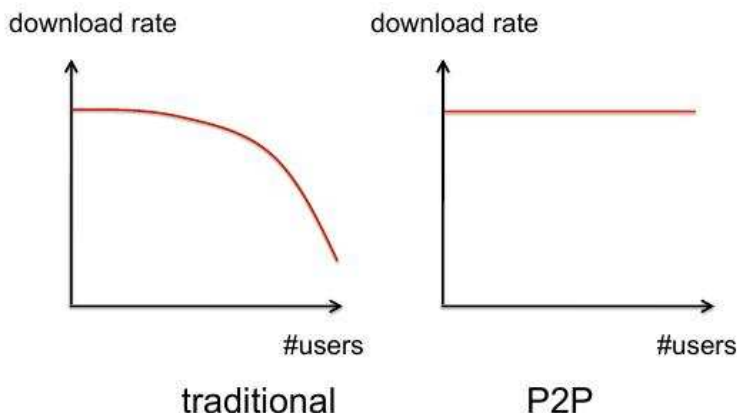


Figure 11.1: P2P innovation

11.3 Implementations of P2P Protocols

Early implementations of P2P technology was Napster and Gnutella.

Napster (1999 - 2001)

1. Centralized database
2. Users download music from each other

Gnutella (2000 -)

1. Get list of IP addresses of peers from set of known peers (no server)
2. To get a file: Query message broadcast by peer A to known peers
3. Query response: sent by B if B has the desired file (routed back to requestor)
4. A can then download directly from B

In both these protocols, the downloader’s strategies were not taken into account. If we consider that data upload is costly and every agent gets a positive utility if he downloads some file that is useful to him, then the upload-download dynamics can be represented as a normal form game – where every agent can either upload and download (Share) or never upload but only download (Free Ride). The setup is represented in Table 11.1. Clearly, the incentives with these two naïve protocols lead to scenario where free-riding is

		Person 2	
		Share	Free Ride
Person 1	Share	2, 2	-1, 3
	Free Ride	3, -1	0, 0

Table 11.1: A P2P file sharing game

a dominant strategy for both players (attested by the study of Adar and Huberman (2000), Figure 11.2). In Gnutella, 85% peers were observed free-riding by 2005; it had less than 1% of worldwide P2P traffic by 2013. Few other P2P systems met similar fate. Therefore it is unlikely that peers will share files with such protocols – hence a client with incentivizing properties was needed.

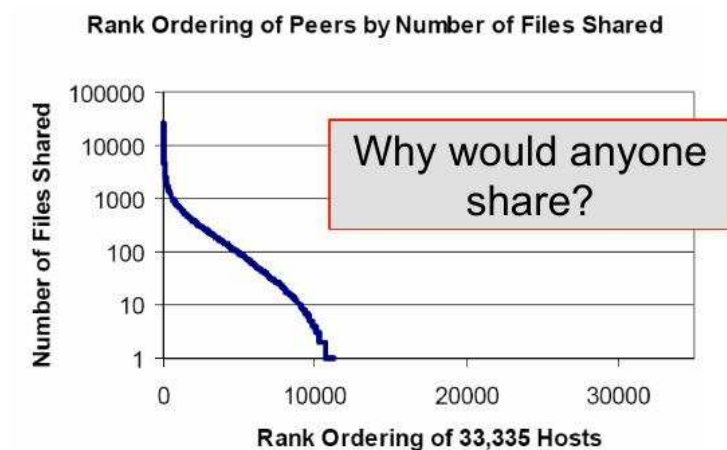


Figure 11.2: Sharing statistics. Image courtesy: Adar and Huberman (2000).

BitTorrent Protocol (2001 -)

BitTorrent was the first protocol that broke this deadlock with incentives for uploading. It breaks the file into multiple pieces and different pieces are delivered to different agents. Among the agents the exchange of the pieces follow a repeated game and the strategy that is followed is an adaptation of the tit-for-tat policy for repeated prisoner's dilemma game. The principle was:

“If you let me download, Ill reciprocate.”

The schematic of BitTorrent is shown in Figure 11.3.




Figure 5.4.: Starting a download process in the BitTorrent protocol: 1) A user goes to a searchable directory to find a link to a .torrent file corresponding to the desired content; 2) the .torrent file contains metadata about the content, in particular the URL of a tracker; 3) the tracker provides a list of peers participating in the swarm for the content (i.e., their IP address and port); 4) the user's BitTorrent client can now contact all these peers and download content.

Figure 11.3: BitTorrent details. Image courtesy: Parkes and Seuken (2017).

BitTorrent Optimistic Unchoking Algorithm Tracker is a centralized entity that controls the traffic, tracks the connection between peers and their speed of upload, download etc.

Reference Client Protocol:

- Set a threshold r of uploading speed (typically the third maximum speed in the recent past).
- If a peer j uploaded to i at a rate $\geq r$, unchoke j in the next period.
- If a peer j uploaded to i at a rate $< r$, choke j in the next period.
- Every three time periods, optimistically unchoke a random peer from the neighborhood who is currently choked, and leave that peer unchoked for three time periods.

The protocol is forcing file sharing to be a repeated game by fragmenting the files, which is a repeated Prisoners' Dilemma. Strategy of the seeder (the reference client – the one that is the default BitTorrent client) is tit-for-tat (TfT). A TfT strategy for a prisoner's dilemma starts with cooperate action, and if the opponent defects at any stage, it keeps of defecting from the next stage until the opponent plays cooperate again.

Though BitTorrent is extremely popular (Approx 85% of P2P traffic in US is BitTorrent), it does have some strategic vulnerabilities.

Attacks on BitTorrent With the BitTorrent protocol, an adversary may consider the following actions to decide how this can be gamed.

- How often to contact tracker?
- Which pieces to reveal?
- How many upload slots, which peers to unchoke, at what speed?
- What data to allow others to download?
- Possible goals: minimize upload, maximize download speed.

11.3.1 BitThief

BitThief does not perform any chokes or unchokes of remote peers, and it never announces any pieces. In other words, a remote peer always assumes that it interacts with a newly arrived peer that has just started downloading. Thus, a BitThief client is able to download without ever uploading any time. However, if it stays in one neighborhood for long, the other peers figure out that it is never uploading and will choke it. Therefore, the client asks the tracker more frequently for new peers and grows the neighborhood rapidly.

This protocol seems to have an easy fix. One can modify the tracker to block such peers that asks for new peers more rapidly.

Ref: Locher et al., "Free Riding in BitTorrent is Cheap", HotNets 2006

11.3.2 Strategic Piece Revealer

This is another attack on BitTorrent that picks the piece of the file to share in order to maximize its importance in its neighborhood. Peer A is ‘interested’ in peer B if B has a piece that A does not. Strategic Piece Revealer (SPR) client strategically shares only those files that are the most common piece the reciprocating peer does not have and saves its own ‘rare’ pieces. The BitTorrent reference client always use “rarest-first” to request, i.e., it will announce the rarest piece first and share that if requested. Therefore, by sharing the most common piece an SPR client will protect its monopoly over the rare pieces and keep others interested. The improvement in terms of the reduction in download time is shown in Figure 11.4.




Figure 11.4: Strategic Piece Revealer.

Ref: Levin et al., “BitTorrent is an Auction: Analyzing and Improving BitTorrents Incentives”, SIGCOMM 2008

11.4 Summary

- P2P demonstrates importance of game-theory in computer systems.
- Early systems were easily manipulated.
- BitTorrent’s innovation was to break files into pieces, enabling TitForTat.
- Still some vulnerabilities, but generally BitTorrent is a very successful example of incentive-based protocol design.

Notes

This lecture has been adapted from the relevant lecture notes of CS186 Harvard (Instructor: David C. Parkes) and the content and images are from the unpublished book of Parkes and Seuken, “Economics and Computation”, (2017).

Lecture 12: August 29, 2017

Lecturer: Swaprava Nath

Scribe(s): Ankit Bhardwaj

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

12.1 Introduction

In the complete information games that we have seen so far, the only uncertainties were over the probabilistic choice of actions for different players (e.g. for mixed strategies or behavioral strategies). Even in that case, the payoffs for different paths of a game were known deterministically to every player (was actually a common knowledge). In incomplete information games, as we shall see, that is not the case.

Definition 12.1 (Incomplete Information Games) *An incomplete information game is a game where players do not deterministically know which game they are playing. All players may have private information or types.*

Consider the following example.

Example 12.2 (A Soccer Game) *Suppose that there is a match scheduled between two competing clubs. The actions available to both the teams is the kind of game they decide to play. Let their choices be either to play attacking / aggressive or to play defensive. Let both these actions be denoted by A and D respectively.*

*Now, based on external random factors like weather at the day of game, unexpected player injuries and such, each team may have a hidden agenda. Say, when the teams are in a favourable position, they aim to **Win (W)** while in case of unfavourable position they aim to settle for a **Draw (D)**.*

Note: *The agenda of each team is their private information and is unknown to the other team. It is also beyond the control of the players (in this case, the teams). This private information which is a random realization is known as the type of the player. The players can privately observe them but cannot choose them.*

Consider the following utilities for different type profiles.

WW			WD			DW			DD		
	A	D		A	D		A	D		A	D
A	1,1	2,0	A	2,1	3,0	A	1,2	1,1	A	0,0	1,0
D	0,2	0,0	D	1,1	1,0	D	0,3	0,1	D	0,1	-1,-1

In this example, there are 4 possible type profiles of the players: WW is the type profile when both teams intend to win and the payoffs from the actions chosen by the players are depicted in the first matrix above. Similarly, the games corresponding to the other type profiles are depicted in the other matrices. Thus, there are 4 different payoff matrices, one corresponding to each type profile. However, even though both players know all the 4 matrices, they do not know which matrix the game they are in. Thus, they do not deterministically know which game they are playing.

We assume that the players and actions available to every player remain the same in different type profiles of the game – the utility is the only thing that changes across type profiles. We consider a special class of the incomplete information games, known as the Bayesian game, where type profiles are chosen from a common prior distribution P .

12.2 Bayesian Games

Formally, a Bayesian game is represented by a 5-tuple.

$$\langle N, (A_i)_{i \in N}, (\Theta_i)_{i \in N}, P, (\Gamma_\theta)_{\theta \in \times_{i \in N} \Theta_i = \Theta} \rangle$$

Where,

N = $\{1, \dots, n\}$, set of players.

A_i : action set of player i .¹

Θ_i : set of types for player i ; e.g. Win / Draw in Example 12.2.

P : common prior distribution over $\Theta = \times_{i \in N} \Theta_i$ with the restriction that $P(\theta_i) > 0 \forall \theta_i \in \Theta_i, \forall i \in N$.²

Γ_θ := $\langle N, (A_i)_{i \in N}, (u_i(\theta))_{i \in N} \rangle$, a normal form game for the type profile θ .

u_i : $\Theta \times A \mapsto \mathbb{R}$, where $A = \times_{i \in N} A_i$, normal form game utility function for each type profile.

12.2.1 Stages in Bayesian Games

- $\theta = (\theta_i, \theta_{-i})$ is chosen according to P .
- Each player observes his / her own θ_i .
- They pick action $a_i \in A_i$.
- Player i 's payoff $u_i((a_i, a_{-i}); (\theta_i, \theta_{-i}))$ is realized.

12.2.2 Strategies and Utilities

Strategy in a Bayesian game is a plan to map a state/type into action.

Pure Strategy $s_i : \Theta_i \mapsto A_i$

Mixed Strategy $\sigma_i : \Theta_i \mapsto \Delta(A_i)$

Definition 12.3 (Ex-ante Utility) *Ex-ante utility of player i is the expected utility before she observes her own type, i.e.,*

$$U_i(\sigma_i, \sigma_{-i}) = \sum_{\theta \in \Theta} P(\theta) \cdot U_i((\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i})), \theta).$$

Where,

$$U_i((\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i})), \theta) = \sum_{(a_1, \dots, a_n) \in A} \left(\prod_{j \in N} \sigma_j(\theta_j, a_j) \right) u_i((a_i, a_{-i}); (\theta_i, \theta_{-i})). \quad (12.1)$$

In Bayesian games, we assume that once player i observes her type, she establishes a belief according to Bayes rule (hence called a Bayesian game) on P as follows.

$$P(\theta_{-i} | \theta_i) = \frac{P(\theta_i, \theta_{-i})}{\sum_{\tilde{\theta}_{-i} \in \theta_{-i}} P(\theta_i, \tilde{\theta}_{-i})}.$$

Note: This is where the positive marginal is crucial.

Definition 12.4 (Ex-interim Utility) *Ex-interim utility of player i is the expected utility after she observes her own type, i.e.,*

$$U_i(\sigma|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) \cdot U_i(\sigma(\theta), \theta).$$

Where $U_i(\sigma(\theta), \theta)$ is as defined in Equation 12.1.

Ex-interim utility is most commonly used for analysis of incomplete information games. The above two utilities are related as follows.

$$U_i(\sigma) = \sum_{\theta_i \in \Theta_i} P(\theta_i) U_i(\sigma|\theta_i).$$

Example 12.5 (Two player bargaining game) *Consider a game between two players where player 1 is the seller and player 2 is a buyer. Player 1's type is the minimum price at which he is willing to sell a commodity. Player 2's type is the maximum amount that the player is willing to pay. For simplicity, let the types of both players be integers. Hence,*

- $N = \{1, 2\}$.
- $\Theta_1 = \Theta_2 = \{1, 2, \dots, 100\}$.

Assume that both players bid a number in $\{1, 2, \dots, 100\}$. If the bid of the seller is less than or equal to the bid of the buyer, the sale happens. Else, there is no trade. Hence the action sets of the players are, $A_1 = A_2 = \{1, 2, \dots, 100\}$. Assume the beliefs of players and the payoffs are

- $P(\theta_2|\theta_1) = \frac{1}{100} \quad \forall \theta_2 \in \Theta_2, \forall \theta_1 \in \Theta_1$
- $P(\theta_1|\theta_2) = \frac{1}{100} \quad \forall \theta_1 \in \Theta_1, \forall \theta_2 \in \Theta_2$
- $u_1(a_1, a_2; \theta_1, \theta_2) = \begin{cases} \frac{a_1 + a_2}{2} - \theta_1 & \text{if } a_2 \geq a_1 \\ 0 & \text{otherwise} \end{cases}$
- $u_2(a_1, a_2; \theta_1, \theta_2) = \begin{cases} \theta_2 - \frac{a_1 + a_2}{2} & \text{if } a_2 \geq a_1 \\ 0 & \text{otherwise} \end{cases}$

The beliefs $P(\theta_2|\theta_1)$ and $P(\theta_1|\theta_2)$ are consistent with the prior $P(\theta) = \frac{1}{10000}$, $\forall \theta \in \Theta$, where $\Theta = \Theta_1 \times \Theta_2$.

Example 12.6 (Sealed-Bid Auction) *In this game, we have an auctioneer (who is not a player) willing to sell a item via an auction and two buyers (these are the competing players) who place sealed bids (secret to each other) on the item. The player who has bid more is given the item for the amount he/she has bid for.*

The type of each player is the value they attach to the item. Assume that the values belong to $[0, 1]$. Hence $\Theta_1 = \Theta_2 = [0, 1]$. Bids also belong to $[0, 1]$, hence this is same as the players' action set. The allocation functions are

$$o_1(b_1, b_2) = \begin{cases} 1 & \text{if } b_1 \geq b_2 \\ 0 & \text{otherwise} \end{cases}$$

$$o_2(b_1, b_2) = \begin{cases} 1 & \text{if } b_2 > b_1 \\ 0 & \text{otherwise} \end{cases}$$

The probability distribution over this continuous range is given by

$$f_1(\theta_2|\theta_1) = 1, \quad \theta_2 \in [0, 1]$$

$$f_2(\theta_1|\theta_2) = 1, \quad \theta_1 \in [0, 1]$$

which are consistent with the joint distribution, $f(\theta_1, \theta_2) = 1 \quad (\theta_1, \theta_2) \in [0, 1]^2$.

The utility representation for the game is

$$u_i(b_1, b_2; \theta_1, \theta_2) = o_i(b_1, b_2) \cdot (\theta_i - b_i), \quad i = 1, 2.$$

Note: In this utility representation we assume that the valuation and bid currency are in the same metric (basically, we assume that the valuation/satisfaction of buyer for a item can be equated with money). Such utility representations are called as quasi-linear utility representation.

Lecture 13: August 30, 2017

Lecturer: Swaprava Nath

Scribe(s): Garima Shakya

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

13.1 Recap

In the previous lecture, for Bayesian games, two different types of utilities were discussed: ex-ante utility and ex-interim utility. Ex-ante utility is the utility of any player before observing own type of profile and is expressed as

$$U_i(\sigma_i, \sigma_{-i}) = \sum_{\theta \in \Theta} P(\theta) U_i(\sigma(\theta), \theta).$$

Where, $\sigma(\theta) = (\sigma_1(\theta_1), \sigma_2(\theta_2), \dots, \sigma_n(\theta_n))$. Hence

$$U_i(\sigma_i, \sigma_{-i}) = \sum_{\theta \in \Theta} P(\theta) \sum_{(a_1, a_2, \dots, a_n) \in A} \left(\prod_{j \in N} \sigma_j(\theta_j, a_j) \right) u(a_1, \dots, a_n, \theta_1, \dots, \theta_n). \quad (13.1)$$

And, while calculating ex-interim utility the player knows own type of profile and is expressed as:

$$U_i(\sigma|\theta_i) = \sum_{\theta_{-i} \in \Theta_i} P(\theta_{-i}|\theta_i) U_i(\sigma(\theta), \theta) \quad (13.2)$$

The relation between the two utilities is expressed as

$$U_i(\sigma_i, \sigma_{-i}) = \sum_{\theta_i \in \Theta_i} P(\theta_i) U_i(\sigma|\theta_i). \quad (13.3)$$

13.2 Equilibrium Concepts

Definition 13.1 (Nash Equilibrium) In a Bayesian game with prior P , (σ^*, P) is a Nash equilibrium if

$$U_i(\sigma_i^*, \sigma_{-i}^*) \geq U_i(\sigma_i', \sigma_{-i}^*), \quad \forall \sigma_i', \forall i \in N \quad (13.4)$$

Therefore, for player i , playing σ_i^* is a best response if other players play σ_{-i}^* before observing her own type.

Definition 13.2 (Bayesian Equilibrium) In a Bayesian game with prior P , (σ^*, P) is a Bayesian equilibrium if

$$U_i(\sigma_i^*, \sigma_{-i}^*|\theta_i) \geq U_i(\sigma_i', \sigma_{-i}^*|\theta_i), \quad \forall \theta_i \in \Theta_i, \forall \sigma_i', \forall i \in N \quad (13.5)$$

Therefore, for player i , playing σ_i^* is a best response if other players play σ_{-i}^* after observing her own type.

Observe that σ'_i in Equation 13.5 can be replaced with pure actions WLOG, i.e. Therefore,

$$U_i(\sigma_i^*, \sigma_{-i}^* | \theta_i) \geq U_i(a_i, \sigma_{-i}^* | \theta_i), \quad \forall \theta_i \in \Theta_i, \forall a_i \in A_i, \forall i \in N. \quad (13.6)$$

This is because if the inequality holds for every pure action $a_i \in A_i$, then it must hold even when such actions are mixed probabilistically.

13.3 Equivalence of the two equilibrium concepts

Theorem 13.3 *In finite Bayesian games (σ^*, P) is a Bayesian equilibrium iff it is a Nash equilibrium.*

Proof: (\Rightarrow): Suppose (σ^*, P) is BE. Then

$$U_i(\sigma_i^*, \sigma_{-i}^* | \theta_i) \geq U_i(\sigma'_i, \sigma_{-i}^* | \theta_i), \quad \forall \sigma'_i, \forall i \in N, \forall \theta_i \in \Theta_i \quad (13.7)$$

Now, the ex-ante utility of player i at $(\sigma_i^*, \sigma_{-i}^*)$ is

$$\begin{aligned} U_i(\sigma_i^*, \sigma_{-i}^*) &= \sum_{\theta_i \in \Theta_i} P(\theta_i) U_i(\sigma_i^*, \sigma_{-i}^* | \theta_i) && \text{(from Eqn. 13.3)} \\ &\geq \sum_{\theta_i \in \Theta_i} P(\theta_i) U_i(\sigma'_i, \sigma_{-i}^* | \theta_i) && \text{(from Eqn. 13.7)} \\ &= U_i(\sigma'_i, \sigma_{-i}^*). \end{aligned}$$

Hence $((\sigma^*, P))$ is a Nash equilibrium.

(\Leftarrow): Suppose (σ^*, P) is a Nash equilibrium. Assume for contradiction that (σ^*, P) is not a Bayesian equilibrium.

Then, $\exists a_i \in A_i$, some $\theta_i \in \Theta_i$, some $i \in N$ such that,

$$U_i(a_i, \sigma_{-i}^* | \theta_i) > U_i(\sigma_i^*, \sigma_{-i}^* | \theta_i). \quad (13.8)$$

Consider the following strategy $\hat{\sigma}_i$ of i

$$\begin{aligned} \hat{\sigma}_i(\theta'_i) &\equiv \sigma_i^*(\theta'_i), \quad \forall \theta'_i \in \Theta_i \setminus \{\theta_i\}, \\ \hat{\sigma}_i(\theta_i, a_i) &= 1 \\ \hat{\sigma}_i(\theta_i, b_i) &= 0, \quad \forall b_i \in A_i \setminus \{a_i\}. \end{aligned}$$

Hence, the ex-ante utility of player i at $(\hat{\sigma}_i, \sigma_{-i}^*)$ is

$$\begin{aligned} U_i(\hat{\sigma}_i, \sigma_{-i}^*) &= \sum_{\tilde{\theta}_i \in \Theta_i} P(\tilde{\theta}_i) U_i(\hat{\sigma}_i, \sigma_{-i}^* | \tilde{\theta}_i) \\ &= \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i) U_i(\hat{\sigma}_i, \sigma_{-i}^* | \tilde{\theta}_i) + P(\theta_i) U_i(\hat{\sigma}_i, \sigma_{-i}^* | \theta_i) \\ &= \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i) U_i(\sigma_i^*, \sigma_{-i}^* | \tilde{\theta}_i) + P(\theta_i) U_i(a_i, \sigma_{-i}^* | \theta_i) \\ &> \sum_{\tilde{\theta}_i \in \Theta_i \setminus \{\theta_i\}} P(\tilde{\theta}_i) U_i(\sigma_i^*, \sigma_{-i}^* | \tilde{\theta}_i) + P(\theta_i) U_i(\sigma_i^*, \sigma_{-i}^* | \theta_i) && \text{(from Eqn. 13.8)} \\ &= U_i(\sigma_i^*, \sigma_{-i}^*). \end{aligned}$$

Which is a contradiction to $(\sigma_i^*, \sigma_{-i}^*)$ being a Nash equilibrium. Thus, our assumption was incorrect and we have proved the theorem. \blacksquare

13.4 Existence of Bayesian Equilibrium

Theorem 13.4 *Every finite Bayesian game has a Bayesian equilibrium.*

Proof: *Idea:* Transform the Bayesian game into a complete information normal form game treating each type a player. The transformed game is represented by $\langle \bar{N}, (A_{\theta_i})_{\theta_i \in \bar{N}}, (U_{\theta_i})_{\theta_i \in \bar{N}} \rangle$, where

$$\begin{aligned} \bar{N} &= \cup_{i \in N} \Theta_i = \{\theta_1^1, \theta_1^2, \dots, \theta_1^{|\Theta_1|}, \theta_2^1, \theta_2^2, \dots, \theta_2^{|\Theta_2|}, \theta_n^1, \theta_n^2, \dots, \theta_n^{|\Theta_n|}\} && \text{(finite by assumption)} \\ A_{\theta_i} &= A_i, \quad \forall \theta_i \in \Theta_i, \forall i \in N \\ U_{\theta_i}(a_{\theta_i}, a_{-\theta_i}) &= \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i} | \theta_i) U_i(a_i(\theta_i), a_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) \end{aligned}$$

Note: A mixed strategy of player θ_i , σ_{θ_i} is a probability distribution over ΔA_i , which is a mixed strategy of player i at type θ_i , $\sigma_i(\theta_i)$ in the original Bayesian game. Similarly, we can show that a MSNE in the transformed game is a Bayesian equilibrium in the original game. Since, by Nash theorem, MSNE exists in the transformed game (which is finite), Bayesian equilibrium exists in the original game. ■

Lecture 14: September 1, 2017

Lecturer: Swaprava Nath

Scribe(s): Kaustubh Rane

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

14.1 Examples of Bayesian equilibrium

14.1.1 Sealed Bid Auction

In this game, we have a seller (who is not a player) willing to sell a commodity via an auction and two buyers (these are the competing players) who place sealed bids (secret to each other) on the commodity.

Values $\in [0, 1]$: these are the types θ_i , $i \in \{1, 2\}$.

Bids $\in [0, 1]$: these are the actions b_i .

The probability distribution over this continuous range is given by:

$$f_1(\theta_2 | \theta_1) = f_1(\theta_2) = 1, \quad \theta_2 \in [0, 1].$$

$$f_2(\theta_1 | \theta_2) = f_2(\theta_1) = 1, \quad \theta_1 \in [0, 1].$$

The above two are consistent with the common prior: $f(\theta_1, \theta_2) = 1$, $(\theta_1, \theta_2) \in [0, 1]^2$

14.1.1.1 First Price Auction

Highest bidder is winner, who then has to pay the bid.

Utilities :

$$u_1(b_1, b_2, \theta_1, \theta_2) = (\theta_1 - b_1) I\{b_1 \geq b_2\}$$

$$u_2(b_1, b_2, \theta_1, \theta_2) = (\theta_2 - b_2) I\{b_1 < b_2\}$$

Say $b_1 = s_1(\theta_1) = \alpha_1 \theta_1$ and $b_2 = s_2(\theta_2) = \alpha_2 \theta_2$ (Assuming bid to be a fraction of true valuation) where s_1, s_2 are respective strategies.

Player 1's problem :

$$\begin{aligned} & \max_{\sigma_1} \mathbb{E}[U_1(\sigma_1, \sigma_2^* | \theta_1)] \\ &= \max_{b_1 \in [0, \alpha_2]} \int_0^1 f(\theta_2) (\theta_1 - b_1) I\{b_1 \geq \alpha_2 \theta_2\} d\theta_2 \\ &= \max_{b_1 \in [0, \alpha_2]} (\theta_1 - b_1) \cdot \frac{b_1}{\alpha_2} \end{aligned}$$

Differentiating w.r.t b_1 to maximize, and using $b_1 \in [0, \alpha_2]$, we get

$$b_1 = \min\{\theta_1/2, \alpha_2\}$$

Similarly, $b_2 = \min\{\theta_2/2, \alpha_1\}$.

Thus, $((\frac{\theta_1}{2}, \frac{\theta_2}{2}),$ uniform prior) is a Bayesian equilibrium.

14.1.1.2 Second Price Auction

The player who has the highest bid wins and pays the second highest bid.

Utilities :

$$u_1(b_1, b_2, \theta_1, \theta_2) = (\theta_1 - b_2) I\{b_1 \geq b_2\}$$

$$u_2(b_1, b_2, \theta_1, \theta_2) = (\theta_2 - b_1) I\{b_1 < b_2\}$$

Player 1's problem :

$$\begin{aligned} & \max_{\sigma_1} \mathbb{E}[U_1(\sigma_1, \sigma_2^* | \theta_1)] \\ &= \max_{b_1 \in [0, \alpha_2]} \int_0^1 f(\theta_2)(\theta_1 - \alpha_2 \theta_2) I\{b_1 \geq \alpha_2 \theta_2\} d\theta_2 \\ &= \max_{b_1 \in [0, \alpha_2]} \int_0^{b_1/\alpha_2} (\theta_1 - \alpha_2 \theta_2) d\theta_2 \\ &= \max_{b_1 \in [0, \alpha_2]} \theta_1 \cdot b_1/\alpha_2 - \alpha_2/2 \cdot b_1^2/\alpha_2^2 \end{aligned}$$

Differentiating w.r.t b_1 to maximize, we get

$$b_1 = \theta_1, \text{ and similarly for player 2, } b_2 = \theta_2.$$

Thus, $((\theta_1, \theta_2),$ uniform prior) is a Bayesian equilibrium.

Arbitrary prior: For a non-uniform prior, we consider the same maximization problem for player 1.

$$\begin{aligned} & \max_{b_1 \in [0, \alpha_2]} \int_0^{b_1/\alpha_2} f(\theta_2)(\theta_1 - \alpha_2 \theta_2) d\theta_2 \\ &= \max_{b_1 \in [0, \alpha_2]} \theta_1 \cdot F(b_1/\alpha_2) - \alpha_2 \int_0^{b_1/\alpha_2} \theta_2 f(\theta_2) d\theta_2 \\ &= \max_{b_1 \in [0, \alpha_2]} \theta_1 \cdot F(b_1/\alpha_2) - b_1 \cdot F(b_1/\alpha_2) + \alpha_2 \int_0^{b_1/\alpha_2} F(\theta_2) d\theta_2 \quad (\text{integrating by parts}) \end{aligned}$$

Differentiating w.r.t b_1 to maximize, we get

$$\theta_1 \cdot f(b_1/\alpha_2)/\alpha_2 - F(b_1/\alpha_2) - b_1 \cdot f(b_1/\alpha_2)/\alpha_2 + F(b_1/\alpha_2) = 0.$$

Thus we get

$$b_1 = \theta_1, \text{ and similarly for player 2, } b_2 = \theta_2.$$

This condition is the same we got with uniform prior. Hence, every prior will have an equilibrium where the bid is to reveal the true type. Second price auction is therefore called a prior free auction.

14.2 Mechanism Design

In Game Theory, we take an agent's approach and the guarantees are predictive.

In Mechanism Design, we take a designer's approach and the guarantees are prescriptive.

Some examples are:

- Matching (student - university so that nobody breaks their current allocation)
- Auction (Combinatorial)
- Spectrum, IPL
- Voting

14.2.1 Setup

N	$= \{1, 2, \dots, n\}$	Set of players/agents
X		Set of outcomes
Θ_i		Set of types of i (Private information of i)
u_i	$: X \times \Theta_i \rightarrow \mathbb{R}$	Private value model (One's utility dependent only on his type, after fixing the outcome)
u_i	$: X \times \Theta \rightarrow \mathbb{R}$	Interdependent value model (Utility depends on everyone's types, after fixing the outcome)

14.2.2 Examples

Voting $X = \{a, b, c, \dots\}$ – Set of Outcomes – the set of candidates.

