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.*

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.

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 $\vec{x}_i = (x_{i1}, \dots, x_{ik})$, let $\text{med}(\vec{x}_i)$ be the median of these locations.
- Return the median of $(\text{med}(\vec{x}_1), \dots, \text{med}(\vec{x}_n))$.

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.

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 $\alpha < 6/5$.