

AGTA Tutorial 6 Solutions

Lecturer: Aris Filos-Ratsikas

Tutor: Charalampos Kokkalis

March 18, 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.

Solution 1. Consider a social choice setting with three candidates A, B , and C . Let us assume without loss of generality that in case of a tie in which A receives the largest amount of points, A is the winner. Consider the following preference profile with three voters:

Voter 1: $A \succ B \succ C$ Voter 2: $C \succ B \succ A$ Voter 3: $B \succ A \succ C$

All candidates receive 1 point, so the winner by the discussion above is A . Now consider the misreport of Voter 2 in which the voter reports:

Voter 2: $B \succ C \succ A$

Now candidate B receives 2 points, whereas candidate A receives 1 point and candidate C receives 0 points. The winner is clearly candidate B . Since Voter 2 prefers candidate B to candidate A , the voter has manipulated the plurality mechanism. This proves that the mechanism is not truthful.

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 $v_i : M \rightarrow \mathbb{R}$ 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.*)

Solution 2.

- A. Since the valuation functions v_i are injective, this means that each valuation function v_i induces a strict preference ranking \succ_i over the candidates. Furthermore, for any such ranking, there is a valuation function that induces it, and therefore the cardinal social choice setting is a generalisation of the setting that we discussed in class.

- B. We will show that any deterministic truthful mechanism is ordinal. Assume by contradiction that there exists a deterministic truthful mechanism that is not ordinal. This means that there are two different valuation function v_i and v'_i of some voter i such that $f(v_i, v_{-i}) = a \neq b = f(v'_i, v_{-i})$, while at the same time $a \succ_i b \rightarrow a \succ'_i b$, where \succ_i and \succ'_i are the ordinal preference rankings induced by v_i and v'_i respectively.

Assume first that $b \succ_i a$. In that case, voter i with true value v_i would have an incentive to misreport v'_i and force the mechanism to output b instead, a more preferable outcome, contradicting truthfulness. Similarly, assume that $a \succ_i b$ with also implies $a \succ'_i b$. Then, voter i with true value v'_i would have an incentive to misreport v_i and force the mechanism to output a instead, a more favourable outcome. In each case, truthfulness is contradicted.

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

Solution 3. Assume by contradiction that f is truthful and onto, but it is not Pareto optimal. This means that for some input preference profile \succ , every single voter prefers some candidate A to some candidate B , but f outputs B instead. Let \succ' be the preference profile obtained from \succ in which every voter is ranking candidate A first. By monotonicity, which is implied by truthfulness, since the relative order of A and B has not changed in the rankings of any voter, the winner in \succ' must still be candidate B .

Now, by the fact that f is onto, there exists some preference profile \succ^{onto} such that $f(\succ^{\text{onto}}) = A$. From this profile, obtain preference profile \succ'^{onto} by moving A on top of each voters' ranking. By monotonicity, since A was winning before, it must still be winning. Finally, from profile \succ'^{onto} , move all the candidates (step by step) such that $\succ'^{\text{onto}} = \succ'$. By monotonicity, since the winner on \succ'^{onto} was A , and we only “moved around” candidates below A for every voter, the winner on \succ' must still be A . But earlier we argued that the winner on \succ' was B ; we have obtained a contradiction.

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 location.

- Show that the expected maximum cost of the mechanism is within $3/2$ of the optimal maximum cost.

Solution 4.