θ_i is a linear order over the candidates.

e.g., Let $\theta_1 = a \succ b \succ c$.

v_i is any vNM utility which is consistent with θ_i .

$\Rightarrow v_1(a) > v_1(b) > v_1(c) \quad u_1(a, v_1) = v_1(a)$.

Single Object Allocation $x \in X$ is a tuple (a, p) – allocation and payment.

$p_i \in \mathbb{R}$ (price charged).

$a = (a_1, a_2, \dots, a_n)$ (Whom to allocate).

$a_i \in \{0, 1\} \quad \sum_i a_i \leq 1$ (not given to more than one person).

$\theta_i \in \mathbb{R}$ (Satisfaction if the object is obtained by i).

$u_i(x, \theta_i) = u_i((a, p), \theta_i) = a_i \theta_i - p_i$ quasi-linear payoff/utility (Linear in terms of payment).

Public Project $x = (a, p)$, $x \in X$ is a choice of a project a and tax assigned p .

$$\theta_i : A \rightarrow \mathbb{R} \quad \theta_i \in \mathbb{R}^{|A|}.$$

$$u_i(x, \theta_i) = \theta_i(a) - p_i.$$

Lecture 15: 5th September, 2017

Lecturer: Swaprava Nath

Scribe(s): Divyanshu Shende

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

15.1 Recap of Mechanism Design Basics

In the previous class, we discussed that in mechanism design we look at the game from the designer's perspective. Our motive is to design games such that they follow certain properties in their equilibria. Typical examples of where mechanisms are used include auctions, voting and matchings. Notice that in each of these settings, the players have private information. Formally, we model the setting as follows:

Setup

- $N = \{1, 2, \dots, n\}$ – a set of *players*
- X – a set of *outcomes*
- Θ_i – set of *types* of player i . This is used to model the private information of player i .
- $u_i : X \times \Theta_i \rightarrow \mathbb{R}$ – the *utility* that player i gets. This depends on the outcome and also his true type (private value model).

15.2 Social Choice Function

Mechanism design starts with an objective of the *social planner* or *mechanism designer*. The objective could be, e.g., to allocate an object to the individual who values it the most. However, the values of the agents are their private information or types. Therefore, this objective can be represented by a function that maps the type space of all the agents to the outcome space. Formally, this function, which we call the *social choice function* (SCF), is defined as $f : \Theta \rightarrow X$, where $\Theta = \Theta_1 \times \Theta_2 \times \dots \times \Theta_n$. That is, given a type profile $\theta = (\theta_1, \theta_2, \dots, \theta_n)$, the outcome is given by $f(\theta)$.

15.2.1 Examples of Social Choice Functions

Since SCFs reflect the objective of the social planner, it can be very generic. We will later see that the function can be complex depending on the domain of the problems. However, some commonly used SCFs are described below.

1. *Utilitarian*: A SCF $f : \Theta \rightarrow X$ is *utilitarian* if it *always* chooses the outcome that maximizes the sum of utilities of each player. Formally,

$$f(\theta_i, \theta_{-i}) = \arg \max_{a \in X} \sum_{i \in N} u_i(a, \theta_i).$$

In the above, $\sum_{i \in N} u_i(a, \theta_i)$ is also called the *social welfare* of the outcome a . Therefore, an utilitarian SCF chooses the outcome that maximizes the social welfare.

2. *Egalitarian*: A SCF $f : \Theta \rightarrow X$ is *egalitarian* if it *always* chooses the outcome that maximizes the minimum utility that any player gets on that outcome. Formally,

$$f(\theta_i, \theta_{-i}) = \arg \max_{a \in X} \left[\min_{i \in N} u_i(a, \theta_i) \right]$$

Note that in an egalitarian SCF, we look at minimum utility that any player gets under an outcome and then we choose the outcome that maximizes this. This SCF, in some sense, aims to equalize the utilities of the agents.

15.3 Mechanisms

Definition 15.1 (Mechanism) A Mechanism \mathcal{M} is a collection of message spaces, (M_1, M_2, \dots, M_n) and a decision rule $g : \times_{i \in N} M_i \rightarrow X$. It is typically denoted as $\mathcal{M} = (M_1, M_2, \dots, M_n, g)$ or $\mathcal{M} = \langle M, g \rangle$, where $M = \times_{i \in N} M_i$.

The idea of message spaces has been kept abstract on purpose to allow the model to be generic. One can think of each message space as a set of actions that a player can take. Players communicate with the central authority using messages and these messages are used by the decision function g to decide an outcome. There is a special case in which the central authority simply asks the players to report their types. Such a mechanism is called a *direct mechanism* and is defined as follows.

Definition 15.2 (Direct Mechanism) A mechanism $\mathcal{M} = \langle M, g \rangle$ is a direct mechanism if $M_i = \Theta_i, \forall i \in N$ and g is the SCF f . Therefore, $\langle \Theta, f \rangle$ is a direct mechanism.

Typically, the message that a player sends to the central authority depends on the player's type so we can also think of *message mapping* $m_i : \Theta_i \rightarrow M_i$ as a strategy. Therefore, a player's action depends on the type that player receives. Suppose player i has received a type $\theta_i \in \Theta_i$. We now give the following definition:

Definition 15.3 (Weakly Dominant Message) Let $\mathcal{M} = \langle M, g \rangle$ be a mechanism. A message $m_i \in M_i$ is weakly dominant for player i at θ_i if

$$u_i(g(m_i, m_{-i}), \theta_i) \geq u_i(g(m'_i, m_{-i}), \theta_i), \quad \forall m'_i \in M_i, \forall m_{-i} \in M_{-i}.$$

Note that we define a message $m_i \in M_i$ as a weakly dominant *for a player i when his type is θ_i* . A weakly dominant message is optimal for player i (given θ_i) irrespective of what the other players communicate.

Definition 15.4 (Implemented in Dominant Strategies) An SCF $f : \Theta \rightarrow X$ is implemented in dominant strategies by a mechanism $\mathcal{M} = \langle M, g \rangle$ if the following two conditions hold:

1. $\forall i \in N, \exists m_i : \Theta_i \rightarrow M_i$ such that $\forall \theta_i \in \Theta_i, m_i(\theta_i)$ is weakly dominant for player i at θ_i , and
2. $g(m_i(\theta_i), m_{-i}(\theta_{-i})) = f(\theta_i, \theta_{-i}), \forall \theta \in \Theta$.

In the above case, we also say that f is *dominant strategy implementable* (DSI) and $\mathcal{M} = \langle M, g \rangle$ implements f .

Definition 15.5 (Strategyproof or Dominant Strategy Incentive Compatible) A direct mechanism $\langle \Theta, f \rangle$ is strategyproof or dominant strategy incentive compatible (DSIC) if

$$u_i(f(\theta_i, \hat{\theta}_{-i}), \theta_i) \geq u_i(f(\theta'_i, \hat{\theta}_{-i}), \theta_i), \forall \theta_i, \theta'_i \in \Theta_i, \forall \hat{\theta}_{-i} \in \Theta_{-i}, \forall i \in N.$$

Strategyproof says that no matter what the other players do (i.e., whether or not they reveal their types truthfully), revealing your type truthfully is at least as good as every other strategy (recall that a strategy is a message mapping). Formally, $m_i(\theta_i) = \theta_i$ is a weakly dominant message for player $i, \forall \theta_i \in \Theta_i, \forall i \in N$.

15.4 Revelation Principle for DSI SCFs

Theorem 15.6 *If there exists an indirect mechanism that implements a SCF f in dominant strategies, then there exists a direct mechanism to implement f (i.e., f is DSIC).*

Proof: Let f be implemented in dominant strategies by the indirect mechanism $\mathcal{M} = \langle M, g \rangle$. Hence $\forall i \in N, \exists m_i : \Theta_i \rightarrow M_i$ such that $m_i(\theta_i)$ is weakly dominant for player i at $\theta_i, \forall \theta_i \in \Theta_i$. Hence $\forall i \in N$

$$u_i(g(m_i(\theta_i), m''_{-i}), \theta_i) \geq u_i(g(m', m''_{-i}), \theta_i), \forall m' \in M_i, \forall m''_{-i} \in M_{-i}, \forall \theta \in \Theta.$$

And

$$g(m_i(\theta_i), m_{-i}(\theta_{-i})) = f(\theta_i, \theta_{-i}), \forall \theta \in \Theta.$$

Since the above equations hold for all $\forall m' \in M_i, \forall m_{-i} \in M_{-i}, \forall \theta \in \Theta$, they must also hold when $m' = m_i(\theta'_i)$ and $m''_{-i} = m_{-i}(\theta_{-i})$. Plugging this into the equation, we get

$$u_i\left(g(m_i(\theta_i), m_{-i}(\theta_{-i})), \theta_i\right) \geq u_i\left(g(m_i(\theta'_i), m_{-i}(\theta_{-i})), \theta_i\right), \forall \theta'_i \in \Theta_i, \forall \theta \in \Theta, \forall i \in N.$$

Note that $g(m_i(\theta_i), m_{-i}(\theta_{-i})) = f(\theta_i, \theta_{-i})$ and $g(m_i(\theta'_i), m_{-i}(\theta_{-i})) = f(\theta'_i, \theta_{-i})$ since f is DSI. This gives $\forall i \in N$

$$u_i(f(\theta_i, \theta_{-i})) \geq u_i(f(\theta'_i, \theta_{-i})), \forall \theta'_i \in \Theta_i, \forall \theta \in \Theta.$$

This is precisely the definition of strategyproof social choice functions. Therefore, f is DSIC. ■

15.5 Summary

This lecture formalized the setting from a mechanism designer's perspective by defining a social choice function. The SCF determines the outcome that is taken and is a function of the players' reported types. We then defined indirect mechanisms, where players communicate with the central authority via messages and the central authority determines the outcome using a decision rule (g). A special case of this is when the central authority asks players to report their types and the decision rule is simply the SCF (f). We then

defined weakly dominant messages as those which maximize player i 's utility when his type is θ_i , irrespective of how the other players play. We later defined what it means for a SCF to be implemented in dominant strategies and talked about strategyproof SCFs or dominant strategy incentive compatible (DSIC) SCFs. We concluded by stating and proving the Revelation Theorem which says that if f is implemented by an indirect mechanism in dominant strategies, then there is also a direct mechanism that implements f .

Lecture 16: September 6, 2017

Lecturer: Swaprava Nath

Scribe(s): Sandipan Mandal

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

16.1 Recap

In the last class we continued our discussion on *mechanism design*. We first discussed *social choice functions (SCF)*, utility functions for SCFs and their types, defined *mechanism design (MD)* and its types. Then we went on to discuss *weakly dominant strategy* w.r.t. mechanism and conditions under which a SCF is implemented in dominant strategy. It was followed by discussion on *strategyproofness* or *dominant strategy incentive compatibility (DSIC)* and *revelation principle*.

Before going forward let's revisit the definition of DSIC as it will be used later to illustrate a subtle point w.r.t. Bayesian incentive compatibility. The definition is as follows:

Definition 16.1 A direct mechanism $\langle \Theta, f \rangle$ is strategyproof or dominant strategy incentive compatible (DSIC) if

$$u_i(f(\theta_i, \tilde{\theta}_{-i}), \theta_i) \geq u_i(f(\theta'_i, \tilde{\theta}_{-i}), \theta_i), \quad \forall \theta_i, \theta'_i \in \Theta_i, \forall \tilde{\theta}_{-i} \in \Theta_{-i}, \forall i \in N. \quad (16.1)$$

This definition states that irrespective of whether others are reporting their types to the central authority truthfully or not, for any agent i , reporting her type truthfully is a weakly dominant strategy.

16.2 Relating Mechanism Design to Bayesian Games

Suppose the types are generated from a common prior P and type of any player i , θ_i is revealed only to the respective players. We can consider the mechanism design scenario in a Bayesian game setup, where the Bayesian game is as follows.

$$\langle N, (A_i)_{i \in N}, (\Theta_i)_{i \in N}, P, (\Gamma_\theta)_{\theta \in \Theta} \rangle.$$

- N = Set of Players $\{1, 2, \dots, n\}$
- M_i = Message space of Player i (Corresponding to Action set in Bayesian Games)
- Θ_i = Set of type for player i
- P = Common prior over $\times_{i \in N} \Theta_i$
- $\Gamma_\theta = \langle N, (A_i)_{i \in N}, (u_i(\theta))_{i \in N} \rangle$

Consider the function $m_i : \Theta_i \rightarrow M_i$. We can observe that given any type, the function m_i maps to a member of the message space. It is equivalent to pure strategies of players in a normal Bayesian game. Intuitively every player sends some message to central authority to maximize its utility which are equivalent to strategy a player might have taken in the case of Bayesian games.

16.3 SCF implemented in Bayesian Equilibrium

Definition 16.2 A mechanism $\langle M, g \rangle$ implements SCF f in Bayesian equilibrium if the following two conditions hold.

1. $\exists(m_1, m_2, \dots, m_n)$ s.t. $m_i(\theta_i)$ maximizes the ex-interim utility of agent i , $\forall \theta_i \in \Theta_i, \forall i \in N$, i.e.,

$$\mathbb{E}_{\theta_{-i}|\theta_i}[u_i(g(m_i(\theta_i), m_{-i}(\theta_{-i})), \theta_i)] \geq \mathbb{E}_{\theta_{-i}|\theta_i}[u_i(g(m'_i, m_{-i}(\theta_{-i})), \theta_i)], \quad \forall \theta_i \in \Theta_i, \forall i \in N \quad (16.2)$$

2. $g(m_i(\theta_i), m_{-i}(\theta_{-i})) = f(\theta_i, \theta_{-i})$

Observation: If SCF f is implementable in dominant strategies then f will be implementable in Bayesian equilibrium.

16.4 Bayesian Incentive Compatibility (BIC)

Definition 16.3 A direct mechanism $\langle \Theta, f \rangle$ is Bayesian incentive compatible if

$$\mathbb{E}_{\theta_{-i}|\theta_i}[u_i(f(\theta_i, \theta_{-i}), \theta_i)] \geq \mathbb{E}_{\theta_{-i}|\theta_i}[u_i(f(\theta'_i, \theta_{-i}), \theta_i)], \quad \forall \theta_i, \theta'_i \in \Theta_i, \forall i \in N. \quad (16.3)$$

Observation. From Equations 16.1 and 16.3, we can observe that the condition for a direct mechanism to be DSIC is required to hold for all $\theta_{-i} \in \Theta_{-i}$ but for it to be BIC this condition is not required because we are doing weighted average over all possible θ_{-i} in some sense.

16.5 Revelation Principle for BI SCFs

Theorem 16.4 If a Social Choice Function f is implementable in Bayesian Equilibrium then f is Bayesian Incentive Compatible.

The proof of the above theorem is similar to the proof of Revelation Principle for DSI SCFs and hence left as an exercise.

16.6 Arrovian Social Welfare Function (SWF)

Even before considering strategyproofness, how can we aggregate the individual preferences into a social preference?

To further understand the implication of this question consider the following example. Suppose three friends want to watch a movie every day for a week. Seven movies are chosen – they are the alternatives. But everyone wants to watch their favorite movie first and does not want to watch an already watched movie. Hence, everyone has a different preference order over the movies. However the three friends want to watch the movies *together*. So what is the best order or watching the movies together? In other words, how can we have a social ordering of these alternatives? The question leads us to study of *Arrovian Social Welfare Functions* (ASWF).

Definition 16.5 *Arrovian Social Welfare Function* takes as input individual preferences of different agents outputs a social preference for all the agents.

Notation

- Set of alternatives $A = \{a_1, \dots, a_m\}$.
- Agents $N = \{1, 2, \dots, n\}$.
- $aR_i b$: Alternative a is atleast as good as b for agent i .
- Set of all possible ordering \mathcal{R} .

Properties of R_i (ordering)

- *Completeness*: For every $a, b \in A$ either $aR_i b$ or $bR_i a$
- *Reflexivity*: $\forall a \in A, aR_i a$
- *Transitivity*: If $aR_i b, bR_i c$ then $aR_i c$

Lets divide the relation R_i into two parts P_i (asymmetric part) and I_i (symmetric part) i.e.,

- $aP_i b$: Alternative a is strictly better than b for i .
- $aI_i b$: Alternative a is indifferent to b for i .

Definition 16.6 An ordering R_i is linear if $aR_i b$ and $bR_i a$, then $a = b$, i.e., indifference is not allowed.

16.6.1 Arrovian Social Welfare Function Representation

Using the notations we have discussed have Arrovian Social Welfare Function F is represented as $F : \mathcal{R}^n \rightarrow \mathcal{R}$. Since $F(R)$ is an ordering over the alternatives like R_i we can split it into two parts $\hat{F}(R)$ (asymmetric part) and $\bar{F}(R)$ (symmetric part) where $R = (R_1, \dots, R_n)$.

16.6.2 Desirable Properties for Arrovian SWF

Definition 16.7 A Social Welfare Function F satisfies weak Pareto (WP) if

$$\forall a, b \in A, [aP_i b, \forall i \in N] \implies [a\hat{F}(R)b].$$

Definition 16.8 A Social Welfare Function F satisfies strong Pareto if

$$\forall a, b \in A, [aR_i b, \forall i \in N, \exists j, aP_j b] \implies [a\hat{F}(R)b].$$

Clearly, strong Pareto \implies weak Pareto.

Lecture 17: September 8, 2017

Lecturer: Swaprava Nath

Scribe(s): Souradeep Chandra

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

17.1 Recap

We defined the Arrovian social welfare function (ASWF) to be a mapping from the set of all preference profiles of n agents to a single preference profile. Hence it is a function $F : \mathcal{R}^n \rightarrow \mathcal{R}$, where \mathcal{R} is the set of all possible orderings over $|A|$ candidates. Two desirable properties that were listed were

Definition 17.1 A Social Welfare Function F satisfies weak Pareto (WP) if

$$\forall a, b \in A, [aP_i b, \forall i \in N] \implies [a\hat{F}(R)b].$$

Definition 17.2 A Social Welfare Function F satisfies strong Pareto if

$$\forall a, b \in A, [aR_i b, \forall i \in N, \exists j, aP_j b] \implies [a\hat{F}(R)b].$$

Clearly, strong Pareto \implies weak Pareto. The other desirable property in the ASWF setup is *independence of irrelevant alternatives*.

17.2 Independence of Irrelevant Alternatives

This property is the crux of Arrow's impossibility result. It is a property that connects two different preference profiles.

Two preferences of player i , say R_i and $R'_i \in \mathcal{R}$ are said to *agree* over $\{a, b\}$ if for agent i

- $aP_i b \Leftrightarrow aP'_i b$
- $bP_i a \Leftrightarrow bP'_i a$
- $aI_i b \Leftrightarrow aI'_i b$

We denote this using the notation $R_i|_{a,b} = R'_i|_{a,b}$. Two preference profiles R, R' agree if for every $i \in N$, $R_i|_{a,b} = R'_i|_{a,b}$ and is denoted by

$$R|_{a,b} = R'|_{a,b}.$$

Definition 17.3 (Independence of Irrelevant Alternatives) An ASWF F satisfies independence of irrelevant alternatives (IIA) if for all $a, b \in A$

$$[R|_{a,b} = R'|_{a,b}] \implies [F(R)|_{a,b} = F(R')|_{a,b}].$$

Illustration Consider an ASWF F , where given the position of the ranking for every agent, some scores are assigned to the candidates. Formally, say the score vector is $(s_1, s_2, s_3, \dots, s_m)$, $s_i \geq s_{i+1}, i = 1, 2, \dots, m - 1, s_i \geq 0, \forall i \in N$. Finally all scores of a particular candidate are added and the final ranking is based on the decreasing order of these scores. This is one special class of ASWF.

Some well-known scoring rules are described below:

- **Plurality:** In this case we assign top score, i.e 1 to s_1 and 0 to all others, So $s_1 = 1, \text{and } s_2 = s_3 = \dots = s_m = 0$.

Question: Does plurality satisfy IIA?

Consider two preference profiles R and R' . The preferences of 4 voters are as follows.

R	R'
$a \ a \ c \ d$	$d \ c \ b \ b$
$b \ c \ b \ c$	$a \ a \ c \ a$
$c \ b \ a \ b$	$b \ b \ a \ d$
$d \ d \ d \ a$	$c \ d \ d \ c$

Plurality gives a social ordering between a and b as:

$$a \hat{F}^{Plu}(R)b, \text{ and } b \hat{F}^{Plu}(R')a.$$

However, we see that the ordering of a and b remains same for every agent in R and R' . IIA would require that the social ordering remain unchanged, which does not happen for plurality. Thus we conclude that plurality does not satisfy IIA.

- **Borda:** The scoring rule in this case is: $s_1 = m - 1, s_2 = m - 2, \dots, s_{m-1} = 1, s_m = 0$.
- **Veto:** The scoring rule is: $s_1 = s_2 = \dots = s_{m-1} = 1, s_m = 0$. We can check by suitable examples that neither Borda nor veto satisfies IIA.
- **Dictatorial:** A voting rule is dictatorial if it always selects the preference ordering of a distinguished agent, whom we call the *dictator*. Thus it is trivial that a dictatorial voting rule satisfies IIA.

We are now going to present a classic result in social choice.

Theorem 17.4 (Arrow 1950) For $|A| \geq 3$, if an ASWF F satisfies weak Pareto and IIA then it must be dictatorial.

Proof: The proof of the following two lemmas will lead us to eventually prove Arrow’s theorem. Informally we state the basic statements of the lemmas as follows.

1. *Field Expansion Lemma:* if a group $G \subseteq N, G \neq \emptyset$ is *decisive* over a, b , then it is decisive over all pairs of alternatives. Informally, a decisive group is a group such that if every agent in that group agrees on a ranking between a pair of alternatives, that ranking is reflected in the social ranking. Therefore, with this lemma, it is enough to call a group *decisive* since it implies that it is decisive over all pairs of alternatives.
2. *Group Contraction Lemma:* if a group G is decisive, there exists a strict subset of G that is also decisive.

First we define decisiveness formally.

Definition 17.5 Given $F : \mathcal{R}^n \rightarrow \mathcal{R}$. Let $G \subseteq N, G \neq \emptyset$.

1. G is almost decisive over a, b if

$$[aP_i b, \forall i \in G, \text{ and } bP_j a, \forall j \notin G] \implies [a\hat{F}(R)b]$$

2. G is called decisive over a, b if

$$[aP_i b, \forall i \in G] \implies [a\hat{F}(R)b]$$

We will use the notation $\bar{D}_G(a, b)$ to denote that G is almost decisive over a, b and $D_G(a, b)$ to denote that G is decisive over a, b . Clearly, $D_G(a, b) \implies \bar{D}_G(a, b)$.

Lemma 17.6 (Field Expansion) Let F satisfies weak Pareto and IIA then $\forall a, b, x, y, a \neq b, x \neq y$, we have

$$\bar{D}_G(a, b) \implies D_G(x, y).$$

Proof: We consider the following set of exhaustive cases to prove this lemma.

1. $\bar{D}_G(a, b) \implies D_G(a, y)$ where $y \neq a, b$
2. $\bar{D}_G(a, b) \implies D_G(x, b)$ where $x \neq a, b$
3. $\bar{D}_G(a, b) \implies D_G(x, y)$ where $x \neq a, b$ and $y \neq a, b$
4. $\bar{D}_G(a, b) \implies D_G(x, a)$ where $x \neq a, b$
5. $\bar{D}_G(a, b) \implies D_G(b, y)$ where $y \neq a, b$
6. $\bar{D}_G(a, b) \implies D_G(b, a)$
7. $\bar{D}_G(a, b) \implies D_G(a, b)$

Case 1: Given: $\bar{D}_G(a, b)$, we need to show $D_G(a, y)$. Pick arbitrary R such that,

$$aP_i y, \forall i \in G, \text{ need to show } a\hat{F}(R)y.$$

Construct R' as follows.

$$\begin{array}{cc} G & N \setminus G \\ a \succ b \succ y & b \succ a \text{ and } b \succ y \end{array}$$

Where $a \succ b$ denotes a is more preferred than b . For the agents in $N \setminus G$, we ensure that the ranking of a and y remain identical to the ranking of these two alternatives in R . Therefore we have

$$R|_{a,y} = R'|_{a,y}.$$

Now since $aR'_i b, \forall i \in G$ and $bR'_j a, \forall j \notin G$, by definition of $\bar{D}_G(a, b)$ we conclude that $a\hat{F}(R')b$. Since b is preferred over y by all agents in N , WP implies that $b\hat{F}(R')y$. Using transitivity of $F(R')$, we have, $a\hat{F}(R')y$. Since the relative ranking of a and y in R and R' are same, using IIA we get $a\hat{F}(R)y$.

Case 2: Given: $\bar{D}_G(a, b)$, we need to show $D_G(x, b)$. Pick arbitrary R such that,

$$xP_i b, \forall i \in G, \text{ need to show } x\hat{F}(R)b.$$

Construct R' as follows.

$$\begin{array}{ccc} G & & N \setminus G \\ x \succ a \succ b & & x \succ a \text{ and } b \succ a \end{array}$$

For the agents in $N \setminus G$, we ensure that the ranking of x and b remain identical to the ranking of these two alternatives in R . Therefore we have

$$R|_{x,b} = R'|_{x,b}.$$

Now since $aR'_i b, \forall i \in G$ and $bR'_j a, \forall j \notin G$, by definition of $\bar{D}_G(a, b)$ we conclude that $a\hat{F}(R')b$. Since x is preferred over a by all agents in N , WP implies that $x\hat{F}(R')a$. Using transitivity of $F(R')$, we have, $x\hat{F}(R')b$. Since the relative ranking of a and y in R and R' are same, using IIA we get $x\hat{F}(R)b$.

Case 3:

$$\begin{aligned} \bar{D}_G(a, b) &\implies D_G(a, y), y \neq a, b && \text{(by Case 1)} \\ &\implies \bar{D}_G(a, y) && \text{(by definition)} \\ &\implies \bar{D}_G(x, y), \text{ as } x \neq a, y && \text{(by Case 2)} \end{aligned}$$

Case 4:

$$\begin{aligned} \bar{D}_G(a, b) &\implies D_G(x, b), x \neq a, b && \text{(by Case 2)} \\ &\implies \bar{D}_G(x, b) && \text{(by definition)} \\ &\implies \bar{D}_G(x, a), \text{ as } a \neq b, x && \text{(by Case 1)} \end{aligned}$$

Case 5:

$$\begin{aligned} \bar{D}_G(a, b) &\implies D_G(a, y), y \neq a, b && \text{(by Case 1)} \\ &\implies \bar{D}_G(a, y) && \text{(by definition)} \\ &\implies \bar{D}_G(b, y), \text{ as } b \neq a, y && \text{(by Case 2)} \end{aligned}$$

Case 6:

$$\begin{aligned} \bar{D}_G(a, b) &\implies D_G(x, b), x \neq a, b && \text{(by Case 2)} \\ &\implies \bar{D}_G(x, b) && \text{(by definition)} \\ &\implies \bar{D}_G(a, b), \text{ as } a \neq b, x && \text{(by Case 2)} \end{aligned}$$

Case 7:

$$\begin{aligned} \bar{D}_G(a, b) &\implies D_G(b, y), y \neq a, b && \text{(by Case 5)} \\ &\implies \bar{D}_G(b, y) && \text{(by definition)} \\ &\implies \bar{D}_G(b, a), \text{ as } a \neq b, y && \text{(by Case 1)} \end{aligned}$$

In the next class we will prove the *Group Contraction Lemma* to complete our proof of *Arrow's Theorem*. ■

Lecture 18: September 12, 2017

Lecturer: Swaprava Nath

Scribe(s): Kunal Chaturvedi

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

18.1 Recap

In the last class, we discussed about a key property for social welfare functions, namely independence of irrelevant alternatives (IIA) and saw that scoring rules does not satisfy it. We stated the celebrated Arrow's impossibility theorem.

The proof of Arrow's theorem is via two lemmas. We proved the *field expansion lemma* in the previous class. In this class we will prove the *group contraction lemma*.

18.2 Continuing the proof of Arrow's theorem

Theorem 18.1 (Arrow 1950) *For $|A| \geq 3$, if an ASWF F satisfies weak Pareto and IIA then it must be dictatorial.*

Proof:[(Contd.)] To complete the proof we state and prove the *group contraction lemma*

Lemma 18.2 (Group Contraction) *Let the group $G \subseteq N, G \neq \emptyset$ be decisive. Then $\exists G' \subset G$ which is also decisive.*

Proof: If $|G| = 1$, the lemma trivially holds. For $|G| \geq 2$, consider two subsets of G , namely $G_1 \subset G$ and $G_2 = G \setminus G_1$. Construct a preference profile R as follows.

G_1	G_2	$N \setminus G$
a	c	b
b	a	c
c	b	a

Note that for all $i \in G$, $aP_i b$. Since G is decisive, we get

$$a\hat{F}(R)b. \tag{18.1}$$

Consider two possible cases for the social ordering F between alternatives a and c .

Case 1: $a\hat{F}(R)c$.

Consider G_1 . We see that, by construction

$$aP_i c, \forall i \in G_1, \text{ and } cP_j a, \forall j \notin G_1.$$

Consider any arbitrary preference ordering R' where the above condition holds. Since, F satisfies IIA, we conclude that $a\hat{F}(R)c$. Hence $\bar{D}_{G_1}(a, c)$. Using field expansion lemma, we get that G_1 is decisive.

Case 2: $\neg(a\hat{F}(R)c) \Rightarrow cF(R)a$

Also from Equation 18.1, we have $a\hat{F}(R)b$. Therefore, by transitivity, $c\hat{F}(R)b$.

Consider G_2 . By construction

$$cP_ib, \forall i \in G_2, \text{ and } bP_jc, \forall j \notin G_2.$$

Consider again any arbitrary preference ordering R' where the above condition holds. Since, F satisfies IIA, we conclude that $c\hat{F}(R)b$. Hence $\bar{D}_{G_2}(c, b)$. Using field expansion lemma, we get that G_2 is decisive. ■

Let us finish the proof of Arrow's theorem. By weak Pareto, N is decisive. By Lemma 18.2, $\exists i \in N$ such that $\{i\}$ is decisive. As we have a singleton set which is decisive, we conclude that i is the dictator. ■

Observation: For a given F , the dictator i is unique.

