Social Networks and Online Markets

Homework 1

Due: 21/6/2022, 23:59

Instructions

You must hand in the homeworks electronically and before the due date and time.

The first homework has to be done by each **person individually**.

Handing in: You must hand in the homeworks by the due date and time by an email to `birbas@diag.uniroma1.it` that will contain as attachment **(not links to some file-uploading server!)** a .zip or .pdf file with your answers.
After you submit, you will receive an acknowledgement email that your homework has been received and at what date and time. If you have not received an acknowledgement email within 2 days after the deadline then contact Georgios.

The solutions for the theoretical exercises must contain your answers either typed up or hand written clearly and scanned.

For questions you can email Georgios Birmpas: `birbas@diag.uniroma1.it`.

For information about collaboration, and about being late check the web page.

Problem 1. Consider the following setting that we have also seen in the class: We have a set N of n agents, and a set M of m items. Each agent i, has an additive valuation function $v_i(\cdot)$ over the items. Our goal is to produce an allocation $\mathcal{A} = (A_1, \ldots, A_n)$ of the items to the agent. Recall that an allocation is basically a partition of the set of items, i.e., $\forall i,j \in \{1, 2, ..., n\}$ where, $i \neq j$, we have that $A_i \cap A_j = \emptyset$, and in addition $\cup_{\{1,2,...,n\}} A_i = M$.

The designer, desires a mechanism that produces *fair* allocations. So he thinks of the following fairness concepts:

Definition 1 *An allocation $\mathcal{A} = (A_1, \ldots, A_n)$ is envy-free, if for every $i, j \in N$, $v_i(A_i) \geq v_i(A_j)$.*

Definition 2 *An allocation $\mathcal{A} = (A_1, \ldots, A_n)$ is an EFX (envy-free up to any good) allocation, if $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ holds for every pair $i, j \in N$, with $A_j \neq \emptyset$, and for every $g \in A_j$.*

Definition 3 *An allocation $\mathcal{A} = (A_1, \ldots, A_n)$ is an EF1 (envy-free up to one good) allocation, if for every pair of agents $i, j \in N$, with $A_j \neq \emptyset$, there exists an item $g \in A_j$, such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$.*

The idea that he has for designing a fair mechanism is based on the notion of the *envied* agent. An agent i is envied by an agent j, if agent j has more value for the bundle of items of agent i than he has for his own. With that in mind he proposes the following mechanism:

1. The agents are ordered in an arbitrary way
2. The items are ordered in an arbitrary way
3. For every agent i, set $A_i = \emptyset$
4. Find an agent i that is not envied by anyone, and give him the item with the lowest index from the set of available ones (as long as there is one). This item will be included now to A_i and will no longer be considered available.

5. If every agent is envied by someone, then this means that there are cycles of envy. I.e., there are cycles of agents $k_1 \to k_2 \to k_3 \to ... \to k_l \to k_1$, where $\forall i \in [l]$ we have that $v_{k_i}(A_{k_i}) < v_{k_i}(A_{k_{i+1}})$ (and the last agent in the sequence envies the first). Such cycles can be eliminated by reallocating the current bundles in the following manner: $\forall i \in [l]$, give to agent k_i the bundle of items of agent k_{i+1} (the bundle of the agent that he envies). By doing so, the value of each agent involved in the cycles increases, something that guarantees that cycles will be eliminated in the end (as the value of an agent cannot increase indefinitely), and thus there will be at least one agent that is not envied by anyone after the re-allocations. Therefore, in this step we eliminate the cycles and we return to step 4.

6. Return the final allocation.

Prove that:

1. This mechanism does not always produce envy-free allocations.

2. This mechanism does not always produce EFX allocations.

3. This mechanism always produces EF1 allocations.

4. Consider instances where all the agents have the same ordinal preference over the items (i.e., if they order the items according to how they value them, all of them have the same ordering). Modify the above mechanism and show that for such instances it actually produces EFX allocations. *(Hint. Modify the second step of the mechanism, i.e., the way the items are ordered.)*

Problem 2. Consider the following problem: There is a set of N players and one single digital good. Each agent has a value for the digital good that describes how much he wants it. Since the good is digital, it can be provided to any number of players. However, there is a cost for the provision of the good that is equal to 1 (this is the same regardless of the number of the players that get the good in the end).

The designer wants to design a mechanism that is truthful, covers the cost created due to the provision of the good, and is also economically efficient. So he thinks as follows: Let the players declare their values, and then set the price of the good to $p = \frac{1}{n}$. If the value of **every** player is more than $\frac{1}{n}$, then everyone will get the item and the procedure terminates. However, if there are some players, say k in number, that have value less than $\frac{1}{n}$, then they are removed from the game and they do not get the digital good. If something like this happen, the price of the digital good is updated to $p = \frac{1}{n-k}$, and the same procedure is followed again with the new price and the new set of players (the set of players that remain). This goes on, until either we end up with a set of players where everyone has a value higher than the current price (this will be the set of the winners), or we end up with an empty set of players. The designer calls this mechanism, the *sharing mechanism.*

If a player is a winner (gets the digital good) then his utility is $u_i = v_i - p$, while if he is removed from the game (he does not get the item) his utility is $u_i = 0$. Finally, as mentioned in the beginning, the mechanism has to be efficient, so the objective in this case is to minimize the harm to the society which is defined as follows: $HS = C(S) + \sum_{i \in N \setminus S} v_i$, where $C(S)$ is the cost of the provision of the good to set S (the set of winners). In case $S = \emptyset$, then $C(S) = 0$.

