15896/15483 Truth, Justice, and Algorithms

Fall 2018 Procaccia and Psomas HW 4

Due on November 19 at 11:59PM

Grading For this homework, you will let me know how to grade your assignment. The mechanism by which you will do so is as follows: you must allocate a budget of 100 faux-dollars among the four problems. The amount that you bid on each problem will determine the weight of that problem, where the weight of a question is the square root of the amount of money you bid on that problem. Your overall score for each question is the product of the fraction of the question you got correct times the weight of that question. Therefore, the maximum score for this homework assignment is 20 points, which is obtained if you split your budget equally among the four problems and get all of them completely correct. Final scores out of 20 will be scaled to be out of 100.

For instance, let's say that you bid 10 dollars on question 1, 20 dollars on question 2, 30 dollars on question 3, and 40 dollars on question 4. If you got 20/25 points on question 1, 15/25 points on question 2, 25/25 points on question 3, and 10/25 points on question 4, your final score is $(20/25)\sqrt{10}+(15/25)\sqrt{20}+(25/25)\sqrt{30}+(10/25)\sqrt{40}=13.22/20=66.10/100$.

Please specify your bids very clearly in your writeup. If no bids are specified, each question will be weighted equally, as usual.

1. **The Gibbard-Satterthwaite Theorem** (25 points: 10/5/10)

   We saw in class a proof sketch of the Gibbard-Satterthwaite Theorem for the special case of strategyproof and neutral voting rules with $m \geq 3$ and $m \geq n$. That proof relied on two key lemmas of strong monotonicity and Pareto optimality. In this problem, you will prove the two lemmas and formalize the theorem's proof for this special case.

   Prove the following statements.

   (a) **Strong Monotonicity**: Let f be a strategyproof voting rule, $\vec{\sigma}=(\sigma_1,\ldots,\sigma_n)$ be a preference profile, and $f(\vec{\sigma})=a$. If $\vec{\sigma}'$ is a profile such that $[a \succ_i x \Rightarrow a \succ'_i x]$ for all $x \in A$ and $i \in N$, then $f(\vec{\sigma}')=a$.

   (b) **Pareto Optimality**: Let f be a strategyproof and onto voting rule. Furthermore, let $\vec{\sigma}=(\sigma_1,\ldots,\sigma_n)$ be a preference profile and $a,b \in A$ such that $a \succ_i b$ for all $i \in N$. Then $f(\vec{\sigma}) \neq b$.
   Hint: use part (a).

   (c) Let m be the number of alternatives and n be the number of voters, and assume that $m \geq 3$ and $m \geq n$. Furthermore, let f be a strategyproof and neutral voting rule. Then f is dictatorial.
   Important note: There are many proofs of the full version of the Gibbard-Satterthwaite Theorem; *here the task is specifically to formalize the proof sketch we did in class*.

   Answer:

(a) Suppose first that $\sigma_i = \sigma_i'$ for $i > 1$. Let $f(\sigma_1', \sigma_{-1}) = b$. From strategyproofness it follows that $a \succeq_1 b$, and hence from the assumption of the lemma, $a \succeq_1' b$. Strategyproofness also implies that $b \succeq_1' a$, and because preferences are strict it follows that $a = b$. The lemma now follows after repeating this argument while changing the preferences for only $i = 2$, then $i = 3$, etc.

(b) Suppose that $f(\vec{\sigma}) = b$. Since f is onto there is a profile $\vec{\sigma}'$ such that $f(\vec{\sigma}') = a$. Let $\vec{\sigma}''$ be such that for all $i \in N$, $a \succ_i'' b \succ_i'' x$ for all $x \in A \setminus \{a, b\}$, and the rest of the alternatives are ranked identically to σ_i. By monotonicity, $b = f(\vec{\sigma}) = f(\vec{\sigma}'')$ and $a = f(\vec{\sigma}') = f(\vec{\sigma}'')$, which is a contradiction. Hence $f(\vec{\sigma}) \neq b$.

(c) For this part of the proof it is convenient to define the preferences of each $i \in N$ via a utility function u_i such that for $x, y \in A$, $x \succ_i y$ if and only if $u_i(x) > u_i(y)$. Therefore, $f(\vec{u})$ is well defined. We will also denote $A = \{a_1, \ldots, a_m\}$.