18.3 Social Choice Setting

Arrow's theorem tells us that we cannot hope to find a meaningful aggregation of preferences that satisfies some very basic desired properties. There are restrictions put on the preferences, e.g., single peaked preferences, where there are non-dictatorial results.

However, we will consider a different path of relaxing the conditions. The ASWF setting asks for a social ordering which may be too much to satisfy. An alternative way to formulate the aggregation problem is to consider the setting of "social choice" functions where the outcome is a single alternative instead of a ranking. Hence a *social choice function* is a map $f : \mathcal{P}^n \mapsto A$, where \mathcal{P} is the set of linear orders, i.e., strict preferences. The set A represents the set of alternatives.

A representative case of this kind of social setting is voting.

18.3.1 Examples of Voting Protocols

This is a list of some voting protocols.

1. **Plurality:** Every voter votes for his/her most favorite candidate and the candidate with highest number of votes is the winner. **Example(s):** Voting system in India, USA, Britain, Canada etc.
2. **Plurality with runoff (two stages):** In this case the top two candidates from the first round of voting advance to the second round of voting. In the second round, the highest voted candidate wins overall. **Example(s):** French presidential election, Rajya Sabha election in India.
3. **Approval Voting:** Each voter casts a single vote for as many candidates as he wants. The candidate with the most votes is the winner. **Example(s):** Approval rating is used by the Mathematical Association of America.
4. **Scoring Rule:** Scores are assigned to candidates according to the vector (s_1, s_2, \dots, s_m) . The highest scoring candidate wins. Some scoring rules which we have seen in the previous lecture are.
 - (a) Borda Count
 - (b) Veto Rule

(c) Plurality is a special case of this class.

We calculate scores as we calculated before but now we only consider the highest scorer as the winner.

5. **Maximin:** Candidates are visualized as vertices in a directed graph and are considered pairwise to assign points. Edges point from the winner to the loser in pairwise runoff. Edges are given weights according to the pairwise winning margin. The candidate with the largest margin of weights wins.
6. **Copeland:** Candidate with maximum pairwise wins is the winner (unlike maximin, here the edge weights do not matter).

None of the voting rules is Pareto superior than another. Given some social objective for the social planner, different voting rules perform better or worse. Every voting rule has some merits, and therefore they survived. However, in future lectures, we will consider a general set of preferences where none of these voting rules perform well.

Condorcet Paradox: Pair-wise runoffs can lead to paradoxical situations. Suppose we have three candidates a, b, c and three voters whose preferences are as follows.

Voter 1	Voter 2	Voter 3
a	c	b
b	a	c
c	b	a

In this case no candidate beats everyone else in pairwise elections.

Definition 18.3 (Condorcet Consistency) *If there exists a candidate which is preferred over every candidate in pairwise runoffs, then he should be the winner.*

Note: Condorcet consistency does not hold for all voting schemes. Example : scoring rules.

Lecture 19: September 13, 2017

Lecturer: Swaprava Nath

Scribe(s): Amudala Dheeraj Naidu

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

19.1 Recap

In the last lecture we started motivating the social choice setting with some real examples like voting. As the setting has changed, we need to redefine the axioms because we observe that IIA does not make sense as there is no social ordering anymore. In this lecture we define some new axioms and we notice some similarity between these axioms and the axioms defined in the social welfare setting. Recall the definition of social choice function.

Definition 19.1 *A social choice function f is a mapping from the set of strict preference profiles (\mathcal{P}^n) to the set alternatives A , i.e.,*

$$f : \mathcal{P}^n \mapsto A.$$

The output of the social choice function is an alternative instead of an ordering.

19.2 Definitions

We define the following properties in the social choice setting.

Definition 19.2 (Pareto Domination) *In a preference profile P , an alternative a is Pareto dominated by b if $bP_i a$, $\forall i \in N$.*

We see that Pareto domination is defined for a particular preference profile. Our next property is a property of social choice function which states that if an alternative a is dominated, then the social outcome should not be equal to a . Formally:

Definition 19.3 (Pareto Efficiency) *An SCF f is Pareto efficient (PE) if $\forall P \in \mathcal{P}$, if a is Pareto dominated in P then $f(P) \neq a$.*

An SCF f is unanimous if the most preferred alternatives of all agents are same in a preference profile, then the social choice outcome in that profile must be equal to that alternative. Formally,

Definition 19.4 (Unanimity) *An SCF f is unanimous (UN) if $\forall P$ such that $P_1(1) = P_2(1) = \dots = P_n(1) = a$, then $f(P) = a$.*

In the above definition the subscripts correspond to the agent and arguments in the parentheses correspond to the rank of the alternative in the preference of the agent.

From the definitions of Pareto efficiency and unanimity we can claim that PE implies UN. With a little abuse of notation, we denote the set of all PE SCFs as **PE** and UN SCFs as **UN**.

Claim 19.5 $PE \subset UN$.

Proof: Suppose $F \in PE$. Pick an arbitrary P such that $P_1(1) = P_2(1) = \dots = P_n(1) = a$, i.e., the ‘if’ condition of UN holds. Clearly, this implies that $aP_i b, \forall b \in A \setminus \{a\}, i \in N$. Since every $b \in A \setminus \{a\}$ is Pareto dominated by a , $F(P) \neq b$ by PE. Hence we conclude that $F(P) = a$. Hence F is UN. ■

Note that the containment is strict.

Strict Example: When the top alternatives for all the agents in a preference profile are not the same, UN does not enforce any outcome. In such a case, one can choose an alternative that is strictly dominated to make the SCF not PE.

The next property is onto-ness, which states that for every alternative there exists a preference profile whose social outcome is that alternative.

Definition 19.6 (Onto) An SCF F is onto (*ONTO*) if $\forall a \in A, \exists P \in \mathcal{P}^n$ such that $f(P) = a$.

Similar to our earlier notation, we denote the set of all onto SCFs as **ONTO**.

Claim 19.7 $UN \subset ONTO$.

Proof: Since \mathcal{P} contains all possible preferences over the alternatives, for every $a \in A$, there exists preference orders where a is on top. Consider a preference profile P where every agent has this same preference order having a on top – UN implies that $F(P) = a$. Hence F is onto. ■

Strict example: Consider an SCF defined as follows. Consider a specific sequence of the alternatives, say WLOG (a_1, a_2, \dots, a_m) . For every profile with all agents having the same top alternative, say a_k the SCF picks a_{k+1} , for $k = 1, \dots, m - 1$, and a_1 for $k = m$. This is clearly ONTO, since all alternatives are chosen by the SCF for some profile, but not UN.

19.2.1 Truthfulness

In the setting of social choice, there is no notion of IIA. Rather we will consider the property of truthfulness. It is similar to the definition of dominant strategy incentive compatibility where it is defined in terms of utility representation. However, since here we talk in terms of ordinal preference profiles, the definition will be adapted accordingly.

Definition 19.8 An SCF is manipulable if $\exists i \in N, P_i, P'_i \in \mathcal{P}, P_{-i} \in \mathcal{P}^{n-1}$ such that $f(P'_i, P_{-i})P_i f(P_i, P_{-i})$.

Therefore, an SCF is manipulable when an agent at some profile P is strictly better off by reporting a preference P'_i rather than her true preference P_i .

The SCF f is *non-manipulable* or *strategyproof* (SP) if it is not manipulable by any agent at any profile.

19.2.2 Characterization of strategyproofness

There exists a notion similar to IIA in the context of social choice functions that connects together different preference profiles, for which we need to define dominated sets.

Definition 19.9 (Dominated Set) *The dominated set of an alternative a at a preference ordering P_i is defined as $D(a, P_i) = \{b \in A : aP_i b\}$*

Hence, the dominated set is the collection of all such alternatives that are less preferred than a for a given preference ordering P_i .

Example: consider the following preference profile P_i

$$P_i = \begin{matrix} a \\ b \\ c \\ d \end{matrix} \quad \text{then } D(b, P_i) = \{c, d\}. \tag{19.1}$$

Now we define a structural property that characterizes strategyproofness.

Definition 19.10 (Monotonicity) *An SCF f is monotone (MONO) if the profiles P and P' be such that $f(P) = a$ and $D(a, P_i) \subseteq D(a, P'_i), \forall i \in N$, then $f(P') = a$.*

Illustration: consider two preference profiles P and P' as shown below and suppose that we have $f(P) = a$

$$\begin{array}{ccc} & P & P' \\ \hline & a & a \\ \cdot & & \cdot \quad a \quad a \\ \cdot & a \quad a & \cdot \quad \cdot \quad \cdot \\ \cdot & \cdot \quad \cdot & \cdot \quad \cdot \quad \cdot \\ \cdot & \cdot \quad \cdot & \cdot \quad \cdot \quad \cdot \end{array} \tag{19.2}$$

That is, the relative position of a is getting weakly better from P to P' , which implies that the alternatives that a was dominating in P are expanding in P' . Monotonicity requires that the social choice outcome should remain the same, i.e., $f(P') = a$.

Theorem 19.11 *An SCF f is SP if and only if f is MONO.*

Note: The technique used in the proof is used for proving other results in social choice.

Proof:

Part 1: SP \Rightarrow MONO: Let f be a strategyproof SCF. Consider two profiles P and P' such that $f(P) = a$ and $D(a, P_i) \subseteq D(a, P'_i) \forall i \in N$. To show that f is monotone, we break the transition from P to P' into n stages such that in each stage the preference of exactly one agent changes as follows.

$$\begin{array}{ccccccc} (P_1, P_2, \dots, P_n) & \rightarrow & (P'_1, P_2, \dots, P_n) & \rightarrow & (P'_1, P'_2, P_3, \dots, P_n) & \rightarrow & (P_1, P_2, \dots, P'_{k-1}, P'_k, \dots, P_n) \\ P = P^{(0)} & & P^{(1)} & & P^{(2)} & \dots & P^{(k)} \\ & & & & & & \downarrow \\ & & & & & & (P'_1, P'_2, \dots, P'_n) \\ & & & & & & P^{(n)} = P' \end{array} \tag{19.3}$$

We claim that the social outcome should remain the same in these transitions.

Claim 19.12 $f(P^{(k)}) = a, \forall k = 1, 2, \dots, n.$

Since $P^{(n)} = P'$, this claim proves that f is monotone.

Proof: We prove the claim by contradiction. Suppose $\exists k$ such that

$$\begin{aligned} f\left(P^{(k-1)}\right) &= a, \quad \text{where } (P'_1, P'_2, \dots, P'_{k-1}, P_k, \dots, P_n) = P^{(k-1)} \text{ and,} \\ f\left(P^{(k)}\right) &= b \neq a, \quad \text{where } (P'_1, P'_2, \dots, P'_k, P_{k+1}, \dots, P_n) = P^{(k)}. \end{aligned} \quad (19.4)$$

Given the if part of monotone property, i.e., by moving from P to P' the relative position of a is weakly increasing there are three different possible cases.

- Case 1: $aP_k b$ and $aP'_k b$: then at profile $P^{(k)}$ player k is better off by reporting P_k – thereby securing the outcome to be a .
- Case 2: $bP_k a$ and $bP'_k a$: then at profile $P^{(k-1)}$ player k is better off by reporting P'_k – thereby securing the outcome to be b .
- Case 3: $bP_k a$ and $aP'_k b$: then player k will manipulate in both profiles – at profile $P^{(k-1)}$, k is better off by reporting P'_k – thereby securing the outcome to be b , and at profile $P^{(k)}$, k is better off by reporting P_k – thereby securing the outcome to be a .

All the three cases are contradictions to f being strategyproof. Hence we have proved the claim. \blacksquare

Part 2: SP \Leftarrow MONO: We prove this as !SP \Rightarrow !MONO. Say for contradiction, $\exists f$ which is not strategyproof but is monotone. Non-strategyproofness implies that f is manipulable, i.e., $\exists i, P_i, P'_i, P_{-i}$ such that $f(P'_i, P_{-i}) P_i f(P_i, P_{-i})$. Denote the social outcomes as $f(P'_i, P_{-i}) = b$ and $f(P_i, P_{-i}) = a$.

Construct another preference profile P'' such that $P'' = (P''_i, P_{-i})$ and also P''_i has the alternatives b and a in the top two positions as shown below.

$$\begin{array}{l} P'' = (P''_i, P_{-i}) \\ \quad b \quad \cdot \\ \quad a \quad \cdot \quad P''_i(1) = b, \text{ and } P''_i(2) = a. \\ \quad \cdot \quad \cdot \\ \quad \cdot \quad \cdot \end{array} \quad (19.5)$$

Now we look at the following two transitions and apply the definition of monotonicity.

- Transition $P \rightarrow P''$: note that in P , for player i , $bP_i a$ and in P''_i , since b and a takes the top two positions respectively, it is clear that a 's position has weakly improved. For the other agents, the preferences did not change. Hence

$$D(a, P_j) \subseteq D(a, P''_j), \forall j \in N.$$

Since f is monotone, we have $f(P'') = a$.

- Transition $P' \rightarrow P''$: Since b is at the top position in P''_i , and for the other agents, the preferences did not change, we have

$$D(b, P'_j) \subseteq D(b, P''_j) \forall j \in N.$$

Monotonicity of f gives $f(P'') = b \neq a$ which is a contradiction.

From Parts 1 and 2, we have the theorem. ■

We have seen that $PE \subset UN \subset ONTO$. However, the next result shows that they are equivalent under strategyproofness.

Lemma 19.13 *If f is monotone and onto then it is Pareto efficient.*

Proof: Assume for contradiction that f is monotone and onto but not Pareto efficient. Hence

$$\exists a, b, P \text{ such that } bP_i a, \forall i \in N \text{ but } f(P) = a.$$

Since f is onto

$$\exists P' \text{ such that } f(P') = b.$$

Construct P'' as follows.

$$\begin{array}{c} P'' \\ \hline b \\ a \\ \vdots \end{array} \quad P_i(1) = b, \text{ and } P_i(2) = a \forall i \in N. \tag{19.6}$$

From the definition of monotonicity, we see that

- for $P \rightarrow P''$, we see that $f(P) = a$ and a 's relative position gets weakly better from P to P'' for all $i \in N$, then by monotonicity $f(P'') = a$.
- similarly for $P' \rightarrow P''$, we see that $f(P') = b$ and b 's relative position gets weakly better from P' to P'' for all $i \in N$, by monotonicity $f(P'') = b$. This is a contradiction. ■

The conclusions of this lecture is therefore captured in Figures 19.1 and 19.2 – the properties discussed coincides under strategyproofness.




Figure 19.1: Representation of all the properties of Social choice function




Figure 19.2: The big picture of all axioms of social choice setting

Lecture 20: September 15, 2017

*Lecturer: Swaprava Nath**Scribe(s): Aayush Ojha*

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

20.1 Recap

In the last lecture we defined Pareto efficiency (PE), unanimity (UN), and ontiness (ONTO) for social choice functions. We also showed that $PE \implies UN \implies ONTO$. Then we defined monotonicity and showed that that is equivalent to strategyproofness. In this lecture, we will look at the Gibbard-Satterthwaite theorem which states if number of alternatives are more than 2 then any onto and monotone social choice function will be dictatorial.

20.2 Gibbard-Satterthwaite theorem

We define dictatorial social choice function.

Definition 20.1 *A social choice function is **dictatorial** if it always selects the first preference of a distinguished agent, which is the dictator.*

First we look at a result from last lecture which will be used in showing the Gibbard-Satterthwaite theorem.

Theorem 20.2 *A SCF, f is onto and strategyproof $\Leftrightarrow f$ is Unanimous and strategyproof $\Leftrightarrow f$ is Pareto efficient and strategyproof.*

We will now formally state the theorem.

Theorem 20.3 (Gibbard (1973), Satterthwaite (1975)) *Suppose $|A| \geq 3$. If f is onto and strategyproof then f is dictatorial.*

Note that, due to Theorem 20.2, this theorem can equivalently be stated w.r.t. unanimity or Pareto efficiency.

20.2.1 Discussions

20.2.1.1 Restricted Preferences

We assume that all preference profile i.e. all possible order for each agent is possible. In a setting where these preferences are restricted, Gibbard-Satterthwaite theorem may not hold. An example: single-peaked preferences.

20.2.1.2 $|A| = 2$

If number of alternatives is two we can construct a social choice function which is onto and strategyproof but not dictatorial. Plurality with a fixed tie breaking rule is strategyproof, onto and non-dictatorial.

20.2.1.3 Indifferences in preferences

If indifferences are allowed among various alternative, then generally Gibbard-Satterthwaite theorem does not hold. In proof we will use some specific profile constructions. If these profile construction are possible Gibbard-Satterthwaite theorem holds.

20.2.1.4 Cardinalization

Also note that Gibbard-Satterthwaite theorem will hold true even when agents provide real number utilities for each alternative as long as the ordinal order is maintained. Thus Gibbard-Satterthwaite theorem holds on cardinalization as long as the ordinal order is maintained by utilities.

20.2.2 Proof of Theorem 20.3

We will look at proof provided by [Sen01]. For simplicity, we will prove the theorem only for the case when number of agents, $n = 2$. Let $N = \{1, 2\}$.

Lemma 20.4 *Let $|A| \geq 3$ and $N = \{1, 2\}$. If f is onto and strategyproof then for every preference profile P , $f(P) \in \{P_1(1), P_2(1)\}$.*

Proof: First we look at case where first preference of both agent is same i.e. $P_1(1) = P_2(1)$. As f is unanimous using theorem 20.2, $f(P) = P_1(1) = P_2(1)$. Let $P_1(1) = a \neq b = P_2(1)$ and $c \in A$ such that $f(P) = c \neq a, b$. We create following four preference profiles.

P_1	P_2	P_1	P_2'	P_1'	P_2'	P_1'	P_2
a	b	a	b	a	b	a	b
-	-	-	a	b	a	b	-
-	-	-	-	-	-	-	-

The preference P_2' is created by putting a at second preference in P_2 and shifting other alternative by at most one place. Similarly, P_1' is created in similar way from P_1 by placing b at the second position. Note that, $f(P_1, P_2') \in \{a, b\}$ because f is PE (an equivalent condition due to Theorem 20.2) and a Pareto dominates every other alternative except b . But if $f(P_1, P_2') = b$, agent 2 can manipulate by reporting P_2' in place of P_2 where preferences for first agent is P_1 . Since f is strategyproof, it implies that $f(P_1, P_2') = a$. Similarly, we can argue that $f(P_1, P_2') = b$.

Consider the transition from (P_1, P_2') to (P_1', P_2') , position of a weakly improves. Thus using monotonicity of f , we conclude $f(P_1', P_2') = a$. However, if we consider the transition from (P_1', P_2) to (P_1', P_2') , position of b weakly improves. Thus by monotonicity of f , we conclude $f(P_1', P_2') = b$, which is a contradiction. Hence we have the lemma. ■

The above lemma reduces the social choice outcome to the top alternatives of the agents. The next lemma will prove the Gibbard-Satterthwaite theorem for 2 agents.

Lemma 20.5 Let $|A| \geq 3$, $N = \{1, 2\}$, and f is onto and strategyproof. Let $P : P_1(1) = a \neq b = P_2(1)$ and $P' : P_1'(1) = c \neq d = P_2'(1)$. Then

$$\begin{aligned} f(P) = a &\implies f(P') = c, \quad \text{and,} \\ f(P) = b &\implies f(P') = d. \end{aligned}$$

Proof: We will only show that if $f(P) = a$ then $f(P') = c$, since the other case is symmetric and the same proof works. To show this we consider the following exhaustive cases.

1. $c = a$ and $d = b$
2. $c \neq a, b$ and $d = b$
3. $c \neq a, b$ and $d \neq b$
4. $c = a$ and $d \neq a, b$
5. $c = b$ and $d \neq a, b$
6. $c = b$ and $d = a$

Case 1: $c = a$ and $d = b$: Suppose for contradiction $f(P') = d = b$. We construct preference profiles as follows.

P_1	P_2	P_1'	P_2'	\hat{P}_1	\hat{P}_2
a	b	a	b	a	b
-	-	-	-	b	a
-	-	-	-	-	-

Consider transition from (P_1, P_2) to (\hat{P}_1, \hat{P}_2) . Preference for a improves for both agents and $f(P_1, P_2) = a$. Thus by monotonicity, $f(\hat{P}_1, \hat{P}_2) = a$. Next consider transition from (P_1', P_2') to (\hat{P}_1, \hat{P}_2) . Preference for b improves for both agents and $f(P_1', P_2') = b$. Thus by monotonicity, $f(\hat{P}_1, \hat{P}_2) = b$. But $a \neq b$. This gives us a contradiction. Therefore, $f(P') = c = a$.

Case 2: $c \neq a, b$ and $d = b$: Suppose for contradiction $f(P') = d = b$. We construct preference profiles as follows.

P_1	P_2	P_1'	P_2'	\hat{P}_1	P_2
a	b	c	b	c	b
-	-	-	-	a	-
-	-	-	-	-	-

First we consider transition from (P_1', P_2') to (\hat{P}_1, P_2) . Notice that this transition satisfies all constraints of *Case 1*. Hence, $f(\hat{P}_1, P_2) = b$.

Consider preference profile (\hat{P}_1, P_2) . At this profile if agent 1 reports P_1 instead of \hat{P}_1 , the outcome is a which she prefers more than the current outcome b , as $f(\hat{P}_1, P_2) = b$ and $f(P_1, P_2) = a$. This is a contradiction to f being strategyproof. Therefore, $f(P') = c$.

Case 3: $c \neq a, b$ and $d \neq b$: Suppose for contradiction $f(P') = d \neq b$. We construct preference profiles as follows.

P_1	P_2	P_1'	P_2'	\hat{P}_1	\hat{P}_2
a	b	c	d	c	b
-	-	-	-	a	-
-	-	-	-	-	-

We first consider transition from (P_1', P_2') to (\hat{P}_1, \hat{P}_2) . This transition follows the constraints of *Case 2*. Hence, $f(\hat{P}_1, \hat{P}_2) = b$.

Next, we consider transition from (P_1, P_2) to (\hat{P}_1, \hat{P}_2) . This transition also follows the constraints of *Case 2*. Hence $f(\hat{P}_1, \hat{P}_2) = c$. But, $b \neq c$. We have a contradiction. Therefore, $f(P') = c$.

Case 4: $c = a$ and $d \neq a, b$: Suppose for contradiction $f(P') = d \neq a, b$. We construct preference profiles as follows.

P_1	P_2	P_1'	P_2'	\hat{P}_1	\hat{P}_2
a	b	a	d	a	b
-	-	-	-	-	-
-	-	-	-	-	-

We first consider transition from (P_1', P_2') to (\hat{P}_1, \hat{P}_2) . This transition follows constraints of *Case 2* (swap agent 1 and agent 2). Hence $f(\hat{P}_1, \hat{P}_2) = b$.

Next, we consider transition from (P_1, P_2) to (\hat{P}_1, \hat{P}_2) . This transition follows constraints of *Case 1*. Hence $f(\hat{P}_1, \hat{P}_2) = a$. But, $b \neq a$. We have a contradiction. Therefore, $f(P') = a$.

Case 5: $c = b$ and $d \neq a, b$: Suppose for contradiction $f(P') = d \neq a, b$. We construct preference profiles as follows.

P_1	P_2	P_1'	P_2'	\hat{P}_1	\hat{P}_2
a	b	b	d	a	d
-	-	-	-	-	-
-	-	-	-	-	-

We first consider transition from (P_1', P_2') to (\hat{P}_1, \hat{P}_2) . This transition follows constraints of *Case 4* (swap agent 1 and agent 2). Hence $f(\hat{P}_1, \hat{P}_2) = d$.

Next, we consider transition from (P_1, P_2) to (\hat{P}_1, \hat{P}_2) . This transition also follows constraints of *Case 4*. Hence $f(\hat{P}_1, \hat{P}_2) = a$. But, $d \neq a$. We have a contradiction. Therefore, $f(P') = c = b$.

Case 6: $c = b$ and $d = a$: Suppose for contradiction $f(P') = d = a$. We construct preference profiles as follows.

P_1	P_2	P_1'	P_2'	\hat{P}_1	P_2'	\tilde{P}_1	P_2'
a	b	b	a	b	a	x	a
-	-	-	-	x	-	-	-
-	-	-	-	-	-	-	-

We assume $x \neq a, b$. As $|A| \geq 3$ such a x will always exist.

We first consider transition from (P_1', P_2') to (\hat{P}_1, P_2') . This transition follows constraints of *Case 1* (swap agent 1 and agent 2). Hence $f(\hat{P}_1, P_2') = a$.

Next, we consider transition from (P_1, P_2) to (\tilde{P}_1, P_2') . This transition follows constraints of *Case 3*. Hence $f(\tilde{P}_1, P_2') = x$.

Consider preference profile (\hat{P}_1, P_2') . At this profile $f(\hat{P}_1, P_2') = a$. If instead of \hat{P}_1 agent 1 report its preference to be \tilde{P}_1 , outcome will be $f(\tilde{P}_1, P_2') = x$, which is more preferred by agent 1 than a at \hat{P}_1 . This is a contradiction to f being strategyproof. Therefore, $f(P') = c = b$.

This shows that $f(P) = a \implies f(P') = c$. To show that $f(P) = b \implies f(P') = d$ we can use similar arguments.

The remaining part of the proof (for more than two agents) uses induction on the number of agents and is skipped here. An interested reader may look up [Sen01]. ■

References

- [Sen01] Arunava Sen. Another direct proof of the Gibbard-Satterthwaite Theorem. *Economics Letters*, 70(3):381–385, 2001.

Lecture 21: October 3, 2017

Lecturer: Swaprava Nath

Scribe(s): Piyush Bagad

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

21.1 Recap and need for domain restriction

In the previous lectures, we saw a characterization theorem due to Gibbard and Satterthwaite that is widely regarded as a negative result.

Theorem 21.1 (Gibbard (1973), Satterthwaite (1975)) *Let the set of alternatives A be such that $|A| \geq 3$. If the social choice function $f : \mathcal{P}^n \mapsto A$ is onto and strategyproof then f is dictatorial.*

Note that the GS theorem needs unrestricted preferences. One of the reasons of such a restrictive result is that since the domain of the SCF is large, it leaves more opportunities for a potential manipulator. Revisiting the major assumptions used by theorem 21.1 can help us better understand this:

- (*Unrestricted preferences*) For $|A| = m \geq 3$, $|\mathcal{P}| = m!$, i.e., all possible linear orderings are available to be chosen by the agents.
- (*Strategyproofness*) We require the SCF to satisfy

$$\forall i \in N, \forall P'_i, P_i \in \mathcal{P}, \forall P_{-i} \in \mathcal{P}^{n-1}, \\ f(P_i, P_{-i}) P_i f(P'_i, P_{-i}) \text{ or } f(P_i, P_{-i}) = f(P'_i, P_{-i}). \quad (21.1)$$

If we now reduce the domain of the SCF from the set of linear orderings (over A), \mathcal{P}^n , to some subset of \mathcal{P} , the SCFs that are truthful with domain \mathcal{P}^n will continue to be truthful on the subset. However, we can hope to find more SCFs that are truthful (and potentially non-dictatorial) on the new restricted domain. In this and the following lectures, we will look at some of these restricted domains.

21.2 Restricted domains

In this course, we will be studying the following three domain restrictions:

1. Single-peaked preferences
2. Divisible object allocation
3. Quasilinear preferences (sometimes, also referred to as ‘mechanisms with money’)

It is worth stating that each of these subdomains has interesting *non-dictatorial* but *strategyproof* SCFs defined on it. We will be first dealing with single-peaked preferences.

21.3 Single-peaked preferences

In this restricted domain, we set a single common order over the alternatives. Once the common order is chosen, we allow only those preferences from \mathcal{P} which have a single peak with respect to that common order. We will soon be defining such preferences formally for a particular common order.

Note: The common order is a fixed order relation over the alternatives and is common to all agents. It is fixed before the agents pick their preferences.

21.3.1 Motivating examples

1. Facility location: placing a public facility like a school, hospital, or post-office in a city.
2. Political ideology: individual political opinions that can be either left, center, or right.
3. Temperature sensing: an individual is most comfortable at a specific temperature and anything hotter or colder is less preferred.

Each of these examples have a natural ordering over the alternatives. We may place the alternatives on the real line \mathbb{R} according to the common ordering $<$. The ordering need not always be an order relation on the real numbers, but can be any order relation that is *transitive* and *anti-symmetric*. For simplicity, we will consider only alternatives having an order on the real line.

21.3.1.1 An illustration

Let $A := \{a, b, c\}$ be the set of facilities ordered on a real line s.t. $a < b < c$ (referring to the locations of a, b, c). Note, here $|\mathcal{P}| = 6$ and the following table of valid and invalid preferences illustrates that for the set of single-peaked preferences \mathcal{S} , $|\mathcal{S}| = 4$. Also, clearly $\mathcal{S} \subset \mathcal{P}$

Table 21.1: Preferences




Figure 21.1: Invalid preference ordering in single-peaked domain.

We now formally define single-peaked preferences.

Definition 21.2 (Single-peaked preference ordering) A preference ordering P_i (strict order over A) for agent i is single-peaked w.r.t. the common order $<$ over the alternatives if there exists an alternative in the agent's preference denoted by $P_i(1)$ such that

- $\forall b, c \in A$ with $b < c \leq P_i(1)$, we have cP_ib , and
- $\forall b, c \in A$ with $P_i(1) \leq b < c$, we have bP_ic .

The alternative $P_i(1)$ is called the peak of the preference of agent i .

Let \mathcal{S} be the set of single-peaked preferences. As illustrated in the earlier example of this section, $\mathcal{S} \subset \mathcal{P}$. The SCF on this restricted domain is $f : \mathcal{S}^n \mapsto A$.

