UPenn NETS 412: Final Practice Problems

April 2018

Problem 1

For this problem, you should know the VCG mechanism as well as Myerson's Lemma characterizing truthful choice functions.

A telecom company wants to decide which neighborhood to give fiber internet access. There are m neighborhoods $1, \ldots, m$. Each user i is located in exactly one neighborhood and has value v_i if her neighborhood gets fiber and value 0 otherwise, where $v_i \in \mathbb{R}_{\geq 0}$. Let N_j be the set of users located in neighborhood j. The telecom company wants to design a truthful mechanism (X, P) for this problem. Here truthful means dominant strategy incentive compatible (DSIC).

For each of the following choice functions X, does there exist a payment rule P such that the mechanism is DSIC? Why or why not? If there does, for a bonus point each, give an example of such a payment rule.

Part a Suppose the telecom company wants to select the neighborhood with the largest sum of values. That is, $X(v) = \arg\max_j \sum_{i \in N_j} v_i$.

Solution. This is the social-welfare maximizing allocation, since social welfare is exactly the sum of values of people in the neighborhood that get fiber. So VCG is a truthful mechanism that selects it, and it can be made truthful with VCG payments. It is not required to explicitly write down the VCG payment function for full credit.

Part b Each neighborhood j costs the telecom company a different amount C_j to build infrastructure. So the company wants to select the neighborhood with the largest sum of values minus cost of investment. That is, $X(v) = \arg\max_j \left(\sum_{i \in N_j} v_i\right) - C_j$.

Solution. The key point is to recognize and show that this is a single-parameter domain.[1] (If you didn't get that, stop reading now and try to solve it with that hint!)

Each alternative can be described as one unit of service for everyone in the neighborhood that gets fiber and zero for others, and each agent has valuation v_i per unit of service. So it is a single-parameter domain. So by Myerson's Lemma, an allocation rule is truthful if and only if it is monotone, meaning that every time we fix all bids of all but i, we have that the more i agent bids, the more "units" they get.

Now, this is a monotone choice rule: if we fix the bids of every other agent but i, if i increases her bid, she never gets less "amount" and sometimes gets more. More specifically, if her neighborhood is already getting fiber, then if she bids higher, she still gets fiber; and if her

[1] You could actually use the Groves mechanism here, but later subparts will need the single-parameter approach so we use it here too.

neighborhood is not getting fiber, bidding higher can only possibly switch her neighborhood to getting it.

Since the choice rule is monotone, it can be made truthful by Myerson's Lemma.

The truthful payment turns out to be that each bidder in the winning neighborhood pays the smallest value she could have bid and still had her neighborhood selected; others pay zero.

If you did not yet solve the rest of the parts of the problem, stop reading the solutions here, go back, and try to solve them!

Part c The telecom company wants to give fiber to the neighborhood that is most heterogeneous in valuation, for some reason. So it picks the neighborhood where the difference between highest and lowest value is smallest, i.e. $X(v) = \arg\max_j \left(\max_{i \in N_j} v_i\right) - \left(\min_{i \in N_j} v_i\right)$.

Solution. This allocation rule is not monotone: we can easily see an example where one agent is the lowest bidder in her neighborhood which is not currently getting fiber, but by decreasing her bid, she causes her neighborhood to become the most heterogeneous one and get fiber. So a decrease in her bid actually causes the chosen neighborhood to switch to hers, gaining a higher allocation. So by Myerson's Lemma, it cannot be made truthful.

Part d The telecom company wants to give fiber to the neighborhood that is most homogeneous in valuation, for some reason. So it picks the neighborhood where the difference between highest and lowest value is smallest, i.e. $X(v) = \arg\min_j \left(\max_{i \in N_j} v_i\right) - \left(\min_{i \in N_j} v_i\right)$.

Solution. Again, this is not monotone: consider an agent with the highest valuation by far in her neighborhood, which is not currently being selected. By decreasing her bid, she may cause the choice to switch to her neighborhood. So by Myerson's Lemma, it cannot be made truthful.