1. Prove that the allocation rule of the sharing mechanism is monotone, and that the sharing mechanism is truthful.

2. Prove that the sharing mechanism cannot provide an approximation to the harm to society that is better than $\log |N|$.

3. Consider the version of the problem where now there is a set K of k digital goods, and the provision of **each of them** has a cost that is equal to 1. In this version of the problem, each agent i has now a valuation function that is defined as follows: $v_i(R) = \min\{\sum_{j \in R} v_i(j), B_i\}$, i.e., the value that agent i has for a set $R \subseteq K$ of digital goods that he wins, is the minimum between the sum of his values for the digital goods in R, and a positive number B_i. In this scenario, the designer decides to do the following: He asks every agent to submit the value that they have for each digital good, and then runs the sharing mechanism for each digital good separately. Prove that under this procedure, it is not always the best strategy for an agent to truthfully report his values. *(Hint: This can be shown with a simple example of instances with only 2 agents and 2 digital goods.)*

Problem 3. Consider the following games:

- We have a set N of n agents and a set M of m of companies that resolve tasks. Each agent $i \in N$ has a task of weight w_i that wants to complete. To do so, he can choose a company, and assign his task to it. Now each company is associated with a function that denotes its load of work, given the tasks that it has to complete. In particular, given a strategy profile $\mathbf{s} = (s_1, s_2, \dots, s_n)$ with the strategies of the agents, the load function of company j is $l_j(\mathbf{s}) = \sum_{i:s_i=j} w_i$. On the other hand, each agent i is associated with a cost that depends on the load of the company that he chooses. Specifically, given the strategy profile $\mathbf{s}$ of the agents, the cost of agent i is $c_i(\mathbf{s}) = l_{s_i}(\mathbf{s})$. The goal of each agent i is to minimize his cost. Prove that
$$\phi(\mathbf{s}) = \frac{1}{2} \sum_{j=1}^{m} l_j(\mathbf{s})^2$$
is a potential function, thus this game has a pure Nash equilibrium.

- Consider the following graph game. Given a graph $G = (V, E)$, each agent i is assigned to a node, and each edge $e = (i, j)$ has a weight w_e. An agent i can choose between two strategies, in particular for every i, we have $s_i \in \{-1,1\}$. Finally, given a strategy profile $\mathbf{s} = (s_1, s_2, \dots, s_n)$, the utility of agent i is defined as $u_i = \sum_{e=(i,j) \in E} w_e \cdot s_i \cdot s_j$. Prove that
$$\phi(\mathbf{s}) = \sum_{j<i \in V} w_e \cdot s_i \cdot s_j$$
is a potential function, thus this game has a pure Nash equilibrium.

Problem 4. Answer the following questions:

- Give examples of preference profiles for which Borda and Veto can be manipulated. Explain how the manipulation results in a preferable outcome for the voter who misreports her preference ordering.

- Consider the cardinal social choice setting, in which there is a set of voters N, a set of candidates M (possibly infinite) and each voter has a valuation function $v_i : M \to \mathbb{R}$ assigning a numerical score to the candidates. Assume that a mechanism takes the numerical scores that the agents declare, and outputs a single candidate as the winner. Show that the Gibbard-Satterthwaite theorem extends to the cardinal social choice setting (*Hint: Show that every deterministic truthful mechanism is ordinal, i.e. it only uses the orderings induced by the valuation functions.*)

Problem 5. Consider the following setting. A set of n people in a room are trying to decide the temperature of a conference room. Each person i has a most preferred temperature t_i and her displeasure from a chosen temperature t is defined as $c_i = |t - t_i|$, i.e. her displeasure increases linearly as the chosen temperature moves away from her ideal temperature.

We consider two objectives, the *total displeasure* $\sum_{i=1}^{n} c_i$ and the *maximum displeasure* $\max_{i=1}^{n} c_i$ that we are trying to minimize.

- For each objective, what is the temperature (with respect to the variables t_i) that minimizes the objective?
- Prove that the approximation ratio of Dictatorship is at least $n - 1$ for the total displeasure and at least 2 for the maximum displeasure.
- Give a deterministic truthful mechanism (social choice function) that achieves an approximation ratio of 1 for the total displeasure. Your answer should prove that the mechanism is truthful and that it guarantees an optimal outcome always.
- Prove that no deterministic truthful mechanism can have an approximation ratio smaller than 2 for the maximum cost (*Hint:* Use a profile with only two people and argue by contradiction that such a truthful mechanism exists. This implies something for the choice of the temperature. Consider a deviation of one of the two people to that chosen temperature and argue that it would violate truthfulness).
- Consider the following randomized mechanism: With probability 1/4 output the smallest among most preferred temperatures t_ℓ, with probability 1/4 output the largest among most preferred temperatures t_r and with probability 1/2 output the $(t_r - t_\ell)/2$.
  - Show that the mechanism is *truthful-in-expectation*, i.e. no person can decrease her expected displeasure by misreporting her true preferred temperature.
  - Show that that the expected maximum displeasure of the mechanism is within 3/2 of the optimal maximum displeasure.