Definition 21.3 (Manipulability of an SCF) An SCF f is said to be manipulable if $\exists i \in N, P'_i, P_i \in \mathcal{S}, P_{-i} \in \mathcal{S}^{n-1}$ such that

$$f(P'_i, P_{-i}) P_i f(P_i, P_{-i}).$$

Definition 21.4 (Strategyproof) An SCF f is said to be strategyproof if it is not manipulable. Thus, by definition, it implies

$$\begin{aligned} \forall i \in N, \forall P'_i, P_i \in \mathcal{S}, \forall P_{-i} \in \mathcal{S}^{n-1}, \\ f(P_i, P_{-i}) P_i f(P'_i, P_{-i}) \quad \text{or} \quad f(P_i, P_{-i}) = f(P'_i, P_{-i}). \end{aligned} \quad (21.2)$$

Note that in contrast to Eq. 21.1, now the preferences $P'_i, P_i \in \mathcal{S}, P_{-i} \in \mathcal{S}^{n-1}$ and $\mathcal{S} \subset \mathcal{P}$. This implies that we have less number of conditions to satisfy for strategyproofness.

The question is, how does an SCF defined on the restricted domain of single-peaked preferences circumvent the Gibberd-Sattherthwaite theorem? We illustrate this using an example social choice function.

21.3.2 Strategyproof SCFs

Definition 21.5 Define SCF $f : \mathcal{S}^n \mapsto A$, where \mathcal{S} is the set of single-peaked preferences w.r.t. the common order $<$ as, for $P \in \mathcal{S}^n$

$$f(P) = \min_{i \in N} \{P_i(1)\}.$$

Where minimum is taken w.r.t. the order relation $<$. Hence the SCF picks the left-most peak among the peaks of the agents.

Claim 21.6 The SCF f defined above is strategyproof and non-dictatorial.

Proof: For proving strategyproofness, first consider the agent having the peak preference as the left-most alternative. Clearly, she has no reason to misreport her preference order since her peak is chosen.

Now, consider any other agent i – she has a peak to the right of the left-most peak. In other words, for $i \in N$ we have $f(P) < P_i(1)$. The only possible manipulation she could do to change the outcome is to report her peak to be further left of $f(P)$, therefore changing the preference profile from $(P_i, P_{-i}) \rightarrow (P'_i, P_{-i})$ s.t.

$$\begin{aligned} P'_i(1) &\leq f(P) < P_i(1), \\ f(P'_i, P_{-i}) &= \min_{i \in N} \{P_1(1), P_2(1), \dots, P'_i(1), \dots, P_n(1)\} = P'_i(1). \end{aligned}$$

If $P'_i(1) = f(P)$, clearly $f(P_i, P_{-i}) = f(P'_i, P_{-i})$. Else, we will have

$$P'_i(1) < f(P) < P_i(1).$$

Since P_i is a single-peaked preference, thus, we get $f(P_i, P_{-i}) \geq P_i f(P'_i, P_{-i})$. Hence, by definition 21.4, we conclude that f is strategyproof.

Since the identity of the agent having the left-most peak is not fixed before reporting the preferences, f is non-dictatorial. ■

Using similar arguments, we can prove that an SCF that picks the k^{th} alternative from the left will also be *strategyproof* and *non-dictatorial*, $\forall k \in \{1, 2, \dots, |A|\}$. In particular, it holds true for the right-most alternative and median ($k = \lfloor n/2 \rfloor$).

Definition 21.7 (Median Voter SCF) An SCF $f : S^n \mapsto A$ is said to be a Median Voter SCF if $\exists B = (y_1, y_2, \dots, y_{n-1})$ s.t. $f(P) = \text{median}(B, \text{peaks}(P))$, $\forall P \in S^n$. The points in B are called as “peaks of phantom voters” or “phantom peaks”.

Note: B is fixed for a given f and does not change with $P \in S^n$.

21.3.2.1 Advantage of using phantom voters

All the examples stated above, i.e., left-most peak, right-most peak, median, or any k -th peak from the left etc. can be combined into a single definition through this *median voter SCF*.

Claim 21.8 The SCFs picking the left-most most peak, the right-most peak are median voter SCFs.

Proof: If $A = \{a_1, a_2, \dots, a_{|A|}\}$, let $a = \min_{w.r.t. <} A$, $b = \max_{w.r.t. <} A$
Define $y_1, y_2, \dots, y_{n-1}, z_1, z_2, \dots, z_{n-1} \in S^n$ s.t.

$$\begin{aligned} y_i(1) &= a, \forall i \in \{1, \dots, n-1\}, \\ z_i(1) &= b, \forall i \in \{1, \dots, n-1\}. \end{aligned}$$

For the case of the left-most peak SCF, we can choose $B = (a, a, \dots, a)$ which ensures that the median of points in B and peaks reported by the agents will always result in the minimum of the peaks w.r.t $<$ reported by the agents. For the case of right-most SCF, we can choose $B = (b, b, \dots, b)$ and the proof follows as for the case of minimum. ■

Using similar arguments, we can prove that the other SCFs picking any k^{th} peak from left are also Median Vector SCFs.

Theorem 21.9 (Moulin(1980)) Every median vector SCF is strategyproof.

Proof: We need to consider only the peak preferences of all the agents. So let us denote the preferences denoted only by their peaks, i.e., $P = (P_1(1), \dots, P_i(1), \dots, P_n(1))$ and let $f(P) = a \in A$ is the median of these peaks and the phantom peaks. Consider an agent i

- If $P_i(1) = a$, then there is no reason for i to manipulate.

- If $P_i(1) < a$, then if the agent shifts her preference to further left of a , the median will not change. If she manipulates to report her peak to further right of a , *i.e.* $(P_i, P_{-i}) \mapsto (P'_i, P_{-i})$ s.t. $a < P'_i(1)$, this will imply that $P_i(1) < a < P'_i(1)$, and since P_i is a single-peaked preference, by definition 21.4, $a = f(P_i, P_{-i}) = f(P'_i, P_{-i})$. Thus, i has no profitable manipulation.
- If $a < P_i(1)$, again by exactly symmetrical arguments, i has no profitable manipulation.

Hence, f is strategyproof. ■

Note: Mean does not satisfy the property used in the proof above, since changing the peak $P_i(1)$ of agent i on either sides of a will result in a change in the mean.

21.4 Summary

In this lecture, we revisited the Gibbard-Satterthwaite theorem and due to the restrictiveness of the result, we emphasized the need for domain restriction. Among the major domain restrictions, we looked at the single peaked preferences. In particular, we looked at examples of SCFs that are *strategyproof* but *non-dictatorial* and defined the median voter SCF which characterizes a set of these examples.

Lecture 22: October 4, 2017

Lecturer: Swaprava Nath

Scribe(s): Mayank Singour

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

22.1 Properties of SCF (Recap)

In previous lecture we discussed about the restricted domains (single peaked preferences in particular). Now, we discuss properties in the single-peaked domain, some of which are similar to the properties of unrestricted domain. Denote the set of single peaked preferences over the alternatives in A with \mathcal{S} .

Definition 22.1 (Pareto Efficiency (PE)) $\forall P \in \mathcal{S}^n, \forall b \in A$ if $\exists a \in A$ such that $aP_i b \forall i \in N$, then $f(P) \neq b$.

Definition 22.2 (Unanimity (UN)) $\forall P$ with $P_1(1) = P_2(1) = \dots = P_n(1) = a$, then $f(P) = a$.

Definition 22.3 (Onto-ness (ONTO)) $\forall a \in A, \exists P \in \mathcal{S}^n$ such that $f(P) = a$.

As before, $PE \subseteq UN \subseteq ONTO$.

Claim 22.4 Let p_{\min} and p_{\max} are the leftmost and rightmost peaks of P according to $<$, then f is PE iff $f(P) \in [p_{\min}, p_{\max}]$.

Proof: (\Rightarrow) Suppose f is PE but $f(P) \notin [p_{\min}, p_{\max}]$, then $f(P)$ is either at the left side of p_{\min} or at the right side of p_{\max} . Consider the case $f(P) < p_{\min}$, but then every agent prefers p_{\min} over $f(P)$, a contradiction to PE. Similar argument can be given for $p_{\max} < f(P)$.

(\Leftarrow) if $f(P) \in [p_{\min}, p_{\max}]$, every other alternative $b \neq f(P)$ that belongs to $[p_{\min}, p_{\max}]$ will either be close to p_{\min} (and further from p_{\max}) or vice-versa. Then there exists at least one agent that prefers $f(P)$ more than b . Hence, the ‘if’ condition in the definition of PE is never triggered for every such b . Hence, PE is vacuously satisfied. ■

Definition 22.5 (Monotonicity) If for two profiles P and P' with $f(P) = a$ and $D(a, P_i) \subseteq D(a, P'_i), \forall i \in N$, then $f(P') = a$.

- where $D(a, P_i) = \{b \in A : aP_i b\}$ is the set of alternatives dominated by a under the preference P_i .

22.2 Results on restricted domain of preferences

We will see some results similar to the unrestricted preferences but the proofs will differ as we do not have the flexibility of constructing arbitrary preference profiles. In this lecture, we consider only single peaked preferences.


Theorem 22.6 f is strategyproof (SP) $\Rightarrow f$ is monotone (MONO).

This has exactly the same proof as we did for unrestricted preferences. However the construction of our previous result is not always feasible for the converse of this theorem.

Exercise: Find a counterexample of the converse OR prove the converse.

Theorem 22.7 Let an SCF $f : S^n \mapsto A$ be SP, then f is ONTO $\iff f$ is UN $\iff f$ is PE.

Proof: It is enough to show that if f is SP, f is ONTO $\Rightarrow f$ is PE. We prove this by contradiction. Suppose $\exists a, b$ such that $aP_i b, \forall i \in N$ but $f(P) = b$. Since P_i is single peaked, \exists another alternative $c \in A$ which



is adjacent to b (there is no alternative between c and b) such that $cP_i b, \forall i \in N$. Note that, c could be a itself. This is illustrated in the figure above.

ONTO implies that $\exists P'$ such that $f(P') = c$. Construct P'' such that $P''_i(1) = c$ and $P''_i(2) = b, \forall i \in N$. Note that this profile is always possible to construct in a single-peaked domain. [can you think of a similar construction for proving the converse of theorem 22.6?]

Consider the transition of the profile from $P' \rightarrow P''$, by monotonicity $f(P'') = c$. Now consider the transition of the profile from $P \rightarrow P''$, by monotonicity $f(P'') = b$. This is a contradiction since $c \neq b$. ■

22.3 Anonymity

Define permutation over the agents as $\sigma : N \mapsto N$. We apply a permutation σ to a profile P to construct another profile such that the preference ordering of i in P goes to the permuted agent $\sigma(i)$ in the new profile. We denote the new profile by P^σ .

Example: $N = \{1, 2, 3\}, \sigma : \sigma(1) = 2, \sigma(2) = 3, \sigma(3) = 1$.

P_1	P_2	P_3	P_1^σ	P_2^σ	P_3^σ
a	b	b	b	a	b
b	a	c	c	b	a
c	c	a	a	c	c

Anonymity requires that the social outcome should not depend on agent identity.

Definition 22.8 An SCF $f : \mathcal{S}^n \mapsto A$ is anonymous (ANON) if for every profile P and for every permutation of the agents σ , $f(P^\sigma) = f(P)$.

Note: Dictatorship is not anonymous, but median voter rule is anonymous.

Theorem 22.9 (Moulin 1980) A SP SCF f is ONTO and ANON iff it is a median voter SCF.


Proof: (\Leftarrow) Median voter SCF is SP (see theorem 21.9 of previous lecture), and

- it is ONTO, since we can put all voters peaks at the same alternative/location and set all phantom peaks at zero, then that location is the outcome.
- it is ANON, since we can permute the agents with the peaks unchanged and the outcome will not change.

(\Rightarrow) Given $f : \mathcal{S}^n \mapsto A$ is SP, ONTO and ANON. Define

- P_i^0 : agent i 's preference where the peak is at the leftmost point w.r.t. $<$.
- P_i^1 : agent i 's preference where the peak is at the rightmost point w.r.t. $<$.

The preferences are illustrated in the figure below.



Let y_j 's be the phantom peaks, $j = 1, 2, \dots, n-1$. Pick y_j 's as follows.

$$y_j = f(P_1^0, P_2^0, \dots, P_{n-j}^0, P_{n-j+1}^1, \dots, P_n^1).$$

It does not matter which agents have which peaks because of anonymity.

Claim 22.10 $y_j \leq y_{j+1}$, $j = 1, 2, \dots, n-2$.

Proof: Consider y_j as defined before. Therefore,

$$y_{j+1} = f(P_1^0, P_2^0, \dots, P_{n-j-1}^0, P_{n-j}^1, P_{n-j+1}^1, \dots, P_n^1).$$

Strategyproofness implies $y_j P_{n-j}^0 y_{j+1}$ but P_{n-j}^0 is single peaked with the peak at the leftmost position. Therefore $y_j \leq y_{j+1}$. ■

We will complete this proof in the next lecture. ■

Lecture 23: October 6, 2017

Lecturer: Swaprava Nath

Scribe(s): Debojyoti Dey

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

23.1 Recap

In the previous lecture we showed some important properties of social choice function in the restricted domain of single-peaked preferences. The claims we proved are as follows.

1. Let p_{min} and p_{max} are the leftmost and rightmost peaks according to order relation $<$. Then SCF f is PE if and only if $f(P) \in [p_{min}, p_{max}]$.
2. f is SP $\implies f$ is MONO.
3. Let $f : \mathcal{S}^n \rightarrow A$ is SP. Then f is ONTO $\iff f$ is UN $\iff f$ is PE.

We also defined anonymous (ANON) SCF f which is independent of the permutation of the agents for every preference profile P , that is, $f(P) = f(P^\sigma)$ where P^σ represents σ -permuted preferences of P . We observed that a dictatorial SCF cannot be ANON.

23.2 Characterization of strategyproof SCFs in single-peaked domain

We started proving a characterization result for the median voting rule SCF given as follows.

Theorem 23.1 (Moulin 1980) *A SP SCF f is ONTO and ANON if and only if it is a median voting rule.*




Figure 23.1: Special single peaked preferences over $[0, 1] - P_i^0$ and P_i^1 .

Proof: (Continued from the last lecture)

Consider an arbitrary profile

$$P = (P_1, P_2, \dots, P_n)$$

Let $p_i := P_i(1)$ denote the peak of agent i

We claim that $f(P) = \text{med}(p_1, p_2, \dots, p_n, y_1, \dots, y_{n-1})$.

We can assume WLOG that $p_1 \leq p_2 \leq \dots \leq p_n$ due to ANON. Say $a = \text{med}(p_1, p_2, \dots, p_n, y_1, \dots, y_{n-1})$.

Case 1: a is a phantom peak

Say $a = y_j$ for some $j \in 1, 2, \dots, n-1$. This is a median of $(2n-1)$ points. There are $(j-1)$ phantom peaks to the left of the median (due to the fact that $y_j \leq y_{j+1}$) and $(n-1-j)$ to the right. So, there are $(n-j)$ agent peaks on the left. Hence the following holds,

$$p_1 \leq \dots \leq p_{n-j} \leq y_j = a \leq p_{n-j+1} \leq \dots \leq p_n.$$

Now consider two profiles, $(P_1^0, P_2^0, \dots, P_{n-j}^0, P_{n-j+1}^1, \dots, P_n^1)$ and $(P_1, P_2^0, \dots, P_n^1)$. By definition

$$f(P_1^0, P_2^0, \dots, P_{n-j}^0, P_{n-j+1}^1, \dots, P_n^1) = y_j.$$

Assume that

$$f(P_1, P_2^0, \dots, P_n^1) = b$$

Now we see that

$$f \text{ is SP} \implies y_j P_1^0 b \implies y_j \leq b.$$

But also

$$\begin{aligned} f \text{ is SP} &\implies b P_1 y_j \text{ and it is known that } p_1 \leq y_j \\ &\implies b \leq y_j \end{aligned}$$

Combining the above two implications we get, $b = y_j$. Repeating the argument for the first $(n-j)$ agents, we get

$$f(P_1, P_2, \dots, P_{n-j}, P_{n-j+1}^1, \dots, P_n^1) = y_j.$$

Now consider $f(P_1, \dots, P_{n-j}, P_{n-j+1}^1, \dots, P_{n-1}^1, P_n) = b$ (say). Using the SP property of f , we get

$$\begin{aligned} y_j P_n^1 b &\implies b \leq y_j \\ b P_n y_j \text{ and } y_j \leq p_n &\implies y_j \leq b \end{aligned}$$

Combining the above two implications, $b = y_j$. Repeating the arguments, we finally get,

$$f(P_1, P_2, \dots, P_{n-j}, P_{n-j+1}, \dots, P_n) = y_j = a.$$

which is the median.

Case 2: a is an agent peak

We prove this for 2 agents. The general case repeats the argument.

Claim 23.2 Let $N = \{1, 2\}$, and P, P' be such that, $P_i(1) = P'_i(1) \forall i \in N$, then

$$f(P) = f(P').$$

Proof: Let $a = P_1(1) = P'_1(1)$ and $b = P_2(1) = P'_2(1)$. Also let $f(P) = x$ and $f(P'_1, P_2) = y$. Since f is SP, we have xP_1y and yP'_1x . Since peaks in the two profiles are the same, if x and y fall on the same side of the peak $P_1(1)$ (equivalently $P'_1(1)$) they must be the same. The only other possibility is that x and y fall on the different sides of the peak. We show that this is not possible.

WLOG assume that, $x < a < y$ and $a < b$. We know f is SP+ONTO \iff f is SP+PE and PE requires that $f(P) \in [a, b]$. But $f(P) = x < a$, which is a contradiction.

Repeat this argument for the transition of preference profiles $(P'_1, P_2) \rightarrow (P''_1, P'_2)$. ■

Now consider the profile $P = (P_1, P_2)$ such that $P_1(1) = a$ and $P_2(1) = b$ and y_1 be the phantom peak. By assumption, $med(a, b, y_1)$ is an agent peak. WLOG let the median be a . Assume for contradiction, $f(P) = c \neq a$.

By PE, c must lie within a and b . We consider the two cases, $b < a < y_1$ and $y_1 < a < b$.

Case A: $b < a < y_1$

By PE, $c < a$. Construct P'_1 such that $P'_1(1) = a = P_1(1)$ and $y_1P'_1c$ (possible since y_1 and c are on different sides of the peak $P'_1(1)$). Since $f(P) = c$, $f(P'_1, P_2) = c$ by the previous claim.

Now consider the profile (P'_1, P_2) . We have

$$P_2(1) = b < y_1 < P'_1(1).$$

So the median of $(b, P'_1(1), y_1)$ is y_1 , which is a phantom peak, and hence by our result in Case 1,

$$f(P'_1, P_2) = y_1.$$

By construction of P'_1 ,

$$\begin{aligned} y_1P'_1c &\implies f(P'_1, P_2) = y_1 \\ &\implies f \text{ is not SP.} \end{aligned}$$

This is a contradiction. Hence our assumption $f(P) \neq a$ is wrong in this case.

Case B: $y_1 < a < b$

By PE, $a < c$. Construct P'_1 such that $P'_1(1) = a = P_1(1)$ and $y_1P'_1c$ (possible since y_1 and c are on different sides of the peak $P'_1(1)$). By the previous claim

$$f(P) = c \implies f(P'_1, P_2) = c.$$

Now consider the profile (P'_1, P_2) . We have

$$P'_1(1) < y_1 < b = P_2(1) \implies f(P'_1, P_2) = y_1.$$

But by construction of P'_1 ,

$$\begin{aligned} y_1P'_1c &\implies f(P'_1, P_2) = y_1 \\ &\implies f \text{ is not SP.} \end{aligned}$$

This is a contradiction. Hence our assumption $f(P) \neq a$ is wrong in this case too.

Hence we have proved Case 2 of this theorem for 2 agents. ■

23.3 Conclusion

In this lecture, we have proved the non dictatorial nature of median voter SCF by introducing phantom voters. The phantom voters/peaks are introduced so that the extreme preference conditions can be handled with a “fair” decision. For example, if half the agents are at the extreme left and other half is at the extreme right, a fair distribution of phantom peaks may lead to picking the median somewhere at the center rather than at some extreme point. Note that, median voter SCF is actually a class of voting rules.

Lecture 24: October 10, 2017

Lecturer: Swaprava Nath

Scribe(s): Divyanshu Shende

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

24.1 Recap


The previous classes focused on a domain restriction called *single-peaked preferences* where the agents' preferences over their alternatives had a single peak with respect to a pre-defined ordering over the alternatives (on the real line). This is single-peakedness in a single dimension, which can be extended to multiple dimensions by defining suitable orderings. Also, there can be multiple facilities that need to be located – both this and the multi-dimensional case are active areas of research.

24.2 Task Allocation Problem

Consider a task that must be divided among a set of n agents ($N = \{1, 2, \dots, n\}$). The task is such that the agent's payoff only depends on the amount or *share* of the task that the agent gets. In this case, we may represent the entire task as the unit interval $[0, 1]$ on the real line. We denote by s_i , the share of task that agent i gets. Note that the location of this share on the real line is irrelevant to us. Clearly, for all $i \in N$, we have $s_i \in [0, 1]$ and $\sum_{i \in N} s_i = 1$.

We give each agent a *reward* for the work that the agent does and this depends on the share of the task that the agent is assigned. Moreover, each agent has a particular *cost* that he/she incurs and this cost depends on the share too. For example, for agent i whose share of the task is s_i , the reward could be $w \cdot s_i$ (where w is the wage per unit time and is fixed) and the cost that agent i incurs could be $k_i s_i^2$. We can now define the *utility* that agent i gets when his share of task is s_i to be $u_i(s_i) = \text{reward}(s_i) - \text{cost}(s_i)$. In our example, this would be $u_i(s_i) = w s_i - k_i s_i^2$. Note that in general, w could vary with i (but not s_i).

In the above example, the utility function would look something like this:



24.3 Single-peakedness over the Share of Task

The above utility function can be used to define a preference over the agent's possible shares. Looking at the shape, one might easily start thinking of single-peaked preferences. But there is a difference with the earlier definition of single-peakedness 'over the alternatives'. Firstly, an *alternative* is not just the share of agent i herself, but a tuple (s_1, s_2, \dots, s_n) . So, the set of alternatives is $A = \{(s_1, s_2, \dots, s_n) : s_i \in [0, 1] \text{ and } \sum_{i \in N} s_i = 1\}$. Now, suppose that s_i^* is the optimal share of agent i that maximizes her utility. Note that agent i is indifferent between (s_i^*, s_{-i}) and (s_i^*, s'_{-i}) where $s_{-i} \neq s'_{-i}$. Recall that the indifference between two alternatives is not allowed in single-peaked preferences – the peak should be such that it is strictly preferred over all other alternatives. Therefore, single-peakedness 'over the share of the task' is different from that 'over the alternatives'.

However, there is still some structure to the preferences that we can look at. Observe that agent i has her own *peak* which represents the ideal share of the task that she would like. We would like to retain this structure. So we define \mathcal{S} as the set of all such preferences such that the corresponding utility functions over $[0, 1]$ are "single-peaked". Thus, we now define single-peakedness over the share of an agent's task. We can now define a *social choice function* as $f : \mathcal{S}^n \rightarrow A$.

Consider any $P \in \mathcal{S}^n$. We define $f(P) = (f_1(P), f_2(P), \dots, f_n(P))$, where $f_i(P)$ is agent i 's share of the task. Note that $f(P) \in A$ so we have $\forall i, f_i(P) \in [0, 1]$ and $\sum_{i \in N} f_i(P) = 1$. Also, we let p_i denote the peak of agent i , i.e., p_i is the (unique) peak of the preference P_i of agent i . We now look a familiar property of social choice functions.

24.4 Pareto Efficiency

Definition 24.1 (*Pareto Efficient SCF*) A Social Choice Function f is Pareto Efficient (PE) if, $\forall P \in \mathcal{S}^n, \nexists a \in A$ such that, $a R_i f(P), \forall i \in N$ and $\exists j \in N, a P_j f(P)$.

The above definition says that for any preference profile $P \in \mathcal{S}^n$ there should *not* exist any alternative a such that some agent j strictly prefers a to $f(P)$ and the other agents prefer a at least as much as $f(P)$. This very much resembles the if-condition of the *Strong Pareto* property of SCFs [See Lecture 16]. In informal terms, f is "optimal" in the sense that no other alternative is weakly preferred over $f(P)$, $\forall P \in \mathcal{S}^n$. We can also say that no other alternative *pareto dominates* $f(P)$.

24.4.1 Characterizing Pareto Efficiency

Theorem 24.2 (*PE characterization*) Let f be a social choice function in the above setting. Then, f is PE iff the following conditions hold for all $P \in \mathcal{S}^n$ (recall that p_i 's denote the peaks for each agent):

1. $\sum_{i \in N} p_i = 1 \implies f_i(P) = p_i, \forall i \in N.$
2. $\sum_{i \in N} p_i > 1 \implies f_i(P) \leq p_i, \forall i \in N.$
3. $\sum_{i \in N} p_i < 1 \implies f_i(P) \geq p_i, \forall i \in N.$

Proof: (\implies) Suppose f is PE. Let P be any preference profile in \mathcal{S}^n . Now we consider the three cases.

$$1. \sum_{i \in N} p_i = 1.$$

Suppose $f_i(P) < p_i$ for some i (similar argument will hold for $>$). Then, there exists $j \in N \setminus \{i\}$ such that $f_j(P) > p_j$ (else we would have $\sum_{k \in N} f_k(P) < 1$). Note that if we increase the share of i by ϵ (> 0) and decrease the share of j by ϵ and keep the share of everyone else the same, we still get a valid alternative (for suitable values of ϵ). We define a new alternative in which we choose ϵ so that $f_i(P) < f_i(P) + \epsilon \leq p_i$ and $p_j \leq f_j(P) - \epsilon < f_j(P)$. Now, since nobody else's share has changed, they are indifferent to the new ordering. However, agents i and j are strictly better off since both of them move closer to their peaks. So, our new alternative *pareto dominates* $f(P)$. Therefore, f is not PE. ($\Rightarrow \Leftarrow$). Therefore, $f_i(P) = p_i, \forall i \in N$.

$$2. \sum_{i \in N} p_i > 1.$$

Suppose $\exists i \in N$ such that $f_i(P) > p_i$. Then, there exists $j \in N \setminus \{i\}$ such that $f_j(P) < p_j$ (else we would have $\sum_{k \in N} f_k(P) > 1$). By the same ϵ -shift argument as last case, we can construct an alternative that *pareto dominates* $f(P)$. Therefore f is not PE. ($\Rightarrow \Leftarrow$) Therefore, $f_i(P) \leq p_i, \forall i \in N$.

$$3. \sum_{i \in N} p_i < 1.$$

Suppose $\exists i \in N$ such that $f_i(P) < p_i$. Then, there exists $j \in N \setminus \{i\}$ such that $f_j(P) > p_j$ (else we would have $\sum_{k \in N} f_k(P) < 1$). Repeat the ϵ -shift argument in the two cases above to construct an alternative that *pareto dominates* $f(P)$. Therefore f is not PE. ($\Rightarrow \Leftarrow$) Therefore, $f_i(P) \geq p_i, \forall i \in N$.

Therefore, if f is PE then the three conditions must hold. This proves the ‘only if’ part of the theorem.

(\Leftarrow) Now, suppose the three conditions hold. We take 3 cases as above.

$$1. \sum_{i \in N} p_i = 1 \text{ and } f_i(P) = p_i, \forall i \in N.$$

In this case, every agent receives gets peaks (by condition 1). Since peaks are most preferred for each agent individually, this outcome Pareto dominates all others. So in this case, f outputs an alternative that is not *pareto dominated* by any other alternative.

$$2. \sum_{i \in N} p_i > 1 \text{ and } f_i(P) \leq p_i, \forall i \in N.$$

To see that $f(P)$ is not *pareto dominated*, we only need to observe the following: since every agent is on the left side of their peaks ($f_i(P) \leq p_i$) and the sum of their shares sums to one, we cannot have an alternative where everyone moves ‘‘closer’’ to their peaks. If an agent moves closer to his/her peak, there must be another agent who moves away and therefore has strictly lower payoff. Therefore, there cannot exist an alternative that Pareto dominates $f(P)$.

$$3. \sum_{i \in N} p_i < 1 \text{ and } f_i(P) \geq p_i, \forall i \in N.$$

This argument in this case is a repetition of the above argument.

Note that in all of the three cases, $\nexists a \in A$ such that a *pareto dominates* $f(P)$. Since the cases are exhaustive, we conclude that f is PE. \blacksquare

24.5 Examples of Social Choice Functions

24.5.1 Sequential Dictator

In the Sequential Dictator (SD) social choice function, we fix a pre-determined order over the agents. Then we move through the ordering giving each agent her most preferred share (i.e., p_i). If we run out of shares of the task, then one agent gets below his preferred share and the agents in the ordering that follow, get 0. If we have reached the last agent in the ordering, we give that agent whatever share of the task is left.

Claim 24.3 *SD is strategyproof.*

Proof: Let i be the first agent in the ordering who does not get her p_i , i.e. $s_i \neq p_i$. Consider two exhaustive cases:

Case 1 ($s_i < p_i$): Here, we must have that all agents after i in the ordering receive 0. Clearly, no one before i would want to change since they get their peaks. For i , reporting a peak that is larger than his true peak p_i is of no use since he cannot get above s_i . If he reports below true p_i he has a chance in the case when the report is below s_i where he gets something even less than s_i which is less preferred to him compared receiving s_i (because of single-peakedness). For all agents after i , they cannot get anything other than 0 since the shares of task are already exhausted. So, truthful reporting weakly dominates every other strategy for all agents in this case.

Case 2 ($s_i > p_i$): This is possible only when agent i is the last agent in the sequence. Note that in this case, the other $(n - 1)$ agents get their peaks so misreporting is strictly worse off for them. Agent i has no choice but to accept whatever is left, since his report does not change his share. Therefore, truthful reporting weakly dominates every other strategy in this case as well.

