UPenn NETS 412: Algorithmic Game Theory
Midterm Practice Problems 2

Instructor: Bo Waggoner

Problem 1

Consider a game with two players with action profiles $A_1 = A_2 = \{1, ..., 10\}$. The players get the same payoff with the function $u(s_1, s_2) = (s_1 + s_2)/2$ if $s_1 \neq s_2$ and $u(s_1, s_2) = s_1 + 1$ if $s_1 = s_2$. Basically, the function will average the two values given, but give a bonus for giving the same numbers.

Part a List all the pure strategy nash equilibria. Justify your answer.

Solution. It is an equilibrium for both players to play 8; for both to play 9; or for both to play 10. Then they each get a payoff of 9, 10, 11 respectively and either player deviating cannot improve, e.g. if both play 8 and one deviates, he can at best play 10 and still only gets a payoff of 9, which is not better.

Nothing else is an equilibrium: If the players play different numbers, then the player with the lower number can improve by switching to copy the other (this increases the average *and* gets the bonus); if both are playing 7 or less, then the other can improve by switching to 10, which raises the average by more than 1 even though it loses the bonus.

Part b Find a mixed strategy Nash equilibrium that is not a pure strategy equilibrium.

Solution. For example, there is one where both players mix between playing 9 and 10. (If you haven't yet found the solution, stop reading and solve it using this hint!)
Suppose the first player plays 9 with probability p and 10 with probability $1 - p$. Then the second player's utility for playing 9 is $p(10) + (1 - p)(9.5)$, and for playing 10 is $p(9.5) + (1 - p)(11)$. Setting these equal, so the second player is indifferent to 9 and 10, we get $p = 0.75$. Symmetrically, if the second player mixes with probability $p = 0.75$, then the first player is indifferent between playing 9 and 10. So both players playing 9 with probability 0.75 and 10 with probability 0.25 is an equilibrium.

Part c Give a correlated equilibrium that is not a pure or mixed strategy equilibrium.

Solution. For example, with probability 0.5 both play 9, with probability 0.5 both play 10.

Part d Consider the distribution $\mathcal{D}$ where, for each $j = 7, 8, 9, 10$, both players play j with probability $\frac{1}{4}$. Is this a correlated equilibrium? Why or why not?

Solution. No. To show it isn't, we just have to find one recommendation where one player would prefer to deviate. In particular, when both players are recommended action 7, player 1 would rather deviate to playing 10, as in this case she gets 8.5 payoff instead of 8. (More formally, this shows she would rather deviate to the swap function that follows the recommendation except when she is recommended to play 7, where she deviates to 10 instead.)

Part e Consider the above distribution again. Is it a coarse correlated equilibrium? Why or why not?

Solution. Yes. We have to show that each player would rather follow than deviate to any fixed action $1, \ldots, 10$. If they deviate, then deviating to 10 would be the best deviation since it's the largest number and has an equal probability of getting the bonus as $7, 8, 9$. But if a player deviates to 10, she gets

$$\begin{aligned}\frac{1}{4}\left(\frac{7+10}{2}\right) + \frac{1}{4}\left(\frac{8+10}{2}\right) + \frac{1}{4}\left(\frac{9+10}{2}\right) + \frac{1}{4}(11) &= \frac{17+18+19}{8} + \frac{11}{4} \\ &= \frac{27+11}{4} \\ &= \frac{38}{4}.\end{aligned}$$

Whereas if she follows, she gets

$$\begin{aligned}\frac{1}{4}(8) + \frac{1}{4}(9) + \frac{1}{4}(10) + \frac{1}{4}(11) &= \frac{8+9+10+11}{4} \\ &= \frac{38}{4}\end{aligned}$$

which is equal. So she cannot gain by deviating, so it is a coarse correlated equilibrium.

Part f Now consider the same game except the payoff if $s_1 = s_2$ is $s_1 - .5$. List all the pure strategy nash equilibria. Justify your answer briefly.