Part e The telecom wants to give fiber to neighborhoods that have higher value for it, but it wants to diminish the impact of the high outliers. So it wants to maximize the sum of the logarithms of values, i.e. $X(v) = \arg\max_j \sum_{i \in j} \log(v_i)$.

Solution. This is actually monotone: bidding higher cannot decrease the "amount" a bidder gets, and may increase it. That is, if a bidder's neighborhood is already getting fiber, then it continues to get fiber if she bids higher; and if not, bidding higher can only keep it fiberless or possibly switch it to getting fiber.

So by Myerson's Lemma, X can be made truthful. In fact, the truthful payment again turns out to be each agent in the winning neighborhood paying the smallest value she could have bid and still had her neighborhood selected; others pay zero.

Problem 2

For this problem, you should know the proper scoring rule characterization as well as peer prediction.

Priscilla hires you to help her elicit predictions. She wants experts to make predictions $p \in [0,1]$ representing the probability of rain the next day. When she then observes $y \in \{\text{rain}, \text{no rain}\}$, she will score them with some function $S(p,y)$.

She asks you to design the scoring rule S. It must be truthful: for any belief q, the optimal report should be q. Next, she will add some additional requirements.

Part a Suppose in this part that Priscilla requires:

- The expected score for an agent who believes $q = 0$ should be 10.
- The expected score for an agent who believes $q = 1$ should be 10.
- The expected score for an agent who believes $q = 0.5$ should be 20.

Tell Priscilla how to design a scoring rule satisfying her requirements, or explain why it is impossible.

Solution. It is impossible. The expected score function $f(q)$ for any proper (truthful) scoring rule is convex, according to the characterization theorem. But a function that satisfies $f(0) = 10$, $f(0.5) = 20$, and $f(1) = 10$ cannot be convex. (If this is not clear to you, draw the three points on a plot!)

Part b Suppose in this part Priscilla requires that the expected score be larger for agents who believe higher probabilities of rain. That is, if $q > q'$, then an agent who believes q should expect to make more money than she would if she believed q'.

Tell Priscilla how to design a scoring rule satisfying her requirements, or explain why it is impossible.

Solution. It is possible, just by taking a convex, increasing function such as $f(q) = q^2$. Now we can use the scoring rule characterization to construct a proper scoring rule. (It is not necessary for this problem to give the exact definition, but it is $S(p, \text{rain}) = p^2 + 2p(1-p)$ and $S(p, \text{no rain}) = p^2 + 2p(0-p) = -p^2$.)

Part c Now Priscilla changes the rules a bit. She knows the expert has observed a piece of information about the weather, a signal which was either "high" or "low". If the expert observed high, then he believes a probability of 0.75 of rain; if low, then 0.25. Priscilla wants you to design a payment rule that just asks the expert to report either high or low, then scores this report with the observed outcome (rain or no rain). It should be a function taking two arguments, the signal and the observation of whether it actually rained or not. It should be truthful, i.e. the agent should strictly prefer reporting his true signal to lying.

Tell Priscilla how to design a payment function satisfying her requirements, or explain why it is impossible.

Solution. It is possible using the idea of peer prediction (but simpler). We take any strictly proper scoring rule S such as the log scoring rule. Let y be the observed weather outcome. If the expert reports high, we pay him $S(0.75, y)$; if he reports low, we pay him $S(0.25, y)$.

This is truthful because, if he sees the signal high, then he would rather get paid $S(0.75, y)$ than $S(0.25, y)$ because his true belief is 0.75 and S is proper. Similarly, if he sees low, he would rather report low.

Part d Suppose we know the expert's initial belief is $q = 0.5$. However, the expert has a crystal ball available that will enable him to find out for sure whether or not it will rain. The catch is that it costs him 5 units of utility to use the crystal ball. Priscilla wants you to design a scoring rule so that the expert's optimal action is use the crystal ball, then report truthfully what he finds out.

Tell Priscilla how to design a scoring rule satisfying her requirements, or explain why it is impossible.

Solution. The proper scoring rule S will have some associated convex function $f(q)$, by the characterization. Let's see how to design f, which will tell us what S is.