Therefore, SD is strategyproof. ■

Claim 24.4 *SD is Pareto Efficient.*

Proof: Let i be the first agent in the ordering who does not get her p_i , i.e. $s_i \neq p_i$. The $(i - 1)$ agents before i get their most preferred share and the ones after i get nothing (if i is not the last agent in the ordering). By the previous proof, we know that i gets more than his p_i iff i is the last agent in the ordering. In this case, decreasing i 's share would mean increasing at least one of the $i - 1$ agents' shares who currently have their peaks. All of those agents would be strictly worse off on this, so any such outcome cannot Pareto dominate SD(P). Now, in case i gets less than p_i , giving more to i (or anyone after i in the ordering) would mean taking away from the agents who have their peaks. They would be strictly worse off and therefore such outcome cannot Pareto dominate SD(P). Therefore, no outcome can Pareto dominate SD(P) and SD is Pareto Efficient. ■

The SD social choice function is SP and PE but still seems “unfair” since it favours certain agents over others (the ones that come earlier in the ordering). As the name implies, the rule gives a preference over “dictators”. We saw that in single-peaked preferences, a property called *anonymity* helped us increase the class of Strategy-proof functions beyond the Dictatorial bounds that we proved in the GS Theorem. Therefore, we define such a property here as well.

Definition 24.5 (ANON) *A Social Choice Function f is ANON if for all permutations $\sigma : N \rightarrow N$ and all preference profiles $P \in \mathcal{S}^n$, we have $f_i(P) = f_{\sigma(i)}(P^\sigma)$, where P^σ is the permutation of the P in accordance with σ [as in Lecture 22].*

Note that SD is SP and PE but not ANON.

24.5.2 Proportional

The proportional SCF gives every agent a constant factor c of their peaks. So if agent i 's peak is p_i , then agent i gets $c \cdot p_i$. We must have, $\sum_{i \in N} c \cdot p_i = 1 \implies c = \frac{1}{\sum_{i \in N} p_i}$.

Note that this SCF is clearly ANON since c is fixed once P is fixed and the agent's payoff depends purely on his peak, not his name or id. So if the preferences are permuted in the same way as the agents, the payoffs do now change.

Using the characterization theorem for PE (Theorem 24.2), we can show that this SCF is also PE. To see this consider three exhaustive cases. If $\sum_{i \in N} p_i = 1$ then $c = 1$ and everyone gets p_i . If $\sum_{i \in N} p_i < 1$, then $c > 1$ and everyone gets share $c \cdot p_i > p_i$. Similarly, if $\sum_{i \in N} p_i > 1$ then $c < 1$ and everyone gets share $c \cdot p_i < p_i$. Therefore, the three conditions of the theorem are satisfied and we conclude that the Proportional SCF is Pareto Efficient.

However, proportional SCF is not strategyproof. For instance consider the case when $\sum_{i \in N} p_i > 1$ so that $c < 1$. Everyone gets $c \cdot p_i < p_i$. If agent j *increases* her peak slightly, then the sum increases and c decreases. Suppose s was the initial sum so that $c = 1/s$. On increasing p_j by ϵ , the sum becomes $s' = s + \epsilon$. The initial payoff of agent j is $u_j = c p_j = p_j/s$ and her new payoff is $u'_j = c' p'_j = (p_j + \epsilon)/(s + \epsilon)$. The ratio

$$\frac{u'_j}{u_j} = \left(\frac{p_j + \epsilon}{p_j}\right) \left(\frac{s}{s + \epsilon}\right) = \frac{\left(1 + \frac{\epsilon}{p_j}\right)}{\left(1 + \frac{\epsilon}{s}\right)}.$$

Since $p_j < s$, we must have $\frac{1}{p_j} > \frac{1}{s}$ and for $\epsilon > 0$, $\frac{\epsilon}{p_j} > \frac{\epsilon}{s}$. Plugging this into the equation above, we get that $u'_j > u_j$. Since she was initially to the left of the peak, we can choose ϵ so that $u'_j \leq p_j$ and so misreporting with an increased peak can be strictly better for agent j . Thus, the proportional SCF is not strategyproof.

Lecture 25: October 11, 2017

Lecturer: Swaprava Nath

Scribe(s): Swaprava Nath

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

24.1 Recap

We discussed task allocation domain and the agent preferences were restricted to single-peaked over the share of the task. We wanted properties like Pareto efficiency, anonymity on top of strategyproofness in this domain. In this lecture, we will look at the uniform rule SCF introduced by Sprumont and study its properties.

24.2 Uniform rule SCF

Definition 24.1 (Uniform rule SCF) *In the task sharing domain, the uniform rule SCF $f : S^n \rightarrow A$ is defined as*

- $f_i^u(P) = p_i$, if $\sum_{i \in N} p_i = 1$.
- $f_i^u(P) = \max\{p_i, \mu(P)\}$, if $\sum_{i \in N} p_i < 1$.
- $f_i^u(P) = \min\{p_i, \lambda(P)\}$, if $\sum_{i \in N} p_i > 1$.


Where $\mu(P)$ solves

$$\sum_{i \in N} \max\{p_i, \mu(P)\} = 1,$$

and $\lambda(P)$ solves

$$\sum_{i \in N} \min\{p_i, \lambda(P)\} = 1.$$

Interpretation of μ : Consider the image in Figure 24.1. The vertical bars denote the total share of the task for every agent and the solid lines within the bars denote the peak of the preferences of the agents. Consider a horizontal thread that denotes the share of tasks to every agent. Let it start from the top (shown as dotted line in the figure). Initially it assigns 1 unit of task to every agent – which is infeasible. But now imagine that the thread moves downward (as shown in the figure) and stops whenever some agent's peak is reached – that agent is assigned her peak share and others shares are according to the level of the thread. If the allocation is still infeasible, lower the thread further for all *other* agents. Continue doing this until a feasible allocation is reached. This is exactly the solution given by μ . For visualizing the solution of λ , a very similar exercise can be done with threads starting from bottom and moving upwards.

Figure 24.1: Illustration of μ

Theorem 24.2 (Sprumont (1991)) *An uniform rule SCF is anonymous, Pareto efficient and strategyproof.*

Proof: The anonymity of this rule is obvious, since the rule uses only the peaks of every agent and not the agent identities. Hence a permutation of the agents' preferences will return an allocation where the allocation of the tasks of the agents has been correspondingly permuted.

To show that this is Pareto efficient, we need to verify the conditions that characterize PE as shown in the previous section. We see that

- for $\sum_{i \in N} p_i < 1$, $f_i^u(P) = \max\{p_i, \mu(P)\} \geq p_i$,
- for $\sum_{i \in N} p_i > 1$, $f_i^u(P) = \min\{p_i, \lambda(P)\} \leq p_i$, and
- for $\sum_{i \in N} p_i = 1$, $f_i^u(P) = p_i$.

Therefore the uniform rule satisfies PE.

For strategyproofness, we consider the following cases.

Case 1, $\sum_{i \in N} p_i = 1$: In this case, the allocation rule gives every agent her peak, hence it is strategyproof.

Case 2, $\sum_{i \in N} p_i < 1$: In this case, $f_i^u(P) = \max\{p_i, \mu(P)\} \geq p_i$. The only potential manipulable scenario is when $f_i^u(P) > p_i$, which implies that $\mu(P) > p_i$, i.e., the thread stopped before reaching the peak p_i from above. The only way this allocation can be changed is by reporting a $p'_i > \mu(P)$. Clearly, that takes the allocation of agent i further away from p_i than $\mu(P)$. Since the preferences are single peaked for the agent i over the allocation of task, she strictly prefers the current allocation $\mu(P)$ than p'_i .

Case 3, $\sum_{i \in N} p_i > 1$: the proof here is similar to that of Case 2, with the argument reversed. Hence, we have proved the theorem. ■

The converse of the above theorem is also true, but we skip its proof. Interested reader can refer to the paper by Sprumont [S91].

Theorem 24.3 *An SCF is strategyproof, Pareto efficient, and anonymous if and only if it is an uniform rule.*

24.3 Mechanism Design with Transfers

We now consider a new restricted domain of preferences that allows transferring utility, which is widely used in many real-world domains. The transferred utility is called ‘money’, and for this reason this sub-domain is also called *mechanisms with money*. In particular, we will only consider the transfers in quasi-linear form, as defined in the following section.

24.3.1 Quasi-linear Utility Model

In this setting the social choice function will be denoted by

$$F : \Theta \mapsto X.$$

Where Θ denotes the set of type profiles and X denotes the set of outcomes. We deliberately use the uppercase F to denote the SCF to distinguish this from what we have discussed so far. The types of the agents are their private information. The set of outcomes in this setting is a collection of two objects: (1) the **allocation**, and (2) the **payment**. Formally,

$$\begin{aligned} X &= \{x : x = (a, \pi)\} \\ a &\in A : \text{set of allocations, and} \\ \pi &= (\pi_1, \pi_2, \dots, \pi_n) \in \mathbb{R}^n \text{ the payments.} \end{aligned}$$

Examples of Allocations: A few examples are as follows.

Example 1: A public decision, $a \in A = \{\text{Bridge, Park, Theater, Museum, ...}\}$

Example 2: Private object allocation, e.g., a homogeneous divisible object (cake) denoted by the real interval $[0, 1]$. An allocation of this object gives every agent some portion of it. Hence, $a = (a_1, \dots, a_n)$, where $a_i \in [0, 1], \forall i \in N$ and $\sum_{i \in N} a_i \leq 1$.

Example 3: Single indivisible item allocation. Here the allocation vector is similar: $a = (a_1, \dots, a_n)$, however, $a_i \in \{0, 1\}, \forall i \in N$ and $\sum_{i \in N} a_i \leq 1$. This implies that the item can either go entirely to exactly one agent or may go to none.

Example 4: Allocation of multiple indivisible objects. Let the set of objects be denoted by S . An allocation in this setting is a partition of these objects into $n+1$ groups (A_0 denoting unallocated objects), i.e.,

$$A = \{(A_0, A_1, \dots, A_n) : A_i \subseteq S, \forall i \in N \cup \{0\}, \cup_{i \in N \cup \{0\}} A_i = S, A_i \cap A_j = \emptyset, \forall i \neq j\}.$$

The effect of the allocations are reflected in the valuation of every agent. Valuation v_i of agent i is a function of the allocation a and the *type* θ_i of every agent. This model of valuation is called **independent private values** (IPV).

$$v_i : A \times \Theta_i \mapsto \mathbb{R}.$$

It is called *private* since given the allocation, the valuation depends only on the private information of agent i . A more general model where it can depend on the entire vector of types given the allocation is called *interdependent* valuation.

In IPV setting, the notation of $v_i(a, \theta_i)$ is sometimes shortened to $\theta_i(a)$ or $v_i(a)$ to denote the same thing.

Payment function: This function is given by

$$p_i : \Theta \mapsto \mathbb{R}, \quad \mathbf{p} := (p_1, \dots, p_n).$$

Utility: The utility of an agent in this model is given by the following expression when the type profile is θ and the outcome is $(a, \boldsymbol{\pi})$.

$$u_i((a, \boldsymbol{\pi}), \theta_i) := v_i(a, \theta_i) - \pi_i \quad (\text{quasi-linear utilities})$$

The utility is linear in payment but can potentially be non-linear in the allocation. Since allocation and payment completely determine the outcome in this setting, this utility model is called quasi-linear.

Why is this a domain restriction? Here the outcome set X is given by $A \times \mathbb{R}$. Consider two alternatives (a, π) and (a, π') with $\pi_j = \pi'_j, \forall j \neq i$ and $\pi_i > \pi'_i$. Due to the quasi-linear preferences, for every type θ_i of agent i , $\theta_i(a) - \pi_i < \theta_i(a) - \pi'_i$. Hence, for every quasi-linear preference \succ_i^{QL} of i , $(a, \pi') \succ_i^{\text{QL}} (a, \pi)$. This implies that all possible orders over the alternatives is *not* admissible in this setting. Recall that the Gibbard-Satterthwaite setting demands that *all possible ordering over the outcomes must be admissible* in the domain. Only then the conclusion of dictatorship holds. Put in the current setting, it should mean that (a, π) *can be* more preferred than (a, π') by an agent. This is certainly not true in the quasi-linear setting – in particular, no agent can have a type where (a, π) is placed above (a, π') . This subtle domain restriction opens up the possibility of a lot of mechanisms to be strategyproof.

References

- [S91] Y. SPRUMONT, “The Division Problem with Single-Peaked Preferences: A Characterization of the Uniform Allocation Rule,” *Econometrica*, Vol. 59, No. 2 (Mar., 1991), pp. 509-519.

Lecture 26: October 13, 2017

Lecturer: Swaprava Nath

Scribe(s): Sudhir Kumar

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

26.1 Recap

In the last lecture we discussed the uniform rule SCF and proved that the uniform rule SCF is PE, ANON, and SP and mentioned that the converse is also true, i.e., a PE, ANON, and SP SCF must necessarily be an uniform rule. We also looked at another restricted domain of quasi-linear preferences. In this domain, an SCF $F : \Theta \mapsto X$, where X is the set of the tuples $(a, \boldsymbol{\pi})$ with a being the allocation and $\boldsymbol{\pi} := (\pi_1, \dots, \pi_n)$ being the vector of payments, can be decomposed as

$$F \equiv (f, \mathbf{p}).$$

Here the function $f : \Theta \mapsto A$ is called the *allocation function* and $\mathbf{p} := (p_i, i \in N \mid p_i : \Theta \mapsto \mathbb{R})$ the *payment function*.

We provide some examples of these functions to illustrate the domain.

26.2 Example of allocation functions

- Constant rule: allocation function is constant for all θ .

$$f^c(\theta) = a, \quad \text{for some } a \in A, \forall \theta \in \Theta.$$

- Dictatorial rule: the allocation that maximizes the valuation of a pre-determined player, whom we call the dictator.

$$f^D(\theta) \in \operatorname{argmax}_{a \in A} v_d(a, \theta_d), \text{ for some } d \in N.$$

- Allocatively efficient (AE) rule / Utilitarian rule: allocation that maximizes the *social welfare*

$$f(\theta) \in \operatorname{argmax}_{a \in A} \sum_{i \in N} v_i(a, \theta_i).$$

- Weighted efficient rule: a slight variant of the AE rule

$$f(\theta) \in \operatorname{argmax}_{a \in A} \sum_{i \in N} w_i v_i(a, \theta_i), w_i \geq 0, \forall i \in N, \text{ (not all } w_i \text{'s are zero).}$$

- Max-Min (also called egalitarian or Rawlsian) allocation rule: maximizes the minimum valuation of the agents

$$f^R(\theta) \in \operatorname{argmax}_{a \in A} \min_{i \in N} v_i(a, \theta_i).$$

- Affine maximizer rule: further generalization of the AE rule

$$f_{\text{AM}}(\theta) \in \operatorname{argmax}_{a \in A} \left(\sum_{i \in N} \lambda_i v_i(a, \theta_i) + \kappa(a) \right), \lambda_i \geq 0, \forall i \in N, \text{ (not all } \lambda_i \text{'s are zero)}.$$

26.3 Examples of payment rules

- Weak budget balanced (also called no-deficit or feasible) payment rule:

$$\sum_{i \in N} p_i(\theta) \geq 0, \forall \theta \in \Theta.$$

This payment rule makes sure that the mechanism designer does not need to add money to the system to run the mechanism. There could be a surplus, but never a shortage of total payment.

- No subsidy:

$$p_i(\theta) \geq 0, \forall i \in N, \forall \theta \in \Theta.$$

This payment rule makes sure that every agent is asked to pay in the mechanism.

- Budget balanced:

$$\sum_{i \in N} p_i(\theta) = 0.$$

This ensures that there is neither surplus nor deficit in the payments.

However, in the discussion that follows, we will be prioritizing the allocation rule, and will put less restriction on the payment rule. In particular, we will be interested in finding an allocation that is “truthful” and may put no restriction on payments.

26.4 Incentive compatibility

To denote *truthfulness* of a mechanism, we need to distinguish the true type with the reported type. Denote the reported type with $\hat{\theta} = (\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_n)$. Recall that a direct mechanism is denoted by $((f, \mathbf{p}), \Theta)$, and the allocation and payment function are evaluated over the *reported* types. The utility of agent i is given by

$$v_i(f(\hat{\theta}), \theta_i) - p_i(\theta).$$

Definition 26.1 (Dominant Strategy Incentive Compatibility (DSIC)) A direct mechanism (f, \mathbf{p}) is dominant strategy incentive compatible (DSIC) if,

$$v_i(f(\theta_i, \theta_{-i}), \theta_i) - p(\theta_i, \theta_{-i}) \geq v_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i) - p(\hat{\theta}_i, \theta_{-i}) \\ \forall \theta_i, \hat{\theta}_i \in \Theta_i, \forall \theta_{-i} \in \Theta_{-i}, \forall i \in N.$$

Note: the definition of DSIC implies that reporting types truthfully is a WDSE. If the above mentioned conditions hold then we say that \mathbf{p} implements f in dominant strategies.

Let us illustrate the conditions given by the definition above for a case with two agents and two types.

Example 26.1 Consider $N = \{1, 2\}$, $\Theta_1 = \Theta_2 = \{\theta^H, \theta^L\}$. The allocation rule is $f : \Theta_1 \times \Theta_2 \mapsto A$. If \mathbf{p} implements f in dominant strategies, then for player 1, the conditions are

$$\begin{aligned} v_1(f(\theta_H, \theta_2), \theta_H) - p(\theta_H, \theta_2) &\geq v_1(f(\theta_L, \theta_2), \theta_H) - p(\theta_L, \theta_2), \forall \theta_2 \in \Theta_2, \\ v_1(f(\theta_L, \theta_2), \theta_L) - p(\theta_L, \theta_2) &\geq v_1(f(\theta_H, \theta_2), \theta_L) - p(\theta_H, \theta_2), \forall \theta_2 \in \Theta_2. \end{aligned}$$

Similarly for player 2

$$\begin{aligned} v_2(f(\theta_H, \theta_1), \theta_H) - p(\theta_H, \theta_1) &\geq v_2(f(\theta_L, \theta_1), \theta_H) - p(\theta_L, \theta_1), \forall \theta_1 \in \Theta_1, \\ v_2(f(\theta_H, \theta_1), \theta_H) - p(\theta_H, \theta_1) &\geq v_2(f(\theta_L, \theta_1), \theta_H) - p(\theta_L, \theta_1), \forall \theta_1 \in \Theta_1. \end{aligned}$$

26.5 Impact of DSIC on payments

1. Effect of an additive function to payment: Let (f, \mathbf{p}) is DSIC. We have

$$\begin{aligned} v_i(f(\theta_i, \theta_{-i}), \theta_i) - p(\theta_i, \theta_{-i}) &\geq v_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i) - p(\hat{\theta}_i, \theta_{-i}) \\ &\quad \forall \theta_i, \hat{\theta}_i \in \Theta_i, \forall \theta_{-i} \in \Theta_{-i}, \forall i \in N. \end{aligned}$$

Now consider

$$q(\theta_i, \theta_{-i}) = p(\theta_i, \theta_{-i}) + h_i(\theta_{-i}).$$

Question: Is (f, \mathbf{q}) DSIC?

Answer: Yes it is. Consider

$$\begin{aligned} &v_i(f(\theta_i, \theta_{-i}), \theta_i) - q(\theta_i, \theta_{-i}) \\ &= v_i(f(\theta_i, \theta_{-i}), \theta_i) - p(\theta_i, \theta_{-i}) - h_i(\theta_{-i}) \\ &\geq v_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i) - p(\hat{\theta}_i, \theta_{-i}) - h_i(\theta_{-i}) \\ &= v_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i) - q(\hat{\theta}_i, \theta_{-i}). \end{aligned}$$

The first inequality holds because (f, \mathbf{p}) is DSIC and the equalities hold from the definition of q_i 's.

2. If the allocation is same for two different types of agent i : Consider types θ_i and $\hat{\theta}_i$ of agent i and assume that for θ_{-i}

$$f(\theta_i, \theta_{-i}) = f(\hat{\theta}_i, \theta_{-i}) = a \text{ (say)}$$

Suppose (f, \mathbf{p}) is DSIC. When agent i 's true type is θ_i , we have

$$\begin{aligned} v_i(f(\theta_i, \theta_{-i}), \theta_i) - p(\theta_i, \theta_{-i}) &\geq v_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i) - p(\hat{\theta}_i, \theta_{-i}) \\ p_i(\hat{\theta}_i, \theta_{-i}) &\geq p_i(\theta_i, \theta_{-i}). \quad (\text{since } f(\theta_i, \theta_{-i}) = f(\hat{\theta}_i, \theta_{-i}).) \end{aligned}$$

When agent i 's true type is $\hat{\theta}_i$, similarly we have

$$\begin{aligned} v_i(f(\hat{\theta}_i, \theta_{-i}), \hat{\theta}_i) - p(\hat{\theta}_i, \theta_{-i}) &\geq v_i(f(\theta_i, \theta_{-i}), \hat{\theta}_i) - p(\theta_i, \theta_{-i}) \\ p_i(\theta_i, \theta_{-i}) &\geq p_i(\hat{\theta}_i, \theta_{-i}). \end{aligned}$$

Combining the above two observations we get

$$p_i(\theta_i, \theta_{-i}) = p_i(\hat{\theta}_i, \theta_{-i}).$$

Hence, if the allocation does not change by an unilateral deviation of type of an agent, the payment also does not change.

Lecture 27: October 17, 2017

Lecturer: Swaprava Nath

Scribe(s): Ameya Loya

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

27.1 Pareto Optimality

Definition 27.1.1 A direct mechanism $(f, (p_1, p_2, \dots, p_n))$ is pareto optimal (PO) if at every type profile $\theta \in \Theta$, \nexists an allocation $b \in A$ and a payment vector $(\pi_1, \pi_2, \dots, \pi_n)$ with $\sum_{i \in N} \pi_i \geq \sum_{i \in N} p_i(\theta)$ such that,

$$\begin{aligned} v_i(b, \theta_i) - \pi_i &\geq v_i(f(\theta), \theta_i) - p_i(\theta) && \text{for all } i \in N, \text{ and,} \\ v_j(b, \theta_j) - \pi_j &> v_j(f(\theta), \theta_j) - p_j(\theta) && \text{for some } j \in N. \end{aligned}$$

Hence for a Pareto optimal mechanism, agents' payoffs are maximal at every type profile. Improving the payoff of one agent will result in the reduction of the payoff some other agent.

27.2 Relation between Pareto Optimality and Allocative Efficiency in Quasi-linear Domain

Recall that an allocation function is allocatively efficient (AE) if it maximizes the social welfare, i.e.,

$$f^{AE}(\theta) \in \arg \max_{a \in A} \sum_{i \in N} v_i(a, \theta_i)$$

Theorem 27.1 A mechanism (f, p) is Pareto optimal iff it is allocatively efficient.

Proof: (\Rightarrow) We first prove that if a mechanism is pareto optimal then it is AE. To do so we show that !AE \implies !PO. Since f is not AE $\exists b \in A$ s.t .

$$\sum_{i \in N} v_i(b, \theta_i) > \sum_{i \in N} v_i(f(\theta), \theta_i), \text{ for some } \theta \in \Theta.$$

Define

$$\delta := \sum_{i \in N} v_i(b, \theta_i) - \sum_{i \in N} v_i(f(\theta), \theta_i).$$

By definition, $\delta > 0$. Now define a payment for every $i \in N$

$$\begin{aligned} \pi_i &= v_i(b, \theta_i) - v_i(f(\theta), \theta_i) + p_i(\theta) - \delta/n \\ \Rightarrow (v_i(b, \theta_i) - \pi_i) - (v_i(f(\theta), \theta_i) - p_i(\theta)) &= \delta/n > 0 \end{aligned}$$

Hence the new allocation and payment yields more payoff to every agent. Moreover, we get $\sum_{i \in N} \pi_i = \sum_{i \in N} p_i(\theta)$. Hence, (f, p) is not PO.

(\Leftarrow) To prove the converse, we show that !PO \implies !AE. If (f, p) is not PO, $\exists b, \pi, \theta$ such that the following holds.

$$\begin{aligned} \sum_{i \in N} \pi_i &\geq \sum_{i \in N} p_i(\theta) \\ v_i(b, \theta_i) - \pi_i &\geq v_i(f(\theta), \theta_i) - p_i(\theta), \forall i \in N, \text{ and,} \\ v_j(b, \theta_j) - \pi_j &> v_j(f(\theta), \theta_j) - p_j(\theta), \text{ for some } j \in N. \end{aligned}$$

Summing over the last two inequalities over all $i \in N$, we get

$$\begin{aligned} &\sum_{i \in N} v_i(b, \theta_i) - \sum_{i \in N} \pi_i > \sum_{i \in N} v_i(f(\theta), \theta_i) - \sum_{i \in N} p_i \\ \Rightarrow &\sum_{i \in N} v_i(b, \theta_i) - \sum_{i \in N} v_i(f(\theta), \theta_i) > \sum_{i \in N} \pi_i - \sum_{i \in N} p_i \geq 0. \end{aligned}$$

Which proves that f is not AE. ■

27.3 Implementability of Allocation Rules

We call an allocation rule $f : \Theta \rightarrow A$ *implementable* if $\exists p$ such that (f, p) is DSIC. We show that the efficient rule is implementable. It is implemented by a class of payments known as the Groves class of payments.

Definition 27.3.1 (Groves class of payments) Let $h_i : \Theta_{-i} \rightarrow \mathbb{R}$ be an arbitrary function for every $i \in N$. The Groves class of payments is defined as the payment rules defined as

$$p_i^G(\theta) = h_i(\theta_{-i}) - \sum_{j \neq i} v_j(f^{AE}(\theta), \theta_j).$$

The mechanism (f^{AE}, p^G) is called the *Groves class of mechanisms*.

Definition 27.3.1 Consider four agents and one indivisible item is to be allocated. Value of the agents are 10, 8, 6 and 4 respectively when they receive the item and zero otherwise. Suppose

$$h_i(\theta_{-i}) = 10 \text{ for all } \theta_{-i} \text{ for all } i \in N.$$

Clearly, the efficient allocation is to give the item to agent 1. Since the h_i function is a constant, and the second term of the Groves payment is zero for agent 1 and 10 for every other agent, the payments are 10, 0, 0 and 0 respectively. Therefore this payment rule charges 10 to the winning agent and zero to others.

However, this class also admits very surprising payments. For example, if we consider $h_i(\theta_{-i}) = \sum_{j \neq i} \frac{\theta_j}{2}$, one can find that the Groves payments will be 9, 0, 1, 2 respectively. These are surprising since the agents who do not receive the item are also asked to pay. Though the payments are surprising enough, they satisfy one very important property.

Theorem 27.2 Groves class of mechanisms is DSIC.

Proof: Suppose for agent i , the true type is θ_i and the reported type is $\hat{\theta}_i$. Also, assume that

$$f^{AE}(\theta_i, \theta_{-i}) = a, \text{ and } f^{AE}(\hat{\theta}_i, \theta_{-i}) = b.$$

Consider the utility of agent i when he reports θ_i

$$\begin{aligned} & v_i(f^{AE}(\theta), \theta_i) - p_i^G(\theta) \\ &= v_i(f^{AE}(\theta), \theta_i) - h_i(\theta_{-i}) + \sum_{j \neq i} v_j(f^{AE}(\theta), \theta_j) \\ &= \sum_{j \in N} v_j(f^{AE}(\theta), \theta_j) - h_i(\theta_{-i}) \\ &\geq \sum_{j \in N} v_j(b, \theta_j) - h_i(\theta_{-i}) \\ &= v_i(f^{AE}(\hat{\theta}_i, \theta_{-i}), \theta_i) - h_i(\theta_{-i}) + \sum_{j \neq i} v_j(f^{AE}(\hat{\theta}_i, \theta_{-i}), \theta_j) \\ &= v_i(f^{AE}(\hat{\theta}_i, \theta_{-i}), \theta_i) - p_i^G(\hat{\theta}_i, \theta_{-i}). \end{aligned}$$

Where the inequality follows from the the definition of f^{AE} . Hence we have proved the theorem. ■

27.4 The Vickrey-Clarke-Groves (VCG) Mechanism

An interesting mechanism in the Groves class is the VCG mechanism, named after Vickrey, Clarke, and Groves. This is also called the *pivotal* mechanism. We will discuss later about its pivotality. It is characterized by a specific $h_i(\theta_{-i})$ given as follows.

$$h_i(\theta_{-i}) = \max_{b \in A} \sum_{j \neq i} v_j(b, \theta_j),$$

and hence

$$p_i(\theta) = \max_{b \in A} \left[\sum_{j \neq i} v_j(b, \theta_j) \right] - \sum_{j \neq i} v_j(f^{AE}(\theta), \theta_j).$$

Note that, if the set of allocations A remains unchanged after removing agent i , e.g., in the public good allocation problem, then the payment above is always ≥ 0 . This shows that VCG mechanism for public goods gives *no subsidy* to any agent. As a consequence, it is obvious that in such a setting, the VCG mechanism is also *weakly budget balanced* (no-deficit), since $\sum_{i \in N} p_i(\theta) \geq 0$.

Lecture 28: October 18, 2017

Lecturer: Swaprava Nath

Scribe(s): Dhawal Upadhyay

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

28.1 Revisiting VCG

In VCG mechanism, the allocation and payment rules are given as follows.

$$f^{AE}(\theta) \in \arg \max_{a \in A} \sum_{i \in N} v_i(a, \theta_i)$$

$$p_i^{VCG}(\theta) = \max_{b \in A} \sum_{j \neq i} v_j(b, \theta_j) - \sum_{j \neq i} v_j(f^{AE}(\theta), \theta_j)$$

Here f^{AE} is the allocation rule that maximizes the sum of the valuations (also called social welfare) of all the players. We call this rule *allocatively efficient* or *just efficient*. Consider the utility of player i

$$\begin{aligned} & v_i(f^{AE}(\theta), \theta_i) - p_i^{VCG}(\theta) \\ &= v_i(f^{AE}(\theta), \theta_i) - \max_{b \in A} \sum_{j \neq i} v_j(b, \theta_j) + \sum_{j \neq i} v_j(f^{AE}(\theta), \theta_j) \\ &= \sum_{j \in N} v_j(f^{AE}(\theta), \theta_j) - \max_{b \in A} \sum_{j \neq i} v_j(b, \theta_j) \end{aligned}$$

The first term on the RHS of the last equality is the maximum social welfare. The second term is the maximum social welfare when i is absent. The difference gives us the marginal contribution of i in the social welfare. This is another way of interpreting the utility of an agent in the VCG mechanism.