Solution. Both players playing 10 is an equilibrium, as is one player playing 9 and the other playing 10. In all these cases, both players get payoff 9.5, which is the highest possible (so any switch only hurts them). No other profile is an equilibrium: If both players play the same number which is not 10, either can gain by switching to 10 (as the average is higher *and* they don't lose the penalty 0.5). If both players play different numbers and one is playing 10, the other can gain by switching from their low number to 9. If both players play different numbers and neither is playing 10, then either can gain by switching to 10.

Problem 2

You may consult the definition of the polynomial weights algorithm while solving this practice problem.

Part a Suppose all losses are either 0 or 1, and furthermore, suppose you know there will be some perfect action that always gets loss 0 (however, you don't know which action.) What value would you choose for the parameter ϵ in the algorithm? Justify your answer. What algorithm for following expert advice does this remind you of?

Solution. I would set $\epsilon = 1.0$, so that if an action has a loss of 1, its "weight" gets set equal to zero and we never pick that action again. All actions with no mistakes so far would still have weight 1, including the perfect action. This is similar to the Halving algorithm.

Part b Suppose we set ϵ very small, for example, $\frac{1}{T}$. In a sentence or two: What would the practical difference be in how the PW algorithm behaves and why might it perform worse in some settings? (Your answer can be high-level, not involving calculations.)

Solution. It would make very small changes to the actions' weights, so it would continue to have a good chance of picking a bad action even after it has suffered a lot of loss.

Part c Suppose we set ϵ large, for example, 0.1 regardless of T. In a sentence or two: What would the practical difference be in how the PW algorithm behaves and why might it perform worse in some settings? (Your answer can be high-level, not involving calculations.)

Solution. It would downweight actions' weights too quickly; a few bad days would be enough to totally change the weight of the action relative to the other actions' weights. Similarly, an action that initially suffered a few losses would get a very small weight, and then wouldn't be picked for a while even if it started performing very well after that.

Problem 3

In this game, each player wants to send flow along a shared channel of maximum capacity 1. On the one hand, each player wants to send as much flow as possible along the channel. On the other hand, the channel becomes less useful the closer it gets to its maximum capacity. Each player can choose to send an amount of flow $x_i \in [0, 1]$ along the channel. (That is, the action set for each player i is $A_i = [0, 1]$, and hence is not finite).

Let $x_i \in [0, 1]$ denote the action of player i. For a profile of actions $x \in A$, let $S = \sum_{j=1}^{n} x_i$ be the total flow. Then player i has utility $u_i(x_i, x_{-i}) = x_i(1 - S)$.

Note that if the total flow is larger than one, each player has negative utility!

Part a Show that this game is a potential game with potential function

$$\phi(x) = -S + \frac{1}{2}S^2 + \frac{1}{2}\sum_{j=1}^{n} x_j^2.$$

Hint: One approach is to consider $\frac{du_i}{dx_i}$ *and* $\frac{d\phi}{dx_i}$.

Solution. We have to show that the change in ϕ is the opposite sign as the change in u_i when i changes his action x_i. That is, an increase in u_i always corresponds to a decrease in potential and vice versa.

Note that $\frac{dS}{dx_i} = 1$. We have

$$\begin{aligned}\frac{du_i}{dx_i} &= (1)(1-S) + x_i(-1) \\ &= 1 - S - x_i.\end{aligned}$$

Meanwhile, differentiating term-by-term (remember the chain rule: $\frac{d}{dx_i}S^2 = 2S\frac{dS}{dx_i}$),

$$\frac{d\phi}{dx_i} = -1 + S + x_i.$$

So $\frac{d\phi}{dx_i} = -\frac{du_i}{dx_i}$, so when player 1 changes x_i, u_i and ϕ change by opposite signs. (In fact the changes are equal and opposite, so ϕ is an *exact* potential function.)