If the expert doesn't look into the crystal ball, his belief will be $q = 0.5$, and he will report 0.5 because S is proper, so he will get expected utility $f(0.5)$.

If he looks into the crystal ball, then with probability 0.5 he will see that it will rain, and with probability 0.5 he will see that it will not rain. Either way, he will report his new belief truthfully because S is proper. So his expected utility is $0.5f(0) + 0.5f(1) - 5$ because he lost 5 utility by looking into the ball.

So he improves expected utility by looking into the ball if $0.5f(0) + 0.5f(1) - 5 > f(0.5)$. Pictorally, this says that if f looks like a smiley face, the left and right endpoints have to be on average 5 larger than the point in the middle of the smile.

So Priscilla should use a convex function that is high enough at the two extremes and low in the middle, for example, $f(q) = 40(p - 0.5)^2$. This satisfies $f(0) = f(1) = 10$ while $f(0.5) = 0$. So when we use the proper scoring rule this gives, the expected utility for looking in the crystall ball is $10 - 5 = 5$ while the expected utility for not looking is 0.

Problem 3

Consider a game in which each of two players announces an integer between 0 and $2n$, where n is fixed ahead of time and known to both players. Let a_1 denote the announcement of Player 1 and a_2 denote the announcement of Player 2. The payoffs are determined as follows:

1. If $a_1 + a_2 \leq 2n$, Player 1 receives a_1 and Player 2 receives a_2.
2. If $a_1 + a_2 > 2n$ and $a_1 > a_2$, Player 1 receives $2n - a_2$ and Player 2 receives a_2.
3. If $a_1 + a_2 > 2n$ and $a_1 < a_2$, Player 1 receives a_1 and Player 2 receives $2n - a_1$.
4. If $a_1 + a_2 > 2n$ and $a_1 = a_2$, each player receives n.

Find an equilibrium of this game using iterated elimination of dominated strategies first when $n = 5$ and then when $n = 50$.

Solution. For $n = 5$: If Player 1 announces 6 her payoff is 5 if Player 2 announces 5 or 6, and 6 if Player 2 announces anything else. Likewise for Player 2.

Round 1: If Player 1 announces $a_1 < 6$ she'll get a_1 no matter what Player 2 announces. Therefore any strategy smaller than 6 is dominated by 6. Likewise for Player 2. We can delete all strategies between 0 and 5 for both players.

Round 2: If Player 1 announces 10 she can get at most 5. This is because Player 2 announces a number between 6 and 10. Therefore 10 is dominated by 6 in this reduced game. Likewise for Player 2. We can delete 10 for both players.

Round 3: If Player 1 announces 9 she can get at most 5. This is because Player 2 announces a number between 6-9. Therefore 9 is dominated by 6 in this further reduced game. Likewise for Player 2. We can delete 9 for both players.

Round 4: If Player 1 announces 8 she can get at most 5. This is because Player 2 announces a number between 6-8. Therefore 8 is dominated by 6 in this reduced game. Likewise for Player 2. We can delete 8 for both players.

Round 5: If Player 1 announces 7 she can get at most 5. This is because Player 2 announces a number between 6-7. Therefore 7 is dominated by 6 in this reduced game. Likewise for Player 2. We can delete 7 for both players.

Hence only 6 survives the iterated elimination of strategies for both players. As a result the payoff of each player is 5.

For $n = 50$: If Player 1 announces 51 her payoff is 50 if Player 2 announces 50 or 51, and 51 if Player 2 announces anything else. Likewise for Player 2.

Round 1: If Player 1 announces $a_1 < 51$ she'll get a_1 no matter what Player 2 announces. Therefore any strategy smaller than 51 is dominated by 51. Likewise for Player 2. We can delete all strategies between 0 and 50 for both players.

Round 2: If Player 1 announces 100 she can get at most 50. This is because Player 2 announces a number between 51 and 100. Therefore 100 is dominated by 51 in this reduced game. Likewise for Player 2. We can delete 100 for both players.

Round 3: If Player 1 announces 99 she can get at most 50. This is because Player 2 announces a number between 51-99. Therefore 99 is dominated by 51 in this further reduced game. Likewise for Player 2. We can delete 99 for both players.
...