28.2 Illustration of VCG payments

- **Single indivisible object allocation:** Consider an agent $i \in N$. Every agent has some value when the object is assigned to him, and zero otherwise. Efficiency requires the object must go to the agent who values it the most. So if i 's bid is maximum, he is assigned the object. The payment is given by

$$p_i^{VCG}(\theta) = \max_{a \in A} \sum_{j \neq i} v_j(a, \theta_j) - \sum_{j \neq i} v_j(f^{AE}(\theta), \theta_j).$$

If agent i 's bid is not the highest, i.e., she does not win the object, both the first and second term on the RHS gives the same number since that turns out to be the bid of the winning agent. Hence, agent i pays nothing.

If agent i 's bid is the highest, i.e., she wins the object, the first term on the RHS becomes the second highest bid, and the second term becomes zero since agents except i do not get the object in the efficient allocation. Thus the payment made by the highest bidder is equal to the second highest bid,

and no other agent pays anything in this mechanism. This is precisely the second price auction that we have seen earlier. So, in the case of single object allocation, VCG mechanism is exactly the second price auction.

- **Public project allocation:** In this case, the first term of payment function is the social welfare of all agents except agent i if i were not present. The second term is the social welfare of other agents in i 's presence. This gives another interpretation of the VCG payment: the loss in social welfare of other agents because of agent i 's presence – and agent i is asked to compensate this loss. Consider the following example.

	Football	Library	Museum
A	0	70	50
B	95	10	50
C	10	50	50

The allocation set is $A = \{F, L, M\}$.

Efficient allocation: M .

A pays: $105 - 100 = 5$.

Reason: without A's presence, best allocation would've been F, with total utility of other agents = 105. In A's presence, utility of other agents = $50 + 50 = 100$. Difference = 5.

Similarly, we can find that

B pays: $120 - 100 = 20$

C pays: $100 - 100 = 0$

Observation: payment values of VCG are always positive (irrespective of utility function) – first term is always at least as large as second term.

We also notice that, in the above example, only those agents are charged a positive payment whose presence *changes the outcome*. These agents are called *pivotal* agents. A non-pivotal agent pays zero under VCG. This is why VCG is also called *pivotal* mechanism.

- **Combinatorial Allocation:** Sale of multiple objects. Consider two objects – 1 and 2. The valuations over all possible combinations of these objects are given as

	ϕ	$\{1\}$	$\{2\}$	$\{1,2\}$
v_1	0	8	6	12
v_2	0	9	4	14

Efficient allocation - $\{1\}$ goes to player 2, and $\{2\}$ goes to player 1. Here, type is the value itself, so we can also write $v_i(a, \theta_i) = \theta_i(a)$.

$$p_1^{VCG}(\theta_1, \theta_2) = \max_{a \in A} \sum_{j \neq 1} \theta_j(a) - \sum_{j \neq 1} \theta_j(f(v)),$$

which gives $= 14 - 9 = 5$; Payoff = $6 - 5 = 1$.

$p_2^{VCG}(\theta_1, \theta_2) = 12 - 6 = 6$; Payoff = $9 - 6 = 3$.

Hence every agent gets a non-negative payoff – something we will discuss more later.

It is also instructive to verify if the agents tried to overbid their valuations to obtain both the objects, what payoff would they receive.

28.3 VCG mechanism in combinatorial auction

VCG mechanism has several useful properties:

1. It is DSIC.
2. It is efficient.
3. Payments are non-negative (hence no subsidy).

Combinatorial auction notation:

- $M = \{1, 2, \dots, m\}$: set of objects.
- Set of bundles: $\Omega = \{S : S \subseteq M\}$.
- Type of agent i is $\theta_i : \Omega \mapsto \mathbb{R}$.

Hence $\theta_i(S)$, $S \in \Omega$ is the value of agent i for the bundle S . We assume $\theta_i(S) \geq 0 \forall S \in \Omega$, which means that every item is a ‘good’ (an item that gives negative valuation is called a ‘bad’).

An allocation of objects is given by $X = \{X_0, X_1, \dots, X_n\}$, $X_i \in \Omega$, $X_i \cap X_j = \emptyset$ if $i \neq j$, and $\cup_{i=0}^n X_i = M$. The set of allocations A is the collection of such X s. X_0 is the set of unallocated objects, X_i is the bundle allocated to i . Assume $\theta_i(\emptyset) = 0$.

We also assume that the valuations have no externalities (selfish valuations), ie $v_i(X, \theta_i) = \theta_i(X_i)$ (does not depend on the bundles of other agents).

Claim 28.1 *The payment for an agent who gets no object in the allocation is zero.*

Proof: Say agent i gets no object in the efficient allocation, ie $X \in \arg \max_{x \in A} \sum_{i \in N} v_i(x, \theta_i)$, and $X_i = \emptyset$.

VCG payment’s first term considers allocation excluding agent i . Define $Y \in \arg \max_{y \in A} \sum_{j \neq i} v_j(y, \theta_j)$.

We have seen that if the allocation set remains unchanged, VCG payment is always non-negative, i.e., $p_i^{VCG} \geq 0$ (no-subsidy condition).

$$\begin{aligned}
 p_i^{VCG}(\theta) &= \sum_{j \neq i} v_j(Y, \theta_j) - \sum_{j \neq i} v_j(X, \theta_j) \\
 &\leq \sum_{j \in N} v_j(Y, \theta_j) - \sum_{j \in N} v_j(X, \theta_j) \\
 &\leq 0.
 \end{aligned}$$

The first inequality holds since we are adding $v_i(Y, \theta_i) \geq 0$ and subtracting $v_i(X, \theta_i) = 0$ – together a non-negative quantity. The second inequality holds since X maximizes the sum of the valuations by definition. Hence combining this with the no-subsidy condition, we conclude that $p_i^{VCG}(\theta) = 0$. ■

Lecture 29: October 20, 2017

Lecturer: Swaprava Nath

Scribe(s): Anil Kumar Gupta

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

29.1 Recap

In the last lecture, we studied the VCG mechanism. We proved a claim that the payment for an agent that is not allocated any object in a combinatorial allocation is zero. We continue the analysis of VCG and make another claim.

29.2 VCG mechanism continued

Definition 29.1 (Individual Rationality) A direct mechanism (f, \mathbf{p}) is individually rational (IR) if for every $\theta \in \Theta$

$$v_i(f(\theta), \theta_i) - p_i(\theta) \geq 0, \quad \forall i \in N.$$

This property ensures that by participating in the mechanism, no agent gets a negative utility. Hence a rational agent should voluntarily participate in the game. We now show that in the setting we consider, VCG satisfies this condition.

Claim 29.2 In the allocation of goods, VCG mechanism is individually rational.

Proof: Recall the definitions of X and Y as follows.

$$X \in \arg \max_{a \in A} \sum_{i \in N} v_i(a, \theta_i), \quad Y \in \arg \max_{a \in A} \sum_{j \neq i} v_j(a, \theta_j).$$

The payoff of agent $i \in N$ according to VCG mechanism is given by

$$\begin{aligned} & v_i(X, \theta_i) - \sum_{j \neq i} v_j(Y, \theta_j) + \sum_{j \neq i} v_j(X, \theta_j) \\ &= \sum_{j \in N} v_j(X, \theta_j) - \sum_{j \neq i} v_j(Y, \theta_j) \\ &= \underbrace{\left(\sum_{j \in N} v_j(X, \theta_j) - \sum_{j \in N} v_j(Y, \theta_j) \right)}_{\geq 0} + \underbrace{v_i(Y, \theta_i)}_{\geq 0} \geq 0 \end{aligned}$$

We get the second equality by adding and subtracting $v_i(Y, \theta_i)$. The first term in the last line is non-negative by definition of X and the second term is non-negative due to all allocations are of goods. This concludes the proof. \blacksquare

29.3 Application Domain : Internet Advertising

Internet advertising is defined as the method of delivering promotional marketing messages called ‘ads’ to consumers when a user searches something on a search engine or visits a website. In this lecture, we will consider the methods of *efficiently* placing an ad and deciding the payment to the advertiser.

29.3.1 Advantages of Internet Advertising

Internet advertising is much better than the conventional advertising in newspapers/articles or radio because of the following advantages.

1. **User Data:** Using Internet ads, the advertiser can gain information about which set of buyers are interested in which set of ads. This enables the advertisers to target certain products to the interested users.
2. **Measurable Actions:** Classification of buyers is possible using Internet advertising. The buyers can be classified into a set of groups and be shown ads according to their interest.
3. **Low Latency:** Internet advertising allows the auctions to happen on the fly. The auction takes place just before the ad is to be shown. For example in search engines like Google, bidding takes place after an user enters a set of keywords to search and ads for the organizations that win the auction are shown. This also allows the advertisers to enter any time and enables automation of the whole process.

29.3.2 Types of Ads on Internet

1. Sponsored Search Ads

In this type of ads, the advertisers bid on the keywords that are entered by the user while searching on a search engine. Some details of this method are as follows.

- Some words can have multiple meanings based on the context. Like ‘jaguar’ can mean a car manufacturer or an animal. So to prevent the car ads when the animal jaguar is searched, stop words are used. Stop words are a set of words which if used with a particular word (here, jaguar), the ad (for car jaguar) is not presented.
- One user can search for the same keyword multiple times. So to prevent activity of a malicious user, or to prevent draining of money from the advertiser, or to prevent irritating the user with same ad every time, a cap is set on the number of times an ad can be shown to a user.

2. Contextual Ads

These type of ads are presented to a user based on the context of a page. For example, gmail can present certain type of ads based on the content of the email.

3. Display Ads

It is the classical way of displaying the ads just like the banner ads in the newspapers.

29.3.3 Position Auction

This type of auction is used to sell multiple ad slots in a webpage. Let $N = \{1, 2, \dots, n\}$ denote the set of bidders and $M = \{1, 2, \dots, m\}$ be set of slots available to be auctioned. For simplicity we assume $m \geq n$, i.e.,

every advertiser gets a slot – the question is who gets which slot. Here among the set of slots, 1 represents the ‘best’ position and m represents the ‘worst’ position.

Advertiser value: It is the valuation received by the advertiser from her ad on the Internet. We make the following assumptions in this setting.

1. Only clicks generate the value to the bidders.
2. All the clicks are valued equally no matter which position it came from. This means that the role of position is only in the probability of being clicked. This assumption helps in decoupling the position effect from the value effect.

The expected value of agent i , when her ad is shown at position j is

$$v_{ij} = CTR_{ij} \cdot v_i.$$

Where CTR_{ij} is the *click through rate* (CTR) that denotes the probability that bidder i 's ad gets clicked at position j . This quantity CTR_{ij} is again assumed to be decomposable into two terms: (1) the quality of the ad of the advertiser i , denoted by $CTR_i \in [0, 1]$, that reflects how good the ad is and (2) the effect of position on the CTR, denoted by $pos_j \in [0, 1]$. Therefore

$$CTR_{ij} = CTR_i \cdot pos_j, \quad CTR_i \in [0, 1], pos_j \in [0, 1].$$

Thereby the value of agent i when her ad is in position j is decomposable into the position effect and agent effect: $v_{ij} = pos_j \cdot (CTR_i \cdot v_i)$. Note that both pos_j and CTR_i are measurable and can be estimated by the auctioneer, but v_i is a private information (type) of the agent.

29.3.4 Deciding the Auction Mechanism

The auction mechanism decides the allocation of resources (e.g., the position at which an ad is displayed) and the payment to be made by each advertiser. The following are some standard methods of auctions.

1. **Early position auctions:** In this type of auction, the received bids are sorted in decreasing order of value and the positions are allotted according to the this order – the highest bidder receiving the first position and lowest bidder receiving last position – and they pay some fixed amount according to their position upfront. This type of bidding puts all risk on the advertiser – now an advertiser has to estimate the probability of getting a click since they are paying even before the ad is clicked.
2. **Pay per click model:** In this mechanism, the price charged to the advertiser is proportional to the number of clicks the users made on the ad. This mechanism puts all the risk on the website because a there could be an ad that is at one of the top positions but never gets clicked and hence does produce any revenue.
3. **Rank by expected revenue:** Modern sponsored search auctions uses a middle ground between these two extreme approaches. In this approach, the advertisers are ranked by product of *estimated* CTR and the bid of the advertiser. This mechanism shares the risk between the advertiser and the website.

We use the following notations :

- $eCTR_i$: the value of the click through rate of the Ad i estimated by the search engine
- b_i : the amount that agent i is willing to pay if a click occurs

- $x_i \in M$: the position assigned to agent i s.t. $x_i \neq x_k, i, k \in N$
- $x = (x_1, x_2, \dots, x_n)$: allocation of ads to positions

Here we assume that the ads are allocated according to the decreasing order of their expected revenue, i.e., according to the decreasing order of $eCTR_i \cdot b_i$. Suppose the agent i reports his bid as b_i . Then

$$\hat{v}_i(x) = pos_{x_i}(eCTR_i \cdot b_i)$$

is the reported value of the agent for the allocation x .

Allocation rule: Standard search engines and advertisement websites pick the allocation x^* such that it maximizes the sum of the valuations of the advertisers, i.e.,

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

This method of allocation is exactly same as the VCG allocation.

Lecture 30: October 24, 2017

Lecturer: Swaprava Nath

Scribe(s): Kunal Chaturvedi

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava.cse.iitk.ac.in.

30.1 Recap

In the previous lecture, we proved that VCG is individually rational for allocation of goods. We also looked at internet advertising and position auctions to sell multiple ads on a webpage. The winner determination problem for position auction is restated for the purpose of this lecture.

30.2 Winner Determination Problem (continued)

Definition 30.1 (Winner Determination Problem (WDP)) The optimal allocation x^* of the slots should be determined such that the sum of the reported valuations of the agents is maximized, i.e.,

$$x^* = \arg \max_x \sum_{i \in N} \hat{v}_i(x). \quad (30.1)$$

We now prove a result on the winner determination problem.

Theorem 30.2 If an allocation solves the winner determination problem, then it must be a rank-by-revenue mechanism.

Recall that a rank-by-revenue assigns the slots according to the decreasing order of $eCTR_i \cdot b_i$, where $eCTR_i$ is the estimated quality of agent i 's ad and b_i is his bid.

Proof: Assume for contradiction that allocation x is optimal, i.e., solves the WDP (Eq. 30.1) and bids of agents 1 and 2 are such that, $eCTR_1 \cdot b_1 > eCTR_2 \cdot b_2$, but, 2 is placed above 1. WLOG assume $x_2 = 1$ and $x_1 = 2$. Consider a different allocation x' with all the agents except 1 and 2 getting the same position, and $x'_1 = 1$ and $x'_2 = 2$. Therefore the social welfare is given by

$$\sum_{i \in N} \hat{v}_i(x') = pos_1(eCTR_1 \cdot b_1) + pos_2(eCTR_2 \cdot b_2) + \sum_{k \neq 1,2} pos_{x_k}(eCTR_k \cdot b_k) \quad (30.2)$$

$$\sum_{i \in N} \hat{v}_i(x) = pos_2(eCTR_1 \cdot b_1) + pos_1(eCTR_2 \cdot b_2) + \sum_{k \neq 1,2} pos_{x_k}(eCTR_k \cdot b_k) \quad (30.3)$$

Subtracting equation 30.3 from 30.2, we get

$$\sum_{i \in N} \hat{v}_i(x') - \sum_{i \in N} \hat{v}_i(x) = (pos_1 - pos_2)(eCTR_1 \cdot b_1 - eCTR_2 \cdot b_2).$$

As both factors in the RHS are positive, $\sum_{i \in N} \hat{v}_i(x') - \sum_{i \in N} \hat{v}_i(x) > 0$, which is a contradiction to x being optimal. This concludes the proof of the theorem. ■

Note: An advantage to be noted here is that the winner determination problem is poly-time.

30.3 VCG in Position Auction

After picking the efficient allocation (which is the solution of the WDP), we need payments to implement it in DSIC. The natural candidate for this is VCG payment. VCG payment is used by Twitter, Facebook for sponsored ads.

Given the bids, (b_1, b_2, \dots, b_n) , ordered WLOG such that

$$eCTR_1 \cdot b_1 \geq eCTR_2 \cdot b_2 \geq \dots \geq eCTR_n \cdot b_n.$$

Hence, the efficient allocation x^* is such that $x_i^* = i$ for $i = 1, 2, \dots, n$. The payment used is VCG.

Define $x_{-i}^* \in \arg \max_x \sum_{j \neq i} \hat{v}_j(x)$. Now the payment of agent i according to VCG is given by

$$\begin{aligned} p_i^{VCG}(b) &= \sum_{j \neq i} \hat{v}_j(x_{-i}^*) - \sum_{j \neq i} \hat{v}_j(x^*) \\ &= \sum_{j=i}^{n-1} pos_j(eCTR_{j+1} \cdot b_{j+1}) - \sum_{j=i}^{n-1} pos_{j+1}(eCTR_{j+1} \cdot b_{j+1}) \end{aligned}$$

Therefore

$$p_i^{VCG}(b) = \begin{cases} \sum_{j=i}^{n-1} (pos_j - pos_{j+1})(eCTR_{j+1} \cdot b_{j+1}) & i = 1, 2, \dots, n-1 \\ 0 & i = n \end{cases}$$

This gives us the total expected payment. To convert to payment per click, we need to normalize with $pos_i \cdot eCTR_i$. So payment per click = $\frac{1}{pos_i \cdot eCTR_i} p_i^{VCG}(b)$.

Observation: allocation w.r.t. rank-by-revenue and payment w.r.t. VCG is DSIC if the $eCTR$'s are accurate.

30.4 Generalized Second Price (GSP) in Position Auction

Generalized second price auction is a mechanism used by certain search engines e.g., Google, Bing etc. The allocation rule is same as rank-by-revenue, but the payment is such that every agent pays her next highest expected bid. This payment method is simple and easy to explain to advertisers, and has similarities with second price auction but also has serious limitations. Consider the following example

Expected revenue before position effect	v_i	$eCTR_i$	pos_j
2	10	0.2	1
4	8	0.5	0.2
4.2	6	0.7	0.1

Both VCG and GSP allocate the slots to solve the winner determination problem (i.e., rank-by-revenue allocation) with slot 1 going to player 3, slot 2 to player 2 and slot 3 to player 1.

For VCG, payment of slot 1 (for player 3) is $p_i^{VCG}(b) = \sum_{j=i}^{n-1} (p_j - p_{j+1})(eCTR_{j+1} \cdot b_{j+1}) = 0.8 \times 4 + 0.1 \times 2 = 3.4$.

Similarly payment for slot 2 (player 2) = $0.1 \times 2 = 0.2$. Payment for slot 3 (player 1) is equal to 0.

Utility of slot 1 (player 3) = $4.2 \times 1 - 3.4 = 0.8$.

Utility of slot 2 (player 2) = $4 \times 0.2 - 0.2 = 0.6$.

Utility of slot 3 (player 1) = $0.1 \times 2 + 0 = 0.2$.

For GSP:

Payment for slot 1 (player 3) = $4 \times 0.2 = 0.8$. Payment for slot 2 (player 2) = $2 \times 0.1 = 0.2$. Utility for Slot 1 = $4.2 \times 1 - 0.8 = 3.4$. Utility for Slot 2 = $4 \times 0.2 - 0.2 = 0.6$.

An interesting case to consider is if player 2 overbids. Say player 2 bids 8.6 to change the allocation. So $b_2 = 8.6$. In this case, the payment under VCG becomes = $0.8 \times 4.2 + 0.1 \times 2 = 3.56$. The utility of player 2 (now assigned slot 1) = $4 - 3.56 = 0.44$. We see that utility has decreased for player 2. This is expected because VCG is DSIC.

Under GSP, payment of player 2 = $4.2 \times 0.2 = 0.84$. Utility of player 2 = $4 \times 1 - 0.84$, which is more than that when she reports truthfully. This shows that GSP is not DSIC.

30.5 Desirable properties of VCG

1. VCG is DSIC. Hence there is very low cognitive load on bidders.
2. VCG never runs into deficits (in some settings). It charges the marginal contribution to the other agents as payments and gets its own marginal contribution as payoff.
3. VCG never charges a losing agent.
4. It is individually rational for every agent to participate.

Lecture 31: October 25, 2017

Lecturer: Swaprava Nath

Scribe(s): Garima Shakya

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

31.1 Recap

In the previous lecture, we discussed the advantages of VCG mechanism such as it is DSIC, never runs into deficit, never charges a losing agent and is individually rational for agents to participate. It is the most used mechanism with money, but it is good to know the limitations of VCG.

31.2 Criticisms of VCG

31.2.1 Privacy

VCG mechanism forces the agents to report their valuations truthfully. If the interaction had to happen beyond one round, it may be preferable for the bidders to use a mechanism that uses the minimal information needed for the current round. Because, if all the agents know the valuations of other agents perfectly they can adjust their actions in the following rounds to achieve a higher utility. The privacy concern is also critical for the trust of a new auctioneer, since the auctioneer has the opportunity of introducing fake bidders to extract more money from a rich bidder.

31.2.2 Susceptibility to collusion

Consider an example with three agents and two possible allocations A and B as shown in Figure 31.1, by the VCG mechanism the player 1 and 2 should pay 150 and 50 respectively.

	A	B	Payment
1	200	0	150
2	100	0	50
3	0	250	0

Figure 31.1: p^{VCG} before manipulation

But, the players 1 and 2 together can change the amount of payment they have to make by manipulating their valuations, as shown in the Figure 31.2.

	A	B	Payment
1	250	0	100
2	150	0	0
3	0	250	0

Figure 31.2: p^{VCG} after manipulation

Therefore, VCG is not group strategy proof.

31.2.3 Not frugal

VCG mechanism does not charge an amount close to the seller’s valuation/cost. It guaranteed to bring revenue (since it is a no-deficit mechanism), but the payment could be very large. Consider an example of shortest path, routing a packet or delivery of an item as shown in the Figure 31.3.




Figure 31.3: A network to deliver an item from A to F

Consider each edge as an agent. The efficient allocation for the shortest path is $A \rightarrow B \rightarrow E \rightarrow F$ with cost 4. Notice the effect of payment to $A \rightarrow B$ due to the cost of $A \rightarrow C$. If the cost of $A \rightarrow C$ is 2 then the payment of $A \rightarrow B$ is:

$$P^{AB} = \sum_{j \neq i} v_j(x_{-j}^*) - \sum_{i \in N} v_i(x^*) = 0 - (-1) + (-2) - 0 + (-3) - 0 = -4$$

Now, if cost of $A \rightarrow C$ becomes $x > 2$, then $p^{AB} = -(x + 2)$ and $A \rightarrow B$ will get more money than before even though the cost for that edge remains the same. Clearly the payment is unbounded if x is unbounded.

31.2.4 Revenue Monotonicity Violated

A system maintains the revenue monotonicity, if the revenue increases with the increase in number of agents. But, the VCG mechanism violates the revenue monotonicity. Let us take an example given in the Figure 31.4.

	X	Y	p^{VCG}
1	0	90	0
2	100	0	90

Figure 31.4: The p^{VCG} before adding a third agent

But, if a new agent participates then the p^{VCG} for all the agent changes and the result is shown in Figure 31.5.

	X	Y	p^{VCG}
1	0	90	0
2	100	0	0
3	100	0	0

Figure 31.5: The p^{VCG} after adding the third agent

Hence both split or merging of agents in VCG has problems.

31.2.5 Not fully Budget Balanced

VCG mechanism does not always satisfies the balanced budget, almost always some surplus will be left. This surplus can not be redistributed as that will change the payoff of the agents and it can not be invested to things that can effect the payoffs. Therefore, it has to be destroyed and known as *money burning*.

31.3 Generalization of VCG

In view of the limitations of VCG mechanism, the class of mechanisms is needed to be expanded to take care of criticisms. This in particular considers the budget balance issue. VCG puts equal weightage for every agent which may be relaxed to yield mechanisms satisfying more desirable properties. In the following section, we consider a special class of allocation rules that subsumes the VCG allocation rule.

31.3.1 Affine maximizer allocation rule

Definition 31.1 (Affine Maximizer) An allocation rule f is called affine maximizer (AM) if there exists weights $w_i \geq 0, \forall i \in N$, not all zero, and a function $\kappa : A \rightarrow \mathbb{R}$ such that

$$f^{AM}(\theta) \in \arg \max_{a \in A} \left[\sum_{i \in N} w_i \theta_i(a) + \kappa(a) \right].$$

Observe that this is a superclass of VCG allocation rule. The question we will ask in this section is similar to the Gibbard-Satterthwaite kind of question in voting setup. In GS theorem, we allowed the preferences to be all possible total ordering, while here we will consider valuations that can take any arbitrary value. The question will be what class of allocation rules are implementable in such a domain.

Definition 31.2 (Independence of Non-influential Agents) *An affine maximizer rule with weights w_i , $i \in N$ and κ satisfies independence of non-influential agents (INA) if for all $i \in N$ with $w_i = 0$ we have:*

$$f(\theta_i, \theta_{-i}) = f(\theta'_i, \theta_{-i}), \forall \theta_i, \theta'_i, \theta_{-i}$$

The INA condition is a *tie-breaking condition*. If the affine maximizer returns a set of allocations instead of a single maximizer, this condition ensures that the ties are always broken consistently and is not influenced by an agent whose report is not accounted in the affine maximization problem. If a dictatorial rule has some value for two different alternatives and the tie is broken by the valuation of a non-dictatorial agent, then it violates INA.

Theorem 31.3 *An INA affine maximizer rule is implementable.*

Proof: Consider the payment

$$p_i^{AM} = \begin{cases} \frac{1}{w_i} \left[h_i(\theta_{-i}) - \left(\sum_{j \neq i} w_j \theta_j(f^{AM}(\theta)) + \kappa(f^{AM}(\theta)) \right) \right] & \forall i : w_i > 0 \\ 0 & \forall i : w_i = 0 \end{cases} \quad (31.1)$$

The payoff for agent i , if $w_i > 0$, is given by

$$\begin{aligned} & \theta_i(f^{AM}(\theta)) - p_i^{AM}(\theta_i, \theta_{-i}) \\ &= \frac{1}{w_i} \left[\sum_{j \in N} w_j \theta_j(f^{AM}(\theta)) + \kappa(f^{AM}(\theta)) - h_i(\theta_{-i}) \right] \\ &\geq \frac{1}{w_i} \left[\sum_{j \in N} w_j \theta_j(f^{AM}(\theta'_i, \theta_{-i})) + \kappa(f^{AM}(\theta'_i, \theta_{-i})) - h_i(\theta_{-i}) \right] \\ &= \theta_i(f^{AM}(\theta'_i, \theta_{-i})) - \underbrace{\frac{1}{w_i} \left[h_i(\theta_{-i}) - \left(\sum_{j \neq i} w_j \theta_j(f^{AM}(\theta'_i, \theta_{-i})) + \kappa(f^{AM}(\theta'_i, \theta_{-i})) \right) \right]}_{=p_i^{AM}(\theta'_i, \theta_{-i})} \end{aligned}$$

Therefore,

$$\theta_i(f^{AM}(\theta)) - p_i^{AM}(\theta_i, \theta_{-i}) \geq \theta_i(f^{AM}(\theta'_i, \theta_{-i})) - p_i^{AM}(\theta'_i, \theta_{-i}) \quad (31.2)$$

Hence agent $i : w_i > 0$ is truthful.

and for $i : w_i = 0$, due to INA,

$$f(\theta_i, \theta_{-i}) = f(\theta'_i, \theta_{-i}), \forall \theta_i, \theta'_i, \theta_{-i}$$

and payment is zero. So, the agent is weakly truthful. ■

Similar to GS theorem, we ask what if the valuations are *unrestricted*. An unrestricted valuation is the mapping $\theta_i : A \mapsto \mathbb{R}$ without any additional constraints. The answer is given in the following theorem which we present without proof.

Theorem 31.4 (Roberts 1979) *Let A be finite with $|A| \geq 3$. If the type space is unrestricted, then every onto and implementable allocation rule must be an affine maximizer.*

As before, restricting the space of valuations types gives us more mechanisms that are DSIC. For example, there are mechanisms other than VCG or Groves mechanisms that are truthful in auction settings.

Lecture 32: October 31, 2017

Lecturer: Swaprava Nath

Scribe(s): Sachin K Salim

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

32.1 Single Object Auction Model

Let the type set of agent i be $T_i \subseteq \mathbb{R}$ and $t_i \in T_i$ be the value agent i gets if he wins the object.

An allocation a is a vector of length n , where a_i denotes the probability that i wins the object. The set of allocations is denoted by

$$\Delta A = \{a \in \mathbb{R}_{\geq 0}^n : \sum_{i=1}^n a_i = 1\},$$

and the allocation rule is a function $f : T_1 \times T_2 \times \dots \times T_n \rightarrow \Delta A$ where $f_i(t_i, t_{-i})$ is the probability that i wins the object when type profile is (t_i, t_{-i}) .

Given an allocation $a = (a_1, \dots, a_n)$, the valuation of agent i is given by the *product form* $a_i \cdot t_i$.

The setting and results can be extended to setting where a_i is the amount of an object allocated to agent i . For example, in sponsored search auctions, a_i is replaced by CTR_i . In general, the analysis presented here is similar if the valuation of every agent is in the product form, where one factor in the product comes as the outcome of the allocation rule and the other factor is a scalar type of the agent. This domain of valuations is called “single parameter domain”.

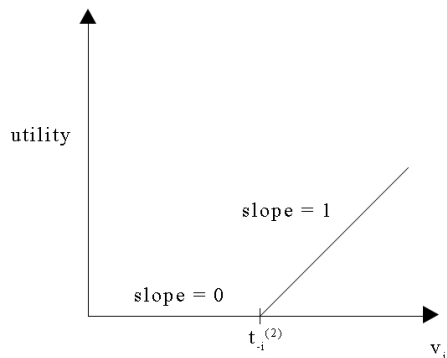

32.1.1 Vickrey (Second Price) Auction

Consider the utility and allocation of the object of agent i w.r.t. the type of the agent as shown in Figures 32.1 and 32.2 respectively. Agent i wins if $v_i > t_{-i}^{(2)}$ and loses if $v_i < t_{-i}^{(2)}$ where v_i is the type of agent i , i.e., the value of the object if i wins. $t_{-i}^{(2)}$ is defined as $\max_{j \neq i} v_j$.