For each $i \in N$, let

$$u_i(a_j) = \begin{cases} n + i - j & i \leq j \leq n \\ i - j & j < i \\ n - j & j > n \end{cases}$$

That is, the ranking of $a_1, \ldots, a_n$ is shifted, and all other alternatives are ranked below them. By Pareto optimality (part b), $f(\vec{u}) = a_j$ for some $j \leq n$. Assume w.l.o.g. that $f(\vec{u}) = a_1$. Let $\vec{u}'$ be defined as follows:

$u_1'(a_1) = n + 2$ and $u_1'(a_n) = n + 1$,
$u_i'(a_n) = n + 2$ and $u_i'(a_1) = n + 1$ for $i > 1$,
$u_i'(a_j) = u_i(a_j)$ otherwise

Hence all voters consider the alternatives a_1 and a_n to be better than the other alternatives. Also note that the ranking of a_1 and a_n is the same in the profiles $\vec{u}$ and $\vec{u}'$; and in $\vec{u}'$, a_1 and a_n are both ranked above other alternatives. Hence by strong monotonicity (part a), $f(\vec{u}') = f(\vec{u}) = a_1$. Finally, define profiles $\vec{u}^k$ for $k = 1, \ldots, n$, where $\vec{u}^1 = \vec{u}'$, and

$$u_i^{k+1}(x) = \begin{cases} u_i^k(x) & i \neq k + 1 \\ u_{k+1}^k(x) & i = k + 1 \text{ and } x \in A \setminus \{a_1\} \\ -m & i = k + 1 \text{ and } x = a_1 \end{cases}$$

By Pareto optimality (part b), $f(\vec{u}^k) \in \{a_1, a_n\}$. But strategyproofness implies that $f(\vec{u}^k) = a_1$, and hence $f(\vec{u}^n) = a_1$. In $\vec{u}^n$, a_1 is ranked at the top by voter 1, and at the bottom by every other voter. Strong monotonicity (part a) implies that a_1 is the winner whenever voter 1 puts a_1 at the top. Neutrality then implies that voter 1 is a dictator.

2. **Distortion** (25 points)

Recall that the participatory budgeting problem as defined in Lecture 16, Slide 13. We consider the special case in which each alternative a has cost $c_a = 1$ (all alternatives have the same (unit) cost), and the total budget $B = 1$. This means that exactly one alternative may be selected.

Under the *threshold approval* input format, each voter approves all alternatives that have utility above a certain threshold (which we choose). The threshold is the same for all voters, and can be selected at random. In principle the aggregation of votes can also be randomized, but deterministic aggregation suffices for this problem.

Prove that in the setting where $c_a = 1$ for all alternatives $a \in A$ and $B = 1$, the distortion associated with threshold approval is $O(\log m)$. Specifically, design a distribution over thresholds and a deterministic aggregation method f, such that for any utility profile, the ratio between the welfare-maximizing solution and the expected social welfare of the outcome under f, where the expectation is taken over the randomness of the threshold, is $O(\log m)$.

Hint: Choose a value $j \in \{1, \dots, \log m\}$ uniformly at random and set a threshold $\ell_j = 2^{j-1}/m$. You can assume for each of exposition that $\log m$ is an integer.

Answer: Consider the rule that chooses a threshold as in the hint and then returns the alternative with the greatest number of approval votes.

Let $\vec{v}$ be the underlying utility profile and let $a^* = \arg\max_{a \in A} \mathrm{sw}(a, \vec{v})$ be the welfare-maximizing alternative. If there exists a threshold ℓ_j that chooses this alternative, then we are done because this happens with probability at least $1/\log m$ and therefore we immediately obtain $O(\log m)$ distortion.

Therefore, assume that there does not exist any threshold under which the rule chooses a^*.

For $a \in A$ and $j \in [\log m]$, let n_j^a denote the number of approval votes a receives when the threshold is $t = \ell_j$, and let $a_j \in A$ be the alternative returned by the rule when $t = \ell_j$. Because our rule returns an alternative with the greatest number of approval votes, we have

$$\forall j \in [\log m], \sum_{k=j}^{\log m} n_k^{a_j} \geq \sum_{k=j}^{\log m} n_k^{a^*} \geq n_j^{a^*}.$$

Now, the expected social welfare achieved by our rule is at least

$$\sum_{j=1}^{\log m} \Pr[t = \ell_j] \cdot \mathrm{sw}(a_j, \vec{v}),$$

which is at least

$$\frac{1}{\log m} \sum_{j=1}^{\log m} \ell_j \left(\sum_{k=j}^{\log m} n_k^{a_j} \right)$$

by the above. Now, define $u_j = 2\ell_j$, which yields

$$\frac{1}{\log m} \sum_{j=1}^{\log m} \ell_j \left(\sum_{k=j}^{\log m} n_k^{a_j} \right) \geq \frac{1}{2\log m} \sum_{j=1}^{\log m} u_j \cdot n_j^{a^*} \geq \frac{1}{2\log m} \cdot \mathrm{sw}(a^*, \vec{v}).$$

Therefore, the distortion of the rule is $O(\log m)$, as desired.