- A.** For the social cost, the location that minimizes the objective is the median top choice, i.e. $\text{med}(x_1, \dots, x_n)$. To see this, observe that any location lower than this would make more than half the agents more unhappy and any location higher than this would make more than half the agents more unhappy, and the rate at which the cost decreases or increases is the same. For the maximum cost, it is not hard to see that the optimal choice is to set $y = (\max_i x_i + \min_i x_i)/2$, i.e. the midpoint between the lowest and the highest ideal location.
- B.** For the social cost, consider a profile where $x_1 = x_2 = \dots = x_{n-1} = 0$ and $x_n = 1$ and let x_n be the choice of Dictatorship. The optimal solution is to set $y = x_1$ for a social cost of 1 whereas Dictatorship achieves a social cost of $n - 1$ for a ratio of $n - 1$. The same example works for the maximum cost as well, where the optimal is the $(x_1 + x_n)/2$ for a maximum cost of $1/2$, whereas the maximum cost of Dictatorship is 1.
- C.** The mechanism that always outputs the median top choice is optimal, as discussed above. It is also easy to see that it is truthful. Indeed, the only way for an agent with ideal location x_i to affect the median is to report some location x'_i which is on the other side of the median, compared to x_i . But in that case the median is going to move even further away from x_i , thus increasing the agent's cost. An exception to this case is when x_i is the median top choice, but in that case the agent already has cost 0 and therefore has no incentive to misreport.
- D.** Consider a profile (x_1, x_2) with two agents, with $x_1 = 0$ and $x_2 = 1$ and assume by contradiction that a deterministic truthful mechanism with approximation ratio strictly smaller than 2 for the maximum cost exists. This implies that the mechanism must output $y \in (0, 1/2] \cup [1/2, 1)$ as otherwise its ratio on this profile would be at least 2, since the maximum cost of the optimal solution is $1/2$. Without loss of generality, assume that $y \in [1/2, x_2)$. Now consider the profile (x_1, x'_2) , where $x'_2 = y$. Since the approximation ratio of our mechanism is smaller than 2 by assumption, the output \hat{y} of the mechanism on input profile (x_1, x'_2) should not be x'_2 . At the same time however, on (x_1, x'_2) the agent with top choice location x'_2 could claim that its ideal location is x_2 , resulting in profile (x_1, x_2) and moving the outcome exactly on her ideal location x'_2 . This is an example of a beneficial manipulation, violating truthfulness and we obtain a contradiction.
- E.** Finally, for the randomized mechanism that selects the two endpoints and the middle point with different probabilities:
- Observe that the only deviations that can actually affect the outcome are those that alter the endpoints of the interval; let $lt(x)$ and $rt(x)$ be the left and the right endpoint of profile x respectively. Let i be a deviating agent, who misreports its most preferred location x_i as being x'_i and let $x' = (x'_i, x_{-i})$. We will assume that the agent affects $lt(x)$ by moving it left or right; the arguments for $rt(x)$ are symmetric. Let $lt(x) - lt(x') = \delta$.
If $\delta < 0$, which means that $lt(x)$ moves to the right, then the deviating agent i must be the one with most preferred location $lt(x)$, as otherwise such a move would not be possible. The agent however does not benefit from such a move as both the left endpoint and the middle would move to the right and thus away from its ideal location.
Assume now that $\delta > 0$, which means that $lt(x)$ moves to the left (and so does $(lt(x) + rt(x))/2$). Let f be the mechanism. We have:

$$\begin{aligned} \text{cost}_i(f(x')) &= \frac{1}{4}(t_i - lt(x) - \delta) + \frac{1}{4}(rt(x) - x_i) + \frac{1}{2} \left| x_i - \frac{lt(x) - \delta + rt(x)}{2} \right| \\ &\geq \frac{1}{4}(x_i - lt(x)) + \frac{1}{4}(rt(x) - x_i) + \frac{1}{2} \left| x_i - \frac{lt(x) + rt(x)}{2} \right| \\ &= \text{cost}_i(f(x)), \end{aligned}$$

which means that the agent does not benefit from the deviation.

- By scaling, we can assume without loss of generality that $\min_i x_i = 0$ and that $\max_i x_i = 1$. The optimal maximum cost is $1/2$, whereas the mechanism achieves $\frac{1}{4} \cdot 1 + \frac{1}{4} \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$ for an approximation ratio of $3/2$. The point here is that while earlier we proved that no truthful deterministic mechanism can achieve an approximation ratio better than 2, we see here that if we allow a weaker truthfulness concept that involves randomization, then better approximation ratios are achievable.