The utility u_i of agent i (Figure 32.1) is defined as

$$u_i = \begin{cases} 0 & \text{if } v_i \leq t_{-i}^{(2)} \\ v_i - t_{-i}^{(2)} & \text{if } v_i > t_{-i}^{(2)} \end{cases}$$

We see that the utility is a convex function. The derivative of the utility function is 0 for $v_i < t_{-i}^{(2)}$ and 1 for $v_i > t_{-i}^{(2)}$, and the function is not differentiable at $v_i = t_{-i}^{(2)}$. Also, its derivative coincides with the probability of winning wherever it exists (Figure 32.2).

Figure 32.1: Utility of agent i Figure 32.2: Allocation of agent i

32.2 Some Results from Convex Analysis

We are interested in functions $g : \mathbb{I} \rightarrow \mathbb{R}$, where $\mathbb{I} \subseteq \mathbb{R}$ is an interval.

Definition 32.1 : A function $g : \mathbb{I} \rightarrow \mathbb{R}$ is convex if for every $x, y \in \mathbb{I}$ and $\lambda \in [0, 1]$,

$$\lambda g(x) + (1 - \lambda)g(y) \geq g(\lambda x + (1 - \lambda)y)$$

Facts:

1. Convex functions are continuous in the interior of its domain; jumps can only occur only at the boundaries.
2. Convex functions are differentiable “almost everywhere” in \mathbb{I} . Formally, there exists a $\mathbb{I}' \subseteq \mathbb{I}$ such that $\mathbb{I} \setminus \mathbb{I}'$ has countable points (has measure zero) and g is differentiable at every point of \mathbb{I}' .

32.2.1 Subgradient

Definition 32.2 For any $x \in \mathbb{I}$, x^* is a subgradient of the function $g : \mathbb{I} \rightarrow \mathbb{R}$ at x if

$$g(z) \geq g(x) + x^* \cdot (z - x), \quad \forall z \in \mathbb{I}$$

Lemma 32.3 Let $g : \mathbb{I} \rightarrow \mathbb{R}$ is a convex function. Suppose x is in the interior of \mathbb{I} and g is differentiable at x . Then $g'(x)$ is the unique subgradient of g at x .

Proof: Consider $z \in \mathbb{I}$ such that $z > x$ (a similar proof works if $z < x$). Consider h such that $(z - x) \geq h > 0$.

$$x + h = \frac{h}{z - x}z + \left(1 - \frac{h}{z - x}\right)x$$

Since g is convex,




Figure 32.3: $g(x)$

$$\begin{aligned} \frac{h}{z-x} g(z) + \left(1 - \frac{h}{z-x}\right) g(x) &\geq g(x+h) \\ \implies \frac{g(z) - g(x)}{z-x} &\geq \frac{g(x+h) - g(x)}{h} \end{aligned}$$

The above result holds for any $h > 0$. So when $h \rightarrow 0$,

$$g(z) - g(x) \geq g'(x) (z - x)$$

Hence $g'(x)$ is a subgradient of g at x .

Now, we need to show uniqueness. Say for contradiction, there exists another subgradient $x^* \neq g'(x)$ at x .

Case 1: $x^* > g'(x)$. By definition,

$$\begin{aligned} g(x+h) - g(x) &\geq x^* h \\ \implies \frac{g(x+h) - g(x)}{h} &\geq x^* > g'(x) \end{aligned}$$

Taking limit as $h \rightarrow 0$

$$g'(x) \geq x^* > g'(x)$$

But this is a contradiction.

Case 2: $x^* < g'(x)$. A similar contradiction can be reached. ■

Lemma 32.4 Let $g : \mathbb{I} \rightarrow \mathbb{R}$ be a convex function. Then for every $x \in \mathbb{I}$, the subgradient of g at x exists.

Fact: For points in $\mathbb{I} \setminus \mathbb{I}'$, the set of subgradients at a point forms a convex set.

$$\text{Define } g'_+(x) = \lim_{\substack{z \rightarrow x \\ z \in \mathbb{I}, z > x}} g'(z) \text{ and } g'_-(x) = \lim_{\substack{z \rightarrow x \\ z \in \mathbb{I}, z < x}} g'(z)$$

The set of subgradients at $x \in \mathbb{I} \setminus \mathbb{I}'$ is $[g'_-(x), g'_+(x)]$

The set of subgradients of g at a point $x \in \mathbb{I}$ is denoted by $\partial g(x)$.

Lemma 32.3 says that $\partial g(x)$ is $\{g'(x)\}$ for $x \in \mathbb{I}'$ and Lemma 32.4 says that it is non-empty for all $x \in \mathbb{I}$.

Lemma 32.5 *Let $g : \mathbb{I} \rightarrow \mathbb{R}$ be a convex function. Let $\phi : \mathbb{I} \rightarrow \mathbb{R}$ such that $\phi(z) \in \partial g(z), \forall z \in \mathbb{I}$. Then $\forall x, y \in \mathbb{I}$ such that $x > y$, we have $\phi(x) \geq \phi(y)$.*

Proof: By definition,

$$g(x) \geq g(y) + \phi(y)(x - y)$$

$$g(y) \geq g(x) + \phi(x)(y - x)$$

Adding the above two equations, $(\phi(x) - \phi(y))(x - y) \geq 0$

$$\implies \phi(x) \geq \phi(y) \quad (\text{because } x > y)$$

■

If g was differentiable everywhere, then we know that

$$g(x) = g(y) + \int_y^x g'(z) dz$$

An extension of this result holds for convex functions with subgradients.

Lemma 32.6 *Let $g : \mathbb{I} \rightarrow \mathbb{R}$ be a convex function. Then for any $x, y \in \mathbb{I}$,*

$$g(x) = g(y) + \int_y^x \phi(z) dz$$

where $\phi : \mathbb{I} \rightarrow \mathbb{R}$ is such that $\phi(z) \in \partial g(z), \forall z \in \mathbb{I}$.

Lecture 33: November 1, 2017

Lecturer: Swaprava Nath

Scribe(s): Neeraj Yadav

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

33.1 Recap

We have studied some basic properties and results dealing with convex functions. We will use two lemmas from the previous lecture (reproduced here for completeness) to prove Myerson's characterization theorem.

Lemma 33.1 Let $g : \mathbb{I} \rightarrow \mathbb{R}$ be a convex function. Let $\phi : \mathbb{I} \rightarrow \mathbb{R}$ such that $\phi(z) \in \partial g(z), \forall z \in \mathbb{I}$. Then $\forall x, y \in \mathbb{I}$ such that $x > y$, we have $\phi(x) \geq \phi(y)$.

Lemma 33.2 Let $g : \mathbb{I} \rightarrow \mathbb{R}$ be a convex function. Then for any $x, y \in \mathbb{I}$,

$$g(x) = g(y) + \int_y^x \phi(z) dz$$

where $\phi : \mathbb{I} \rightarrow \mathbb{R}$ is such that $\phi(z) \in \partial g(z), \forall z \in \mathbb{I}$.

33.2 Monotonicity and Characterization of DSIC Mechanisms for Single Item Auctions

Definition 33.3 (Non-decreasingness) An allocation rule is non-decreasing if for every agent $\forall i \in N$ and $\forall t_{-i} \in T_{-i}$ we have $f_i(t_i, t_{-i}) \geq f_i(s_i, t_{-i}), \forall t_i, s_i \in T_i$ with $t_i > s_i$.

In words, if the types of other agents are held fixed, then the probability of allocation of the object to an agent weakly increases with the increase in her valuation. We now present the main characterization theorem of this section.

Theorem 33.4 (Myerson 1981) Suppose $T_i = [0, b_i], \forall i \in N$ and the valuations are in product form. An allocation rule $f : T \mapsto \Delta A$ and a payment rule (p_1, p_2, \dots, p_n) is DSIC iff

1. The allocation f is non-decreasing, and,
2. Payment is given by

$$p_i(t_i, t_{-i}) = p_i(0, t_{-i}) + t_i f_i(t_i, t_{-i}) - \int_0^{t_i} f_i(x_i, t_{-i}) dx_i.$$

Proof: (\Rightarrow) Consider the utility of agent i when her types are t_i and s_i respectively and other agents' types are t_{-i} .

$$\begin{aligned} u_i(t_i, t_{-i}) &= t_i f_i(t_i, t_{-i}) - p_i(t_i, t_{-i}) \\ u_i(s_i, t_{-i}) &= s_i f_i(s_i, t_{-i}) - p_i(s_i, t_{-i}) \end{aligned}$$

Since (f, p) is DSIC, it must hold that $\forall s_i, t_i, t_{-i}$

$$\begin{aligned} u_i(t_i, t_{-i}) &= t_i f_i(t_i, t_{-i}) - p_i(t_i, t_{-i}) \\ &\geq t_i f_i(s_i, t_{-i}) - p_i(s_i, t_{-i}) \\ &= s_i f_i(s_i, t_{-i}) + f_i(s_i, t_{-i})(t_i - s_i) - p_i(s_i, t_{-i}) \\ &= u_i(s_i, t_{-i}) + f_i(s_i, t_{-i})(t_i - s_i). \end{aligned}$$

Where the first inequality results from the fact that (f, p) is DSIC and the following equalities result from some algebraic manipulation. Define $g(t_i) := u_i(t_i, t_{-i})$, and $\phi(t_i) = f_i(t_i, t_{-i})$. Hence, the above inequality can be written as

$$g(t_i) \geq g(s_i) + \phi(s_i)(t_i - s_i).$$

Therefore, we conclude that $\phi(s_i)$ is a subgradient of g at s_i .

Convexity of g : Pick $x_i, z_i \in T_i$ and define $y_i = \lambda x_i + (1 - \lambda)z_i$ where $\lambda \in [0, 1]$. DSIC implies

$$\begin{aligned} g(x_i) &\geq g(y_i) + \phi(y_i)(x_i - y_i) \\ g(z_i) &\geq g(y_i) + \phi(y_i)(z_i - y_i) \end{aligned}$$

Multiplying the first equation with λ and the second with $(1 - \lambda)$ and adding we get

$$\begin{aligned} \lambda g(x_i) + (1 - \lambda)g(z_i) &\geq g(y_i) + \phi(y_i)(\lambda x_i + (1 - \lambda)z_i - y_i) \\ \Rightarrow \lambda g(x_i) + (1 - \lambda)g(z_i) &\geq g(y_i) \end{aligned}$$

Where the implication holds since $\lambda x_i + (1 - \lambda)z_i - y_i = 0$. This proves that g is convex.

Having proved that g is convex, we can apply Lemma 33.1. As $\phi \equiv f_i(\cdot, t_{-i})$ is a subgradient of $g \equiv u_i(\cdot, t_{-i})$, which is convex, we conclude that $f_i(\cdot, t_{-i})$ is non-decreasing. This proves part 1 of the implication of the theorem.

By Lemma 33.2, we have

$$\begin{aligned} g(t_i) &= g(0) + \int_0^{t_i} \phi(x_i) dx_i \\ \Rightarrow u_i(t_i, t_{-i}) &= u_i(0, t_{-i}) + \int_0^{t_i} f_i(x_i, t_{-i}) dx_i \\ \Rightarrow t_i f_i(t_i, t_{-i}) - p_i(t_i, t_{-i}) &= -p_i(0, t_{-i}) + \int_0^{t_i} f_i(x_i, t_{-i}) dx_i \\ \Rightarrow p_i(t_i, t_{-i}) &= p_i(0, t_{-i}) + t_i f_i(t_i, t_{-i}) - \int_0^{t_i} f_i(x_i, t_{-i}) dx_i. \end{aligned}$$

This proves part 2 of the implication.

(\Leftarrow) For the other direction, it is given that allocation rule is non-decreasing and payment rule is given by

$$p_i(t_i, t_{-i}) = p_i(0, t_{-i}) + t_i f_i(t_i, t_{-i}) - \int_0^{t_i} f_i(x_i, t_{-i}) dx_i$$

We need to prove that (f, p) is DSIC.

Let t_i be the true type and s_i be the reported type of agent i , and the types of the other agents are t_{-i} . The utilities of the agent in the true and reported types are respectively

$$t_i f_i(t_i, t_{-i}) - p_i(t_i, t_{-i}) = t_i f_i(t_i, t_{-i}) - p_i(0, t_{-i}) - t_i f_i(t_i, t_{-i}) + \int_0^{t_i} f_i(x_i, t_{-i}) dx_i$$

$$t_i f_i(s_i, t_{-i}) - p_i(s_i, t_{-i}) = t_i f_i(s_i, t_{-i}) - p_i(0, t_{-i}) - s_i f_i(s_i, t_{-i}) + \int_0^{s_i} f_i(x_i, t_{-i}) dx_i$$

The difference between the two utilities is given by

$$t_i f_i(t_i, t_{-i}) - p_i(t_i, t_{-i}) - [t_i f_i(s_i, t_{-i}) - p_i(s_i, t_{-i})] = (s_i - t_i) f_i(s_i, t_{-i}) + \int_{s_i}^{t_i} f_i(x_i, t_{-i}) dx_i \geq 0$$

The inequality results from the fact that f is non-decreasing. Since s_i, t_i, t_{-i} are arbitrary, we conclude that the mechanism is DSIC. Figure 33.1 illustrates this fact for a specific allocation rule and the corresponding payment.

When the reported type is greater than the true type, i.e., $s_i > t_i$, the RHS of the equality is the shaded region in the second subfigure. Similarly, the RHS of the equality is the shaded region in the third subfigure when $s_i < t_i$.




Figure 33.1: Illustration of the incentive compatibility of the payment rule

■

Lecture 34: November 3, 2017

Lecturer: Swaprava Nath

Scribe(s): Gundeep Arora

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

34.1 Recap

Continuing our discussion from the previous lecture, in which we had defined the notion of monotonicity of an allocation rule and stated the Myerson's theorem.

Theorem 34.1 (Myerson 1981) Suppose $T_i = [0, b_i], \forall i \in N$ and the valuations are in product form. An allocation rule $f : T \mapsto \Delta A$ and a payment rule (p_1, p_2, \dots, p_n) is DSIC iff

1. The allocation f is non-decreasing, and,
2. Payment is given by

$$p_i(t_i, t_{-i}) = p_i(0, t_{-i}) + t_i f_i(t_i, t_{-i}) - \int_0^{t_i} f_i(x_i, t_{-i}) dx_i.$$

In this lecture, we will look at the implications of the payment rule and a number of examples to illustrate the rule and compare this characterization result with that of Roberts' where only one specific class of allocation rules was implementable.

34.2 Discussions on the Payment Rule

Revenue Equivalence Consider the case where we have two different DSIC mechanisms given by (f, p) and (f, q) with the same allocation function f and different payment functions p and q . From the payment function characterization of Myerson's result, we see that the difference between the payments lies only in the constant terms, $p_i(0, t_{-i})$ and $q_i(0, t_{-i})$. Hence, we see that the following difference term is same for all the payments that implement f .

$$\begin{aligned} p_i(t_i, t_{-i}) - p_i(s_i, t_{-i}) &= q_i(t_i, t_{-i}) - q_i(s_i, t_{-i}) \quad \forall s_i \in T \\ \implies p_i(t_i, t_{-i}) &= q_i(t_i, t_{-i}) + h_i(t_{-i}) \end{aligned} \tag{34.1}$$

The term $h_i(t_{-i})$ is independent of i -th agents type t_i . So any two payments for an agent differs by a function that is independent of the type of that agent. Whenever any two payment rules that implement the same allocation rule that can be written in this form, we call the allocation rule to be **revenue equivalent**. The intuition is that for every payment rule the revenue generated is 'almost' equivalent.

Note that this is a much stronger statement than the fact that if any function $h_i(t_{-i})$ is added to the payment, that payment rule also implements the same allocation rule (we saw this before). This result says that for every payment rule that implements the same allocation rule must be related in this form.

Difference with Roberts' characterization theorem It is worth noting the difference of this result with the characterization provided by Roberts' theorem, that gives an explicit formula for the allocation rule. Roberts' result says that when the space of valuations is unrestricted, the allocation rule must be from the affine maximizer class. This is a precise functional form characterization unlike the 'implicit' characterization given by the Myerson's result, which says that the allocation should be non-decreasing, but does not give any structural guarantees.

Corollary 34.2 *An allocation rule is implementable iff it is non-decreasing.*

34.2.1 Some examples of non-decreasing allocation function

1. **Constant allocation:** This allocation is trivially non-decreasing. Payment given by the formula is constant, which can be zero. A special case of this allocation rule is **dictatorial allocation**, where the object is deterministically given to one agent, who is the *dictator*. The payment again is constant.
2. **Vickrey auction:** Here the allocation is efficient and the payment for each agent, fixing the other agents, is non-decreasing.

$$\text{allocation : } f_i(t) = \begin{cases} 1 & t_i > t_{-i}^{(2)} \\ 0 & t_i < t_{-i}^{(2)} \\ \alpha_i & t_i = t_{-i}^{(2)} \end{cases}$$

Where α_i is such that $\sum_i \alpha_i = 1$ and $\alpha_i \geq 0 \forall i \in N$, and $t_{-i}^{(2)} = \max_{j \neq i} t_j$. The characterization theorem says that a rule that breaks the tie in an arbitrary probabilistic way is DSIC. The allocation function can be seen as the subgradient of the utility function, given by:

$$u(v_i) = \begin{cases} 0 & v_i \leq t_{-i}^{(2)} \\ v_i - p_i & v_i > t_{-i}^{(2)} \end{cases}$$

Note that this function is not differentiable at $v_i = t_{-i}^{(2)}$ and therefore the subgradient at that point can be anything in $[0, 1]$. Therefore any $\alpha_i \in [0, 1]$ is a valid candidate for an implementable allocation.

The payment rule (for the winning bidder, say player 1) can be written as:

$$\begin{aligned} p(t_1, t_{-1}) &= 0 + t_1 \cdot f_1 - \int_{t_{-1}^{(2)}}^{t_1} f_1(x_1, t_{-1}) dx_1 \\ &= t_1 - (t_1 - t_{-1}^{(2)}) \\ &= t_{-1}^{(2)} \end{aligned} \tag{34.2}$$

3. **Efficient allocation with reserve:** The allocation here gives the bidder i the item if $t_i > \max(r, t_{-i}^{(2)})$ where r is the reserve price set by the auctioneer. The payment made by the winning bidder is $\max(r, t_{-i}^{(2)})$. Clearly, the allocation rule is non-decreasing. Here the item is not sold if no bid is higher than the reserve price.
4. **A not so common allocation rule:** Consider an allocation rule for two agents $N = \{1, 2\}$ and $A = \{a_0, a_1, a_2\}$ where the allocation a_0 refers to the item being unsold and a_i being the item allotted to player i . Give a type profile (t_1, t_2) the seller computes

$$U(t) = \max\{2, t_1^2, t_2^3\} \tag{34.3}$$

The allocation proceeds as:

$$\begin{aligned}
 a_0 & \text{ if } U(t) = 2 \quad \text{i.e. } 2 > \max\{t_1^2, t_2^3\} \\
 a_1 & \text{ if } U(t) = t_1^2 \quad \text{i.e. } t_1 > \sqrt{\max\{2, t_2^3\}} \\
 a_2 & \text{ if } U(t) = t_2^3 \quad \text{i.e. } t_2 > \sqrt[3]{\max\{2, t_1^2\}}
 \end{aligned} \tag{34.4}$$

It is easy to see that the payments are zero, $\sqrt{\max\{2, t_2^3\}}$, and $\sqrt[3]{\max\{2, t_1^2\}}$ respectively for the three cases above.

34.3 Individual Rationality

Definition 34.3 (Individual Rationality) A mechanism (f, p) is “ex-post” individually rational if

$$t_i \cdot f_i(t_i, t_{-i}) - p_i(t_i, t_{-i}) \geq 0 \quad \forall t_i \in T_i \quad \forall t_{-i} \in T_{-i} \quad \forall i \in N \tag{34.5}$$

Here “ex-post” refers to the idea that even after revealing everyone’s type, participation is weakly preferable.

Lemma 34.4 Suppose (f, p) is DSIC,

1. It is individually rational (IR) iff

$$p_i(0, t_{-i}) \leq 0 \quad \forall i \in N \quad \forall t_{-i} \in T_{-i} \tag{34.6}$$

2. It is IR and gives no subsidy, i.e. $p_i(t_i, t_{-i}) \geq 0 \forall t_i \in T_i$ iff

$$p_i(0, t_{-i}) = 0 \quad \forall i \in N \quad \forall t_{-i} \in T_{-i}. \tag{34.7}$$

Proof: We shall present the proof part by part, and both directions (if and only if) for each

1. **only if direction:** Assume

$$\begin{aligned}
 t_i \cdot f_i(t_i, t_{-i}) - p_i(t_i, t_{-i}) & \geq 0 \\
 \text{for } t_i = 0 \quad p_i(0, t_{-i}) & \leq 0
 \end{aligned} \tag{34.8}$$

if direction : Assume $p_i(0, t_{-i}) \leq 0$

$$t_i \cdot f_i(t_i, t_{-i}) - p_i(t_i, t_{-i}) - t_i \cdot f_i(t_i, t_{-i}) - \int_0^{t_i} f_i(x_i, t_{-i}) dx_i \geq 0 \tag{34.9}$$

Since in this case we assume $p_i(t_i, t_{-i}) \leq 0$, the above inequality holds.

2. **only if direction:** Assume the mechanism DSIC, IR and satisfies no subsidy hence,

$$p_i(0, t_{-i}) \geq 0 \tag{34.10}$$

but using the above proof, we have

$$\begin{aligned}
 p_i(0, t_{-i}) & \leq 0 \\
 \Rightarrow p_i(0, t_{-i}) & = 0
 \end{aligned} \tag{34.11}$$

The other direction is obvious. ■

34.4 Some Non-Vickrey Auctions

1. **Redistribution:** Consider a case where the auction is Groves but not Vickrey auction. In this auction, the highest bidder wins and gets the object but the payment is such that everyone pays what they would have paid in a Vickrey auction, but is compensated some amount given by,

$$p_i(0, t_{-i}) = -\frac{1}{n}z_{-i}^{(2)} \quad (34.12)$$

where $z_{-i}^{(2)}$ = second highest among $\{t_j; j \neq i\}$. WLOG, assume that $t_1 > t_2 > \dots > t_n$. So, we get

$$\begin{aligned} \text{payment of player 1} &= \frac{-1 \cdot t_3}{n} + t_2 \\ \text{payment of player 2} &= \frac{-1 \cdot t_3}{n} \\ \text{payment of every other agent} &= \frac{-1 \cdot t_2}{n} \\ \text{Hence, the sum of payment} &= \frac{2 \cdot (t_2 - t_3)}{n} \end{aligned}$$

The above expression tells us that such an auction is asymptotically budget balanced (as the surplus is redistributed as $n \rightarrow \infty$) while still being DSIC. The allocation however is still deterministic in nature and one can do something better by randomizing the allocation.

2. **Green-Laffont mechanism:** Give object to the highest bidder with probability $(1 - \frac{1}{n})$, and to the second highest bidder with probability $\frac{1}{n}$. Say $t_1 > t_2 > \dots > t_n$, and let

$$p_i(0, t_{-i}) = -\frac{1}{n}z_{-i}^{(2)} \quad (34.13)$$

where $z_{-i}^{(2)}$ = second highest among $\{t_j; j \neq i\}$

$$\text{payment of 1} = -\frac{1}{n} \cdot t_3 + \left(1 - \frac{1}{n}\right) \cdot t_1 - \frac{1}{n} \cdot (t_2 - t_3) - \left(1 - \frac{1}{n}\right) \cdot (t_1 - t_2) \quad (34.14)$$

$$\text{payment of 2} = \frac{-t_3}{n} + \frac{t_2}{n} - \frac{t_2 - t_3}{n} = 0 \quad (34.15)$$

$$\text{payment of every other agent} = -\frac{t_2}{n} \quad (34.16)$$

$$\text{Hence, the sum of payment} = \frac{n-2}{n} \cdot t_2 - (n-2) \cdot \frac{t_2}{n} = 0 \quad (34.17)$$

Hence such an allocation with the given payment mechanism is (exactly) budget balanced.

Lecture 35: November 7, 2017

Lecturer: Swaprava Nath

Scribe(s): Swaprava Nath

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

35.1 Revenue Maximization

We need to benchmark mechanisms w.r.t. their expected revenue. The benchmark is typically done w.r.t. a given prior distribution of the types. Therefore, we move from a prior-free environment to a weaker setup due to two major reasons: (1) for the prior-free setup, we can ask the worst-case revenue earned, which can be arbitrarily bad, hence we cannot achieve any useful result, and (2) in majority of real-world applications, the agents interact repeatedly with the system and it is possible to build a model for their valuation distribution.

Since we move to a Bayesian setting, all the notions of truthfulness and participation constraints need to be redefined.

35.2 Bayesian Incentive Compatibility

Model: Let the type set of agent i be $T_i = [0, b_i], \forall i \in N$. The common prior G is defined over the set of type profiles $T := \times_{i \in N} T_i$, and g denotes the density of the joint distribution.

The conditional distribution over the types of agents except i when the type of agent i is t_i is denoted by $G_{-i}(\cdot | t_i)$. Let $g_{-i}(\cdot | t_i)$ denote the density of this distribution. The conditional density $g_{-i}(t_{-i} | t_i)$ is derived from g using Bayes' rule.

Every mechanism (f, p_1, \dots, p_n) induces an expected allocation and payment rule (α, π) given by

$$\alpha_i(s_i | t_i) = \int_{s_{-i} \in T_{-i}} f_i(s_i, s_{-i}) g_{-i}(s_{-i} | t_i) ds_{-i},$$

$$\pi_i(s_i | t_i) = \int_{s_{-i} \in T_{-i}} p_i(s_i, s_{-i}) g_{-i}(s_{-i} | t_i) ds_{-i}.$$

Where the notation $(s_i | t_i)$ denote that s_i is reported when the true type is t_i . Therefore, the expected utility of agent i is

$$t_i \cdot \alpha_i(s_i | t_i) - \pi_i(s_i | t_i).$$

Note: since the randomization in a Bayesian setting has two following levels, the meaning of 'expected' has to be understood appropriately.

- The first level is w.r.t. the types of other agents, which comes from the common prior in this setting.
- The second level is w.r.t. the randomization of the mechanism (allocation rule) – we have used this in our discussions earlier.

In the context of Bayesian common prior setting, the term ‘expected’ will refer to the expectation w.r.t. **both** the randomization effects. We define the notion of truthfulness in this setting.

Definition 35.1 (Bayesian Incentive Compatibility) A mechanism (f, p) is Bayesian Incentive compatible (BIC) if $\forall s_i, t_i \in T_i, i \in N$

$$t_i \cdot \alpha_i(t_i | t_i) - \pi_i(t_i | t_i) \geq t_i \cdot \alpha_i(s_i | t_i) - \pi_i(s_i | t_i).$$

Consequently, an allocation rule f is Bayesian implementable if $\exists p$ such that (f, p) is BIC. We will use the shorthand $\alpha_i(t_i)$ and $\pi_i(t_i)$ to denote that the true and reported types are both t_i .

35.2.1 Independence of types

We assume that the types are independent, i.e., every agent’s type is drawn independently from a distribution G_i (with density g_i). This assumption is well motivated in the settings where agents cannot communicate with each other.

$$G(s_1, \dots, s_n) = \prod_{i \in N} G_i(s_i)$$

$$G_{-i}(s_{-i} | t_i) = \prod_{j \neq i} G_j(s_j)$$

Now we define a property of the allocation rule that will help us characterize the BIC rules.

Definition 35.2 (Non-decreasing in Expectation) An allocation rule is non-decreasing in expectation (NDE) if $\forall s_i, t_i \in T_i$ such that $s_i < t_i$, we have

$$\alpha_i(s_i) \leq \alpha_i(t_i), \forall i \in N.$$

Non-decreasing allocation rules are always NDE. Therefore, we expanded the space of mechanisms in the Bayesian setting.

Theorem 35.3 (Characterization of BIC rules) A mechanism (f, p) in the independent prior setting is BIC iff

1. f is NDE, and
2. p_i satisfies

$$\pi_i(t_i) := \pi_i(0) + t_i \alpha_i(t_i) - \int_0^{t_i} \alpha_i(s_i) ds_i, \forall t_i \in T_i, i \in N.$$

Proof is left as an exercise.

Note: A BIC mechanism may not be DSIC. We can see this through the following example where the allocation rule is NDE but **not** non-decreasing. Figure 35.1 shows that the types of two agents t_1 and t_2 are uniformly distributed over an unit square. The ‘1’s represent that the object is allocated to agent 1 when




Figure 35.1: An allocation rule that is NDE but not ND.

the types of the two agents fall within that tile. In all the other tiles the object is allocated to agent 2. We see that the expected allocation of agent 1

$$\alpha_1(t_1) = \begin{cases} 0 & \text{for } 0 \leq t_1 < 2/5 \\ 1/5 & \text{for } 2/5 \leq t_1 < 3/5 \\ 2/5 & \text{for } 3/5 \leq t_1 < 4/5 \\ 2/5 & \text{for } 4/5 \leq t_1 \leq 1 \end{cases} \quad (35.1)$$

Hence the allocation of agent 1 is NDE, similarly it can be shown that allocation of agent 2 is NDE too. However, it is clear that the allocation is not ND, since for $t_2 \in [1/5, 2/5)$, agent 1's allocation decreases from 1 to 0.

35.2.2 Interim Individual Rationality

In the Bayesian setup, we need to modify the definition of individual rationality as follows.

Definition 35.4 (Interim Individual Rationality) *A mechanism (f, p) is interim individually rational (IIR) if for every agent $i \in N$, we have*

$$t_i \cdot \alpha_i(t_i) - \pi_i(t_i) \geq 0, \forall t_i \in T_i.$$

The following result summarizes the requirements of a BIC and IIR mechanism.

Lemma 35.5 *A mechanism (f, p) is BIC and IIR iff*

1. f is NDE,
2. p_i satisfies

$$\pi_i(t_i) := \pi_i(0) + t_i \alpha_i(t_i) - \int_0^{t_i} \alpha_i(s_i) ds_i, \forall t_i \in T_i, i \in N,$$

3. $\forall i \in N, \pi_i(0) \leq 0$.

Proof: Conditions (1) and (2) characterize a BIC mechanism (Theorem 35.3). Hence, it remains to be shown that IIR together with (1) and (2) is equivalent to (3).

(IIR + 1 + 2 \Rightarrow 3): apply IIR at $t_i = 0$ in (2) to get $\pi_i(0) \leq 0$.