Round 49: If Player 1 announces 53 she can get at most 50. This is because Player 2 announces a number between 51-53. Therefore 53 is dominated by 51 in this further, further, . . . , further reduced game. Likewise for Player 2. We can delete 53 for both players.

Round 50: If Player 1 announces 52 she can get at most 50. This is because Player 2 announces a number between 51-52. Therefore 52 is dominated by 51 in this further, further, . . . , further reduced game. Likewise for Player 2. We can delete 52 for both players. Hence only 51 survives the iterated elimination of strategies for both players. As a result the payoff of each player is 50.

Problem 4

Solution. See lecture notes.

Part a Define *regret* in an expert-advice setting.

Part b Define *average regret* in an expert-advice setting.

Part c Outline the halving algorithm, and report its guarantees in the case of a perfect expert and an expert making OPT mistakes.

Part d Outline the polynomial weights algorithm, and give the best $\mathcal{O}()$ bound you know for its regret and average regret.

Problem 5

Consider the polynomial weights algorithm. In which of the following situations would you most want a high value for ϵ. justify your response:

1. The performance of the actions is relatively uniform
2. The best action is perfect (loss zero) throughout all rounds
3. The best action starts start off poorly, then finishes well
4. The worst action starts off well, then finishes poorly

Solution. (2), the best action is perfect throughout all rounds. In this case, any time any action makes a mistake, we should downweight it very heavily, because we know that is not the best action. So we want a very high ϵ.

Problem 6

Consider the Top Trading Cycles (TTC) algorithm. Suppose that an individual's favorite item is the one she starts with. Explain why she is indifferent between participating in the game and not, and explain why her participation does not affect the allocation to all of the other players, regardless of their preferences.

Solution. If her favorite item is the one she starts with, she points to herself on the first day, which forms a cycle. That cycle is cleared and she is properly assigned her starting item under TTC. Since she gets her own item, the allocation to the rest of the agents must be a permutation of their own items, so they are not affected by her choice to participate or not.

Problem 7

Give an example with three people on each side of the Stable Marriage Problem in which each person has a total, strict preference order over the people on the opposite side and there exists a *unique* stable matching. Show that your matching is indeed unique and demonstrate how the Deferred Acceptance (Gale-Shapley) algorithm finds it by showing the steps the algorithm takes.

Solution. Let a, b, c be the left people and A, B, C the right. Suppose a and A prefer each other over everyone else, and the same for b and B, and for c and C. Then the only stable matching is for $a - -A$, $b - -B$, and $c - -C$. Anything else, and these pairs would be blocking pairs. In this case, the Deferred Acceptance algorithm has a, b, c propose to their favorites in the first round, then immediately halt.

Problem 8

Consider the following set of voter preferences over three candidates, A, B, and C, where we write $X \succ Y$ to indicate that X is preferred to Y. There are 13 voters.

1. Four people rank $B \succ A \succ C$
2. Four people rank $C \succ A \succ B$
3. Three people rank $A \succ B \succ C$
4. Two people rank $C \succ B \succ A$

Part a If we use the *plurality rule*, which candidate wins and what is the count of the votes?

Solution. A has 3 first-place votes, B has 4, and C has 6. So C wins.

Part b Suppose we use the following rule: each person votes for their favorite candidate. The candidate with the fewest votes is eliminated. Anyone who voted for that eliminated candidate changes their vote to their favorite candidate who has not yet been eliminated. The winner is the candidate who is not eliminated at the end of the process.

What is the count of the votes in the first round, and which candidate is eliminated? What is the count of the votes in the second round, and which candidate is eliminated?

Solution. The first round counts are given above, so A is eliminated with the fewest first-place votes.

If we remove A, the only change in first-place votes is that people in the third category now rank B first. So now B has seven first-place votes while C only has 6, and C is eliminated. (So B wins.)

Part c Is there a Condorcet winner? If so, who?

Solution. A is preferred to B by 7 of the 13 people; A is preferred to C by 7 of the 13 people. So A is a Condorcet winner.