3. **Strategyproof Facility Location** (25 points: 5/10/10)

Consider a facility location game with n agents in which each agent controls k locations, and denote the set of locations that agent i controls as $\vec{x}_i = (x_{i1}, \dots, x_{ik})$. Therefore, the entire location profile is $\vec{x} = (\vec{x}_1, \dots, \vec{x}_n)$.

Let a deterministic mechanism in the multiple locations setting be defined as a function $f : \mathbb{R}^k \times \cdots \times \mathbb{R}^k \to \mathbb{R}$; that is, it takes in a location profile and returns a single location based on all the locations reported by each agent.

The cost of facility location y to an agent i is the sum of distances from y to each of the locations that i controls, or $\text{cost}_i(y, \vec{x}_i) = \sum_{j\in[k]} |y - x_{ij}|$. The social cost of a location y is the sum of costs of each agent for location y:

$$\text{cost}(y, \vec{x}) = \sum_{i\in[n]} \sum_{j\in[k]} |y - x_{ij}|.$$

Consider the following mechanism for the facility location game in the multiple locations setting.

> **MECHANISM 1**
>
> - For each agent i with reported locations $\vec{x}_i = (x_{i1}, \ldots, x_{ik})$, let $\text{med}(\vec{x}_i)$ be the median of these locations.
> - Return the median of $(\text{med}(\vec{x}_1), \ldots, \text{med}(\vec{x}_n))$.

Intuitively, Mechanism 1 creates a new bid for each agent at the median of the locations under its control, and then returns the median of these new bids.

(a) Prove that Mechanism 1 is strategyproof.

(b) Prove that Mechanism 1 is a 3-approximation algorithm for the social cost in the multiple locations setting.

(c) Consider the case of two agents. Prove that for any $\varepsilon > 0$, there exists a k such that any strategyproof deterministic mechanism $f : \mathbb{R}^k \times \mathbb{R}^k \to \mathbb{R}$ cannot have an approximation ratio better than $3 - \varepsilon$ for the social cost in the multiple locations setting.

Answer:

(a) Misreporting agents can only move the final location weakly farther away from their own optimum location. To see this, note that there are two cases for a misreport: the misreport changes the final median, or the misreport does not change the final median. If it does not change the final median, then the agent does not benefit. If it does change the final median, note that the final median can only skew farther away from the empirical minimizer for the agent, as seen in class.

(b) Let the true optimum location be a^* and the location returned by the mechanism be a.

First, we show that the number of total facilities below a is at least $nk/4$ (and, symmetrically, that the number of total facilities above a is at least $nk/4$). The proofs are symmetric; we only provide the former below.

Let $\{y_{ij} : y_{ij} \leq a\}$ be the set of all facilities that are to the left of a. Furthermore, let $\tilde{y}_i$ denote the projected value of agent i. Because Mechanism 1 returns the median of the $\tilde{y}_i$ values, we know that at least half of the individual medians are below a. Further, we know that at least half of each of these individuals' locations are below their median, so we have $|\{y_{ij} : y_{ij} \leq a\}| \geq nk/4$. Symmetrically, we know that $|\{y_{ij} : y_{ij} \geq a\}| \geq nk/4$.

Now, we look at the social cost of the facility location that the mechanism returns. Let $d = |a - a^*|$ and WLOG assume $a < a^*$.

$$
\begin{aligned}
\mathrm{sc}(a,S) &= \frac{1}{nk}\sum_{i,j}|y_{ij}-a| \\
&= \frac{1}{nk}\left(\sum_{i,j:y_{ij}\le a}(a-y_{ij}) + \sum_{i,j:a<y_{ij}\le a^*}(a-y_{ij}) + \sum_{i,j:y_{ij}>a^*}(y_{ij}-a)\right) \\
&\le \frac{1}{nk}\left(\sum_{i,j:y_{ij}\le a}(a-y_{ij}) + \sum_{i,j:a<y_{ij}\le a^*} d + \sum_{i,j:y_{ij}>a^*}(d+(y_{ij}-a^*))\right) \\
&= \frac{1}{nk}\left(\sum_{i,j:y_{ij}\le a}(a-y_{ij}) + |\{i,j: y_{ij}>a\}|d + \sum_{i,j:y_{ij}>a^*}(y_{ij}-a^*)\right) \\
&\le \frac{1}{nk}\left(\sum_{i,j:y_{ij}\le a}(a-y_{ij}) + \frac{3}{4}nkd + \sum_{i,j:y_{ij}>a^*}(y_{ij}-a^*)\right).
\end{aligned}
$$