(IIR + 1 + 2 \Leftarrow 3): $t_i \cdot \alpha_i(t_i) - \pi_i(t_i) = -\pi_i(0) + \int_0^{t_i} \alpha_i(s_i) ds_i \geq 0$ if $\pi_i(0) \leq 0$. ■

35.3 Single Agent Problem

TO understand the dynamics of an optimal auction, we focus on the problem when there is exactly one buyer. The setting here is $T = [0, \beta]$, the mechanism is $(f, p) =: M$ where $f : [0, \beta] \mapsto [0, 1]$ and $p : [0, \beta] \mapsto \mathbb{R}$. The desirable properties are

- Incentive compatibility [DSIC and BIC are equivalent]

$$tf(t) - p(t) \geq tf(s) - p(s), \forall s, t \in T.$$

- Individual rationality [IR and IIR are equivalent]

$$tf(t) - p(t) \geq 0, \forall t \in T.$$

The revenue earned by the mechanism M is given by

$$\Pi^M := \int_0^\beta p(t)g(t)dt.$$

An **optimal mechanism** is a mechanism M^* such that it is IC and IR and satisfies $\Pi^{M^*} \geq \Pi^M, \forall M$.

From IC characterization, we know that

$$p(t) = p(0) + tf(t) - \int_0^t f(s)ds.$$

But from IR requirement, we have $p(0) \leq 0$. Since our goal is to maximize revenue, clearly $p(0) = 0$. Hence the payment reduces to

$$p(t) = tf(t) - \int_0^t f(s)ds.$$

This is completely given by the allocation rule f . Since the prior is G (with density g), the expected revenue is given by

$$\begin{aligned} \Pi^f &= \int_0^\beta p(t)g(t)dt \\ &= \int_0^\beta \left(tf(t) - \int_0^t f(s)ds \right) g(t)dt. \end{aligned}$$

Theorem 35.6 For any implementable allocation rule f , the revenue earned is given by

$$\Pi^f = \int_0^\beta w(t)f(t)g(t)dt$$

where $w(t) = \left(t - \frac{1-G(t)}{g(t)} \right)$, which is known as the virtual valuation of the agent.

Lecture 36: November 8, 2017

Lecturer: Swaprava Nath

Scribe(s): Sandipan Mandal

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

36.1 Recap

In the last lecture, we started discussing the optimal auction for a single object allocation. To build our intuition, we started with a single agent, i.e., there is only one potential buyer for the object and the auctioneer wants to maximize the revenue earned (maximize the payment of that agent subject to some conditions). We mentioned a result which gives a structure for the revenue function. In this lecture, we prove the result, and discuss the solution of *optimal revenue problem* for the single-agent case. Finally, we will extend the result to the multi-agent scenario.

36.2 Optimal Auction for a Single Agent

36.2.1 Struction of the revenue function

Theorem 36.1 For any implementable allocation rule f , the revenue earned is given by

$$\Pi^f = \int_0^\beta w(t)f(t)g(t)dt$$

where $w(t) = \left(t - \frac{1-G(t)}{g(t)}\right)$, which is known as the virtual valuation of the agent.

Proof: Consider the expected payment of the agent

$$\begin{aligned} \Pi^f &= \int_0^\beta p(t)g(t)dt \\ &= \int_0^\beta \left[tf(t) - \int_0^t f(x)dx\right] g(t)dt \\ &= \int_0^\beta tf(t)g(t)dt - \int_0^\beta \left(\int_0^t f(x)dx\right) g(t)dt \\ &= \int_0^\beta tf(t)g(t)dt - \int_0^\beta \left(\int_x^\beta f(x)g(t)dt\right) dx \quad (\text{changing the order of integration of the second term}) \\ &= \int_0^\beta tf(t)g(t)dt - \int_0^\beta \left(\int_x^\beta g(x)dx\right) f(t)dt \quad (\text{exchanging the variables } t \text{ and } x) \\ &= \int_0^\beta \left[tg(t) - \int_t^\beta g(x)dx\right] f(t)dt \end{aligned}$$

$$\begin{aligned}
&= \int_0^\beta [tg(t) - (G(\beta) - G(t))] f(t) dt \\
&= \int_0^\beta [tg(t) - (1 - G(t))] f(t) dt \\
&= \int_0^\beta \left(t - \frac{1 - G(t)}{g(t)} \right) g(t) f(t) dt \\
&= \int_0^\beta w(t) g(t) f(t) dt \qquad \text{(where } w(t) = t - \frac{1 - G(t)}{g(t)} \text{)}
\end{aligned}$$

The second equality follows because IC implies that the payment is given by Myerson's characterization result, and IR together with revenue maximization implies $p(0) = 0$. ■

36.2.2 Optimal revenue problem

Now we want to find the allocation function which maximizes the revenue of the auctioneer. So we need to solve the following optimization problem :

$$\max_{f : f \text{ is non-decreasing}} \Pi^f \qquad \text{(OPT1)}$$

Solving the above optimization problem is difficult, while the unconstrained version (given below) of the above problem is easier.

$$\max_f \Pi^f \qquad \text{(OPT2)}$$

Under the assumption of *monotone hazard rate*, the solutions of *OPT1* and *OPT2* are identical.

Assumption 36.2 *A prior distribution G satisfies the Monotone Hazard Rate (MHR) condition if $\frac{g(x)}{1-G(x)}$ is non-decreasing in x .*

Standard distributions like uniform and exponential satisfy MHR condition. We state a fact concerning MHR condition without proof.

Fact 36.3 *If G satisfies MHR, then \exists a unique x^* s.t.*

$$x^* = \frac{1 - G(x^*)}{g(x^*)} \implies w(x^*) = 0$$

Observe that if G satisfies MHR condition, w is strictly increasing. Therefore for $x < x^*$, w is negative and for $x > x^*$, w is positive. This observation gives us an easy way to solve the unconstrained optimization problem OPT2. As the goal is to maximize the total expected revenue, and for $t < x^*$, the virtual valuation is negative, we want not to allocate the object. Similarly, we want to sell the object for $t > x^*$. When $t = x^*$, the virtual valuation is zero, hence we can allocate the object with any probability and it does not affect the revenue. Therefore the allocation function is given as:

$$f(t) = \begin{cases} 0 & \text{for } t < x^* \\ \alpha \in [0, 1] & \text{for } t = x^* \\ 1 & \text{for } t > x^* \end{cases} \qquad \text{(36.1)}$$

Note that the above allocation function is non-decreasing as well. Hence it is a valid solution to constrained optimization problem OPT1. Therefore, f defined as above is implementable in DSIC. We now state a theorem to summarize the discussions we have had so far.

Theorem 36.4 A mechanism (f, p) under MHR condition is optimal iff the following two conditions hold:

1. allocation f is given by Equation 36.1, and
2. payment is given by $p(t) = f(t)x^*$, $\forall t \in T$.

36.3 Optimal Auction Mechanism for Multiple Agents

36.3.1 Struction of the revenue function

The optimal auction mechanism $M \equiv (f, p)$ is defined as one that is BIC, IIR and satisfies $\Pi^M \geq \Pi^{M'} \forall M'$. We know that

- BIC $\implies f_i$'s are non-decreasing in expectation (NDE) and expected payment $\Pi_i(t_i)$ has a specific formula.
- IIR and revenue optimality $\implies \Pi_i(0) = 0$.

Expected payment made by agent i is given by (assuming $T_i = [0, b_i]$):

$$\begin{aligned}
 \text{payment}_i &= \int_0^{b_i} \Pi_i(t_i) g(t_i) dt_i \\
 &= \int_0^{b_i} \left(t_i \alpha_i(t_i) - \int_0^{t_i} \alpha_i(x_i) dx_i \right) g(t_i) dt_i \\
 &= \int_0^{b_i} \left(t_i - \frac{1 - G_i(t_i)}{g_i(t_i)} \right) g_i(t_i) \alpha_i(t_i) dt_i && \text{(using the proof from the previous section)} \\
 &= \int_0^{b_i} w_i(t_i) \alpha_i(t_i) g_i(t_i) dt_i \\
 &= \int_0^{b_i} w_i(t_i) \int_{T_{-i}} f_i(t_i, t_{-i}) g_{-i}(t_{-i}) dt_{-i} g_i(t_i) dt_i && \text{(substituting the expression of } \alpha_i(t_i) \text{)} \\
 &= \int_T w_i(t_i) f_i(t) g(t) dt && \text{(where } T = \times_{i=1}^n T_i \text{)}
 \end{aligned}$$

Now, we have total revenue earned by the auctioneer as:

$$\Pi^M = \sum_{i \in N} \text{payment}_i = \int_T \left(\sum_{i \in N} w_i(t_i) f_i(t) \right) g(t) dt \tag{36.2}$$

Since the expected revenue is completely determined by the allocation rule f , we can replace Π^M with Π^f .

36.3.2 Optimal Revenue Problem

To find the allocation function which maximizes the revenue of the auctioneer, we need to solve the following constrained optimization problem

$$\max_{f : f \text{ is NDE}} \Pi^f \tag{OPT3}$$

Again we can see solving the above optimization problem is difficult. But solving the unconstrained version (given below) of the above problem is easier.

$$\max_f \Pi^f \quad (\text{OPT4})$$

Observe that Π^f involves taking convex combination of the virtual valuations of every agent (Eq. 36.2). From that expression of Π^f , it is clear that to maximize the integration we need to maximize a convex combination. We know that any convex combination $\sum_i \alpha_i x_i$ where $\sum_i \alpha_i = 1$ and $\alpha_i \in [0, 1]$ is maximized (where α_i s are variables) only when we set $\alpha_j = 1$ for $x_j : x_j \geq x_i, \forall i$ and set $\alpha_i = 0$ for $\forall i \neq j$. However, in the allocation problem we also have the option of not assigning the object to anyone (unsold), which will be optimal if the virtual valuations are negative for every agent. So following this scheme, we have the allocation of any agent i for the unconstrained problem to be the following when $\exists i \in N, w_i(t_i) \geq 0$.

$$f_i(t) = \begin{cases} 1 & \text{if } w_i(t_i) \geq w_j(t_j), \forall j \in N \\ 0 & \text{otherwise} \end{cases} \quad (36.3)$$

With an arbitrary tie-breaking rule. The allocation is $f_i(t) = 0, \forall i \in N$ if $w_i(t_i) < 0, \forall i \in N$.

Clearly the allocation rule is **not** non-decreasing and therefore **not** non-decreasing even in expectation (because some of w_i s can be decreasing w.r.t. t_i s which might cause the corresponding agent to have its allocation reduced from 1 to 0 as result of having her type increased).

Similar to single-agent case, we want the solution of OPT3 to be the same as OPT4. The reason why the two solutions do not match is because we do not have any guarantee for virtual valuation of agents (w_i s). Like in the single-agent case, we make the following assumption on the virtual valuation of agents.

Definition 36.5 (Regular Virtual Valuation) *A virtual valuation w_i is regular if $\forall s_i, t_i \in T_i$ with $s_i < t_i$, then $w_i(s_i) < w_i(t_i)$.*

This assumption regarding virtual valuations is weaker than the MHR condition.

Lemma 36.6 *If the hazard rate $\frac{g_i(t_i)}{1-G_i(t_i)}$ is non-decreasing, then w_i is regular.*

Assumption 36.7 *The virtual valuations of all the agents are regular, i.e., w_i is regular $\forall i \in N$.*

This assumption addresses the issue with virtual valuations mentioned earlier. The implication of this assumption is summarized in the following lemma.

Lemma 36.8 *If every agent's virtual valuation is **regular**, then the solution of the constrained optimization problem is the same as the unconstrained problem.*

Lecture 37: November 10, 2017

Lecturer: Swaprava Nath

Scribe(s): Gundeep Arora

Disclaimer: These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.

37.1 Single object allocation with multiple bidders

Continuing our discussion on revenue optimization in single object allocation from the previous lecture, in which we had devised an optimization problem and characterized that in order to ensure the mechanism to be BIC, f had to be non-decreasing in expectation (NDE). The problem was formulated as:

$$\text{OPT1 : } \max_f \int_T \left(\sum_{i \in N} w_i(t_i) \cdot f_i(t) \right) g(t) dt \quad (37.1)$$

s.t. f is NDE

where $w_i(t)$ is the virtual valuation function, $f_i(t)$ is the allocation function for the type profile t and $g(t)$ is the prior on the type profile of agents. We also proved that if we impose the “regularity” condition on the virtual valuation function, the unrestricted version of the optimization problem above is maximized by an allocation function f^* that is NDE. The virtual valuation w_i was defined as regular if it was an increasing function. We shall now prove the lemma that states the above argument formally.

Lemma 37.1 *Suppose every agent’s virtual valuations are regular. The solution of the constrained problem is same as the unconstrained problem.*

Proof: To obtain a solution for the unconstrained optimization problem, we observe that,

$$\sum_{i \in N} w_i(t_i) \cdot f_i(t) \quad (37.2)$$

is a convex combination of the virtual valuations weighted by the allocation function. This implies that the optimal solution would place all the weight on the $\{i \in N : w_i(t_i) \geq w_j(t_j) \forall j \in N\}$, with ties broken arbitrarily. We have to consider the fact that the virtual valuations can be negative and hence if

$$\begin{aligned} w_i(t) < 0 \quad \forall i \in N &\implies f_i(t) = 0 \quad \forall i \in N \\ \text{else } f_i(t) &= \begin{cases} 1 & w_i(t_i) \geq w_j(t_j) \quad \forall j \in N \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (37.3)$$

We claim that f is non-decreasing. To see this, fix t_{-i} . If $t_i > s_i$, by regularity $w_i(t_i) > w_i(s_i)$. Since the allocation probability can only increase when the virtual valuation increases in the above allocation rule, we conclude that

$$f_i(t_i, t_{-i}) \geq f_i(s_i, t_{-i}) \quad t_i > s_i. \quad (37.4)$$

This proves that the given allocation function is non-decreasing. We had set out to prove that f was NDE, which was a weaker requirement. This concludes the proof. ■

Remark 37.2 f is DSIC since it is non-decreasing.

37.1.1 Payment Rule

We now look at the payment rule in this allocation. By Myerson's characterization and the fact that this is an optimal mechanism, i.e., IR and revenue maximizing, the allocation completely determines the payment. Define

$$\kappa_i^*(t_{-i}) = \inf\{t_i : f_i(t_i, t_{-i}) = 1\} \quad (37.5)$$

The above function lists down the minimum value type/bid of agent i , that ensures allocation to her. It is important to note that it includes the condition that the virtual valuation function of the agent be positive, for the allocation to take place. More formally $\kappa_i^*(t_{-i}) \geq w_i^{-1}(0)$. The payment is given by

$$\begin{aligned} p_i(t) &= t_i f_i(t) - \int_0^{t_i} f_i(x_i, t_{-i}) dx_i \\ &= \kappa_i^*(t_{-i}) f_i(t) \end{aligned} \quad (37.6)$$

Theorem 37.3 *Let virtual valuations be regular. For every type t*

$$\begin{aligned} \text{if } w_i(t) < 0, \forall i \in N &\implies f_i(t) = 0, \forall i \in N \\ \text{else } f_i(t) &= \begin{cases} 1 & w_i(t_i) \geq w_j(t_j) \quad \forall j \in N \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (37.7)$$

The payment is given by, $p_i(t) = \kappa_i^(t_{-i}) f_i(t)$. Then (f, p) is a revenue optimal mechanism.*

While we wanted to obtain an optimal mechanism that is BIC, IIR and randomized, thereby expanding the space of mechanisms, the one we just proved to be optimal is DSIC, IR and deterministic. However, the parameters of these DSIC, deterministic mechanisms need the knowledge of priors (the computation of w_i 's need the priors). This is different from the Vickrey auction which never uses the prior. Such mechanisms are called prior-free mechanisms in literature.

37.2 Some Examples

Let us now look at some examples to illustrate the allocation and payment rules.

1. Consider an auction with two agents (\mathcal{A}_1 and \mathcal{A}_2), with uniform priors over type spaces $\mathcal{T}_1 = [0, 12]$, $\mathcal{T}_2 = [0, 18]$. The virtual valuations are given by,

$$\begin{aligned} w_1(t_1) &= t_1 - \frac{1 - G_1(t_1)}{g(t_1)} \\ &= t_1 - \frac{1 - \frac{t_1}{12}}{\frac{1}{12}} \\ &= 2t_1 - 12 \end{aligned} \quad (37.8)$$

$$\text{similarly, } w_2(t_2) = 2t_2 - 18$$

where $G_1(t_1)$ is the cumulative distribution function of the uniform distribution prior of \mathcal{A}_1 . Now, we shall take some values of the bids t_1, t_2 and look at the allocation function and the payment made by the agent in each case.

t_1	t_2	Allocated to	p_1	p_2	$w_1(t_1)$	$w_2(t_2)$
4	8	unsold	0	0	-4	-2
2	12	\mathcal{A}_2	0	9	-8	6
6	6	\mathcal{A}_1	6	0	0	-6
9	9	\mathcal{A}_1	6	0	6	0
8	15	\mathcal{A}_2	0	11	4	12

2. **Symmetric bidders:** As visible from the previous example, if the bidders were symmetric, i.e., they had the same prior distributions and hence the same CDF, their virtual valuation functions would have been the same.

$$\begin{aligned}
 T_i &= T \quad \forall i \in N \\
 G_i &= G \quad \forall i \in N \\
 \Rightarrow w_i &= w
 \end{aligned} \tag{37.9}$$

Now since, $w(t_i) > w(t_j) \iff t_i > t_j$, the object is allocated to the agent with highest valuation (and unsold if everyone has negative value for the object). The payment is given by

$$p_i(t) = \max\{w^{-1}(0), \max_{j \neq i} t_j\} \tag{37.10}$$

This is exactly the second price auction with a reserve price.

3. **Efficiency and optimality:** Efficiency requires to maximize the social welfare, while optimal mechanism requires to maximize the revenue. We now consider the impact revenue optimization has on the efficiency of the allocation function. Consider an auction with two agents, \mathcal{A}_1 , \mathcal{A}_2 with uniform priors given as, $T_1 = [0, 10]$ and $T_2 = [0, 6]$. The virtual valuation function evaluates to,

$$\begin{aligned}
 w_1(t_1) &= 2t_1 - 10 \\
 w_2(t_2) &= 2t_2 - 6 \\
 2t_1 - 10 &= 2t_2 - 6 \quad \text{Decision boundary for the winning agent} \\
 t_2 &= t_1 - 2
 \end{aligned} \tag{37.11}$$

Now, in the figure presented, there is a region ($t_1 < 5, t_2 < 3$) where the product is not allocated to anyone, even though the types/bids of the agents are positive. This is not efficient.

Efficiency would require the allocation decision boundary to be the line $t_1 = t_2$ (allocate to \mathcal{A}_1 if $t_1 > t_2$), however, for the allocation rule in the optimal mechanism the decision boundary is the line $t_2 = t_1 - 2$, making the allocation in favor of \mathcal{A}_2 even though $t_1 > t_2$ (the shaded area in the figure), which is inefficient.




Figure 37.1: A plot of the bid and the allocation decision boundary of the optimal mechanism.

Lecture 38: November 14, 2017

Lecturer: Swaprava Nath

Scribe(s): Gundeep Arora

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor at swaprava@cse.iitk.ac.in.*

38.1 Efficiency and Budget Balance

In this lecture, we will look at some selected results and their proof sketches that are related to some of the concepts we looked at during the course. The first thing we look at is the interplay of efficiency and budget balance in the quasi-linear setting.

The following result shows that if we want efficiency in ‘rich enough’ quasi-linear setting, the only class of payment rules that implements the efficient SCF is the Groves class of payments.

Recall that the Groves class of payment rules is given by

$$p_i(t) = h_i(t_{-i}) - \sum_{j \neq i} t_j(f^*(t)) \quad (38.1)$$

where $f^*(t) \in \arg \max_{a \in A} \sum_{i=1}^n t_i(a)$

Theorem 38.1 (Green and Laffont (1979), Holmström (1979)) *If the type space is sufficiently rich, every efficient and DSIC mechanism is a Groves mechanism.*

Proof sketch: [this exposition is due to Holmström (1979), see the paper for a complete treatment] Consider for simplicity that there are only two allocations $\mathcal{A} = \{a, b\}$. Let (f, p) be a DSIC and efficient mechanism. The valuation for agent i for the allocation a is given by $t_i(a)$. Hence the social welfare for these two allocations is $\sum_{i \in N} t_i(a)$ and $\sum_{i \in N} t_i(b)$. An efficient allocation rule maximizes the social welfare, hence if $\sum_{i \in N} t_i(a) > \sum_{i \in N} t_i(b)$, it will choose the allocation a . Let us now

- fix the valuations of all agents except i , i.e., t_{-i} is fixed, and
- fix the valuation of i at all the other allocations, i.e., $b, t_i(b)$.

We only change the valuation of agent i for allocation a keeping all the other terms of the social welfare fixed to observe what impact it has on the allocation rule. Clearly, \exists a threshold $t_i^*(a)$, such that

$$\begin{aligned} \forall t_i(a) \geq t_i^*(a) & \quad \text{outcome is } a \\ \forall t_i(a) < t_i^*(a) & \quad \text{outcome is } b \end{aligned} \quad (38.2)$$

We know that if a mechanism is DSIC, and by changing the valuation of an agent if the allocation remains the same, the payment should not change for that agent. So, without loss of generality, payment of an agent can be denoted as a function of the allocation and the types of other agents. Since, here the types of all

agents except i remains fixed, whenever the outcome is a , the payment can be denoted by p_{ia} . Consider a valuation $t_i^*(a) + \epsilon$, $\epsilon > 0$. By definition of DSIC

$$t_i^*(a) + \epsilon - p_{ia} \geq t_i(b) - p_{ib} \quad (38.3)$$

The RHS denotes the utility when agent i underreports his type below $t_i^*(a)$ and the allocation changes to b . Similarly, when the true type of i is $t_i^*a - \delta$, $\delta > 0$, DSIC implies

$$t_i^*(a) - \delta - p_{ia} \leq t_i(b) - p_{ib} \quad (38.4)$$

Here the LHS denotes the utility of agent i when he overreports his type and the allocation becomes a . Now since ϵ, δ are arbitrary, we take limits of them tending to zero, which gives

$$t_i^*(a) - p_{ia} = t_i(b) - p_{ib} \quad (38.5)$$

Now, by the definition of $t_i^*(a)$, it is the threshold where a starts becoming the efficient outcome. Hence the following balance equation holds.

$$\begin{aligned} t_i^*(a) + \sum_{j \neq i} t_j(a) &= t_i(b) + \sum_{j \neq i} t_j(b) \\ \implies p_{ia} - p_{ib} &= \sum_{j \neq i} t_j(b) - \sum_{j \neq i} t_j(a) && \text{(substituting Eq. 38.5)} \\ \implies p_{ix} &= h_i(t_{-i}) - \sum_{j \neq i} t_j(x) \end{aligned}$$

The last implication holds as we have collected all the terms independent of t_i and defined that to be $h_i(t_{-i})$. Note that the difference in payment at two allocations hold for every agent and therefore the payment can only have an additive term $h_i(t_{-i})$ which is independent of t_i . This payment is same as the Groves payment rule. The proof can be extended to any finite number of allocations. ■

Theorem 38.2 (Green and Laffont (1979)) *No Groves mechanism is budget balanced, i.e., $\nexists p_i^G$ such that $\sum_{i \in N} p_i^G(t) = 0, \forall t \in T$.*

Proof sketch:[see Green and Laffont (1979) for the complete proof] We outline the proof idea for two agents and two allocations $\{0, 1\}$ in a public project model, where 0 implies that the project is not undertaken and 1 implies that it is undertaken. For the case when allocation 0 is chosen, all agents have zero value.

For a contradiction, suppose $\exists h_i$ s.t. $\sum_{i \in N} p_i(t) = 0, \forall t \in T$. In the Groves class, the only flexibility we have is in the choice of h_i as the rest is fixed. Consider three numbers, w_1^+, w_1^-, w_2 , where w_1^+, w_1^- are valuations of agent 1 and w_2 is the valuation of agent 2, such that

$$\begin{aligned} w_1^- + w_2 &< 0 && \text{outcome is 0} \\ w_1^+ + w_2 &> 0 && \text{outcome is 1} \end{aligned} \quad (38.6)$$

Now Groves payment at (w_1^+, w_2) satisfies

$$h_1(w_2) - w_2 + h_2(w_1^+) - w_1^+ = 0 \quad \text{(project is undertaken and equality by BB assumption)}$$

Similarly

$$h_1(w_2) + h_2(w_1^-) = 0 \quad \text{(project is not undertaken and equality by BB assumption)}$$

Subtracting one from the other, we get

$$w_2 = h_1(w_1^+) - h_2(w_1^-) - w_1^+.$$

Now the RHS of the above equation is completely independent of w_2 . For any small change in w_2 such that the inequalities of Eq. 38.6 continues to hold, clearly the above equality cannot hold. This is a contradiction. ■ The two results can be summarized in the form of the following corollary.

Corollary 38.3 *If the valuation space is sufficiently rich, no efficient mechanism can be both DSIC and BB.*

38.2 Weakening DSIC for positive results

However, if we weaken the IC notion to Bayesian, we can have positive results. The following mechanism is due to d'Aspremont and Gerard-Varet (1979), Arrow (1979), and is called **dAGVA** mechanism. Under this mechanism, the allocation is still the efficient one. Payment is defined via priors. Define

$$\begin{aligned} \delta_i(t_i) &= \mathbb{E}_{t_{-i}|t_i} \sum_{j \neq i} t_j(a^*(t)) \\ a^*(t) &\in \arg \max_{a \in A} \sum_{i \in N} t_i(a) \end{aligned} \quad (38.7)$$

Payment of the mechanism is given by

$$p_i^{\text{dAGVA}}(t) = \frac{1}{n-1} \sum_{j \neq i} \delta_j(t_j) - \delta_i(t_i) \quad (38.8)$$

This payment implements the efficient allocation rule in the Bayesian Nash equilibrium.

$$\begin{aligned} &\mathbb{E}_{t_{-i}|t_i} [t_i(a^*(t)) - p_i^{\text{dAGVA}}(t)] \\ &= \mathbb{E}_{t_{-i}|t_i} \sum_{j \in N} t_j(a^*(t)) - \mathbb{E}_{t_{-i}|t_i} \left[\frac{1}{n-1} \sum_{j \neq i} \delta_j(t_j) \right] \\ &\geq \mathbb{E}_{t_{-i}|t_i} \sum_{i \in N} t_j(a^*(t'_i, t_{-i})) - \mathbb{E}_{t_{-i}|t_i} \left[\frac{1}{n-1} \sum_{j \neq i} \delta_j(t_j) \right] \\ &= \mathbb{E}_{t_{-i}|t_i} [t_i(a^*(t'_i, t_{-i})) - p_i^{\text{dAGVA}}(t'_i, t_{-i})] \end{aligned}$$

The first inequality holds by definition of a^* , and the rest of the equalities are obtained by reorganizing the payment expression. Also the sum of the payment of all the agents is given by

$$\begin{aligned} \sum_{i \in N} p_i^{\text{dAGVA}}(t) &= \frac{1}{n-1} \sum_{i \in N} \sum_{j \neq i} \delta_j(t_j) - \sum_{i \in N} \delta_i(t_i) \\ &= \frac{n-1}{n-1} \sum_{j \in N} \delta_j(t_j) - \sum_{i \in N} \delta_i(t_i) = 0 \end{aligned}$$

Hence we have proved the following theorem.

Theorem 38.4 *The dAGVA mechanism is efficient, BIC and BB.*

But the dAGVA does not guarantee IIR. Turns out that the above properties along with IIR is impossible to satisfy even in bilateral trading problem, i.e., there is one buyer and one seller and one object, efficient trade happens when the sellers cost is below the valuation of the buyer.

Theorem 38.5 (Myerson and Satterthwaite (1983)) *In the bilateral trading problem, there is no mechanism that BIC, efficient, IIR and BB.*

38.3 Summary

We will now summarize the mechanism design space we discussed in this course via the following figures. The LHS denotes the preference domains and the RHS denotes the space of mechanisms – the arrows correspond to the necessity of a DSIC/BIC mechanisms to be in a certain class of mechanisms.

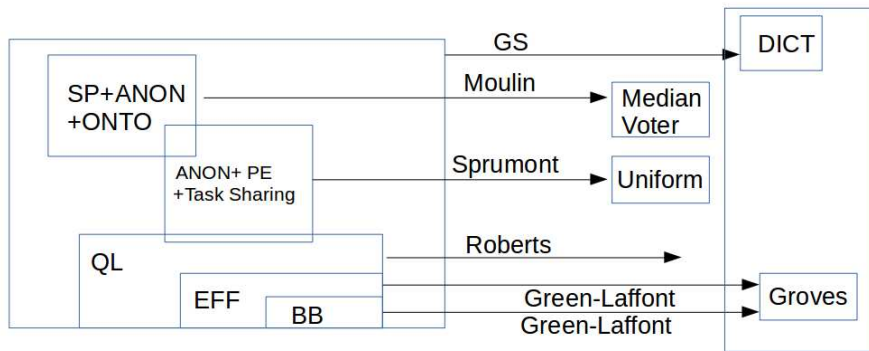


Figure 38.1: The map from valuation type to mechanism type for DSIC mechanisms only (the last arrow goes to \emptyset)




Figure 38.2: The map from valuation type to mechanism type for BIC mechanisms only

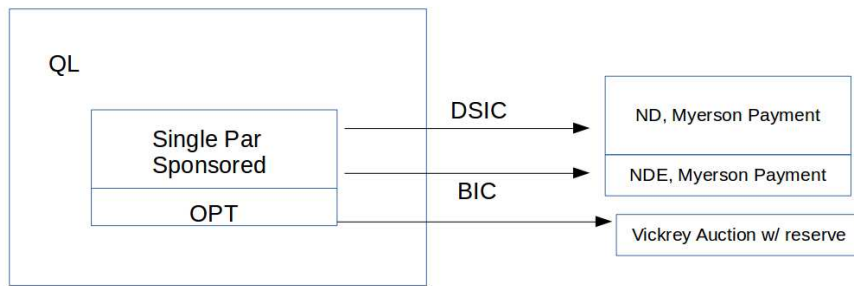


Figure 38.3: The map from valuation type to mechanism property