AGTA Tutorial 6

Lecturer: Aris Filos-Ratsikas Tutor: Charalampos Kokkalis

March 13, 2025

Exercise 1. Consider the *Plurality* voting rule: For every candidate $c \in A$, the candidate receives one point for each voter that ranks c in the first position of its ranking. The winner is the candidate with the most points, breaking ties arbitrarily.

Provide an example showing that *Plurality* is not truthful. The example should show the misreport of a voter that results in that voter receiving higher utility.

Exercise 2. 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.

A. Show that the cardinal social choice setting is a generalization of the social choice setting that we discussed in class, when we assume that the valuation functions v_i are injective, i.e., the valuations do not display ties.

B. Show that the Gibbard-Satterthwaite theorem extends to the cardinal social choice setting above (*Hint: Show that every deterministic truthful mechanism is ordinal, i.e. it only uses the orderings induced by the valuation functions.*)

Exercise 3. Prove the following statement: A social choice rule f is Pareto optimal if it is truthful and onto.

Exercise 4 (Facility Location)**.** Consider the facility location problem that we defined in the lectures. There is a set of n agents and one facility to be built on the real line $\mathbb{R}$. Each agent i has a most preferred location x_i (a *peak*), and its cost from a chosen location y is defined as $c_i = |y - x_i|$, i.e. its cost increases linearly as the chosen facility location moves away from its peak.

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

A. For each objective, what is the location (with respect to the variables x_i) that minimizes the objective? *(already discussed in the lectures).*

B. Prove that the approximation ratio of Dictatorship is at least $n - 1$ for the social cost and at least 2 for the maximum cost.

C. Give a deterministic truthful mechanism (social choice function) that achieves an approximation ratio of 1 for the social cost. Your answer should prove that the mechanism is truthful and that it guarantees an optimal outcome always. *(already discussed in the lectures).*

D. Prove that no deterministic truthful mechanism can have an approximation ratio smaller than 2 for the maximum cost *(already discussed in the lectures).*

E. Consider the following randomized mechanism: with probability 1/4 output the smallest among most preferred locations x_ℓ, with probability 1/4 output the largest among most preferred locations x_r and with probability 1/2 output $(x_r - x_\ell)/2$.

- Show that the mechanism is *truthful-in-expectation*, i.e. no person can decrease her expected cost by misreporting her true preferred temperature.
- Show that that the expected maximum cost of the mechanism is within 3/2 of the optimal maximum cost.

Exercise 5 (More advanced, optional)**.** As we discussed in the lectures, one of the ways to escape the implications of the Gibbard-Satterthwaite theorem is to use *randomisation*: an (ordinal) randomised voting rule is a function $f : (\succ)^n \to \Delta(A)$, where $(\succ)^n$ is the set of all possible preference profiles with n voters and $\Delta(A)$ is the set of probability distributions over the set of candidates A.

For randomised voting rules, the Gibbard-Sattertwaite theorem does not apply. There is however another similar theorem that applies, due to Gibbard [1], refined and presented below in a more convenient form due to [2, 3].

Theorem 1. A truthful (in expectation) randomised voting rule for n voters and m candidates is truthful, anonymous, and neutral, if and only if it is a convex combination (a probability mixture) over voting rules in the following two classes:

- Rules $U^q_{m,n}$: Select a voter uniformly at random from the n voters. Then, out of the m candidates, select the q candidate that this voter ranks at the top with probability $1/q$ each. These are called *mixed unilaterals*.
- Rules $D^q_{m,n}$: Select two out of the m candidates a and b uniformly at random and eliminate all the other candidates. Then, if at least q out of the n voters prefer a to b, then select a, otherwise select either a or b uniformly at random. These are called *mixed duples*.

Informally, a voting rule is anonymous if the identities of the voters do not matter for the outcome, and neutral if the identities of the candidates do not matter for the outcome. It is not important for this exercise to know the formal definitions.

A. Explain how one can use zero-sum games and linear programming to construct a *lower bound* on the approximation ratio of any truthful, anonymous, and neutral randomised voting rule for a fixed number of candidates m. The approach here is similar to that of Exercise 1 in Tutorial 3, so you might want to consult that.

B. Imagine that we had a magic box, or an *oracle*, which, given a certain randomised voting rule as input, returned the worst-case approximation ratio of this voting rule over all instances of the problem, together with an instance on which this worst-case ratio is attained. Explain how we could use this oracle together with the idea in the previous bullet to create an iterative algorithm which gets closer and closer to finding the truthful, anonymous, and neutral voting rule with the best possible approximation ratio (over all instances of the problem).

References

[1] Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–81, 1977.

[2] Salvador Barbera. Majority and positional voting in a probabilistic framework. *The Review of Economic Studies*, 46(2):379–389, 1979.

[3] Aris Filos-Ratsikas and Peter Bro Miltersen. Truthful approximations to range voting. *In proceedings of the 10th International Conference on Web and Internet Economics (WINE)*, 2014.