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 $f : (\succ)^n \rightarrow \Delta(A)$, 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-Satterthwaite 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 $U_{m,n}^q$: 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 $D_{m,n}^q$: 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.

- 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.
- 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).

Solution 5. This exercise is meant to showcase a cool way of using zero-sum games to obtain approximation ratios for truthful mechanisms. Fix the number of candidates m . Then, by the Theorem 1, we know exactly which mechanisms are truthful (in expectation), i.e., any convex combination of mechanisms of the form $U_{m,n}^q$ and $D_{m,n}^q$, for different values of q . For example, when $m = 3$, then the possible mechanisms are $U_{n,3}^1$, $U_{n,3}^2$, $U_{n,3}^3$, $D_{n,3}^{n/2+1}$, $D_{n,3}^{n/2+2}$, $D_{n,3}^{n/2+3}$, \dots , $D_{n,3}^n$. Crucially, there is a finite number of those, and for each one of those, given as input an instance I , we can compute the approximation ratio of the mechanism on the instance.

- A. For our zero sum game, the maximiser will have all the aforementioned mechanisms $U_{m,n}^q$ and $D_{m,n}^q$ as pure strategies. A mixed strategy of the maximiser will be a randomised truthful mechanism; by the characterisation of Theorem 1, in fact every randomised truthful mechanism will be a mixed strategy of the maximiser. The minimiser will have a set of “bad” instances $I_1, I_2, I_3, \dots, I_k$ as the pure strategies. The utility (of the maximiser) when playing a pure strategy against a pure strategy of the minimiser will be the inverse of the approximation ratio of the corresponding mechanism on the corresponding instance. If we solve the zero-sum game for a collection of instances \mathcal{I} of the minimiser, we obtain a lower bound on the approximation ratio of any truthful randomised mechanism. This is only a lower bound rather than an exact bound, because the worst-case instance for the best possible mechanism might not be in the set \mathcal{I} . But we have a guarantee that on \mathcal{I} , no mechanism can do better than the value of the game.
- B. Now assume that we have our magic oracle which inputs a mechanism (as a probability distribution over the components $U_{m,n}^q$ and $D_{m,n}^q$) and outputs (a) the approximation ratio of the mechanism and (b) the worst-case instance on which this approximation ratio is attained. We can iteratively improve out lower and upper bounds as follows.
- (a) We first run our zero-sum game (linear program) using the set of bad instances \mathcal{I} . This gives us an optimal strategy of the maximiser, i.e., a mechanism M .
 - (b) We then use our magic oracle on M to find an approximation ratio ρ . The inverse of this ratio will generally be smaller than the value of the game. The oracle will also give us an instance x on which this ratio is attained.
 - (c) We add x to the set of bad instances \mathcal{I} and re-run the zero-sum game. By definition the minimiser has now become more powerful, so the value of the game shall be smaller. The optimal mixed strategy of the maximiser will be different from before, so we will get a new mechanism M' .
 - (d) We use our magic oracle on M' to obtain a new approximation ratio ρ' and a new instance x' . We can then add x' to the set \mathcal{I} and continue with the same process.
 - (e) As we iterate through the process, the value of ρ that we compute increases and the value of the game decreases. When they are equal, we will have found the best possible approximation ratio.

Note that this magic oracle can actually be constructed using a (non-convex) quadratic program. These are like linear programs, but the objective function is quadratic, not linear (but the constraints are linear). The problem is that since the function is non-convex, we do not have algorithms to solve such a program in polynomial time, so the execution of the oracle is quite a costly operation. But still, using these, one can compute good approximation ratio bounds for up to 5 candidates. See [3] and [4].

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.
- [4] Soroush Ebadian, Aris Filos-Ratsikas, Mohamad Latifian, and Nisarg Shah, N. Computational aspects of distortion. *In proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2024.