Proceeding similarly for the social optimum, we have

$$
\begin{aligned}
\mathrm{sc}(a^*,S) &= \frac{1}{nk}\sum_{i,j}|y_{ij}-a^*| \\
&= \frac{1}{nk}\left(\sum_{i,j:y_{ij}\le a}(a-y_{ij}+d) + \sum_{i,j:a<y_{ij}\le a^*}(a^*-y_{ij}) + \sum_{i,j:y_{ij}>a^*}(y_{ij}-a^*)\right) \\
&\ge \frac{1}{nk}\left(\sum_{i,j:y_{ij}\le a}(a-y_{ij}+d) + \sum_{i,j:y_{ij}>a^*}(y_{ij}-a^*)\right) \\
&\ge \frac{1}{nk}\left(\sum_{i,j:y_{ij}\le a}(a-y_{ij}) + \sum_{i,j:y_{ij}>a^*}(y_{ij}-a^*) + \frac{1}{4}nkd\right).
\end{aligned}
$$

Because two of the terms are identical, the bound follows.

(c) Let $k = 2t+1$. Let player 1 have k locations at position 0 and let player 2 have k locations at position 2^q for some integer q. We first show that the outcome of any strategyproof deterministic mechanism must be close to either end. In particular, the a' that any strategyproof deterministic mechanism returns must either be at most $1/2$ or at least $2^q - 1/2$. Further, note that if the mechanism returns something outside the range of $[0, 2^q]$, this results in a bad approximation ratio.

We prove this by induction. For the case of $q = 0$, note that this is exactly the requirement that the median is located somewhere between the two points, which is true.

From here on, we only consider the case where $a \le 1/2$ (the other proceeds by a symmetric argument). If $a \le 1/2$ when player 1 has k locations at position 0 and player 2 has k locations at position 2^q for some q, we would like to show that $a \le 1/2$ or $a \ge 2^{q+1} - 1/2$ when player 1 has k locations at position 0 and player 2 has k locations at position 2^{q+1}. Consider the setting when player 1 remains the same but player 2 has all her locations at 2^{q+1} instead of 2^q (this doubles the distance between player 1's and player 2's points). Call the new result of the mechanism when run on this profile a'. By strategyproofness, we know that $|2^q - a'| \ge |2^q - a|$ because otherwise player 1 has an incentive to misreport. Therefore, we have $a' \le 1/2$ or $a' \ge 2^{q+1} - 1/2$.

Now, we use this construction: agent 1 has k facilities at 0 and agent 2 has k facilities at 2^q. WLOG, $a \leq 1/2$. Consider the setting in which agent 1 has $t+1$ facilities at 0 and t facilities at 2^q, and agent 2 has $2t+1$ facilities at 2^q. By strategyproofness, the mechanism must choose a median that is also $< 1/2$ or else agent 1 has an incentive to misreport. However, the optimal solution is to place a facility at position 2^q. Therefore, the cost of the mechanism is at least $\frac{3t+1}{4t+2} \cdot (2^q - 1/2)$, but the optimal cost is $\frac{t+1}{4t+2} \cdot 2^q$; the ratio of these tends to 3 as t and q tend to infinity.

4. **Kidney Exchange** (25 points)

We proved in class that when there are at least two players, no deterministic strategyproof kidney exchange mechanism can provide an α-approximation for $\alpha < 2$. Show that no randomized strategyproof kidney exchange mechanism can provide an α-approximation for $\alpha < 6/5$.

Answer: Let f be a randomized strategyproof mechanism. Consider the graphs from class $(1-2-2-1-1-1-2)$.

Since G does not have a perfect matching, we know that $u_1(f(G)) + u_2(f(G)) \leq 6$. This means that either $u_1(f(G)) \leq 10/3$ or $u_2(f(G)) \leq 8/3$.

If $u_1(f(G)) \leq 10/3$, consider the graph G' that consists of G where the first player hides its last two vertices (vertices 5 and 6). By strategyproofness, f can only match both of player 1's vertices with probability at most $2/3$, so it can match both edges with probability at most $2/3$ for a maximum of $5/3$ edges in expectation, but the optimum is 2. The ratio is at least $6/5$.

If $u_2(f(G)) \leq 8/3$, then we consider the graph G'' that consists of G where the second player hides its first two vertices (vertices 2 and 3). By strategyproofness, f can only match player 2's vertex with probability at most $2/3$, so it can match both edges with probability at most $2/3$ for a maximum of $5/3$ edges in expectation, but the optimum is again 2. The ratio is at least $6/5$.