An alternative approach without calculus is to consider two actions x_i, x_i' and show that the change is equal for both functions.

Part b Find a Nash equilibrium of this game. What is the social welfare at this equilibrium? (*i.e.* the sum of utilities of all the players.)

Solution. This is a symmetric game, so let's try and find a symmetric Nash equilibrium (in which all players are playing the same action). Consider player i's action, and let $t = \sum_{j \neq i} x_i$. Player i wants to maximize his utility $u_i(x_i, x_{-i}) = x_i(1 - t - x_i) = -x_i^2 + (1-t)x_i$. By calculus, we know that u_i is maximized by setting $x_i = (1-t)/2$.

This would have to hold for all players, so we get each player is playing some $y = (1-t)/2$ where $t = (n-1)y$. So $y = (1-(n-1)y)/2$, so $2y = 1-(n-1)y$, so $(n+1)y = 1$, so every player is playing $x_i = y = \frac{1}{n+1}$. Each player gets utility $u_i(x) = \frac{1}{n+1}\left(1 - \frac{n}{n+1}\right) = \frac{1}{(n+1)^2}$.

So the social welfare is $\Sigma_{i=1}^{n} u_i(x) = n/(n+1)^2$ (Sidenote: this is $\Theta\left(\frac{1}{n}\right)$.)

Part c What is the optimal social welfare? (*i.e.* what is the social welfare at the profile of actions that maximizes it, regardless of whether or not this profile is an equilibrium.) Intuitively, why isn't this optimum achieved in equilibrium?

Solution. Each player's utility is $x_i(1-S)$, so the total utility is $\sum_{i=1}^{n} x_i(1-S) = S(1-S)$. We can maximize this if $S = \frac{1}{2}$ and the optimal value is $\frac{1}{4}$. (For example, every player plays $x_i = \frac{1}{2n}$.) But it's not achieved in equilibrium because any one player would want to deviate to increasing her flow.

Problem 4

Consider a game in which there are two firms producing identical goods. Let q_1 and q_2 represent the quantity of the good produced by each firm. The price at which the goods can be sold depends on the quantity produced. The price that the firms can charge for their goods is given as $p(q_1, q+2) = a - (q_1 + q_2)$, where a is some constant. Finally, each firm incurs a cost of c for each unit of the good produced. The firms experience no fixed costs, so the cost of producing a quantity of zero is zero. We'll assume $c < a$ and that the good is infinitely divisible, so it makes sense to talk about fractional quantities.

This is called a Cournot game (also Cournot duopoly, or Cournot oligopoloy for more than two players). The model was developed by Cournot, a mathematician who lived in the 19th century, and Cournot published his model over a century before the emergence of the field of game theory. This model and its variants are heavily studied in economics.

Part a Write down the profit function of Firm 1 π_1 as a function of c, q_1, q_2, a, where *profit* is the revenue of Firm 1 minus the cost of producing the quantity q_1 of the good.

Solution.

$$\pi_1 = (a - (q_1 + q_2))q_1 - cq_1 = aq_1 - q_1^2 - a_1 q_2 - cq_1$$

Part b Suppose Firm 2 has fixed its quantity at $\overline{q}_2$. What quantity q_1 should Firm 1 pick to maximize profit? This is Firm 1's best-response to $\overline{q}_2$.

Hint: One way to figure this out is to take a derivative of the profit function.

Solution. By taking a derivative and setting it equal to zero, we have that Firm 1 maximizes its profit by solving

$$0 = \frac{d}{dq_1}\pi_1 = a - 2q_1 - \overline{q}_2 - c$$

which happens when

$$q_1 = \frac{1}{2}(a - \overline{q}_2 - c)$$

Part c By symmetry, Firm 2's best-response function to a choice $\overline{q}_1$ by Firm 1 is identical to that of Firm 1, just with the indices swapped. Find the values of q_1 and q_2 at the Nash equilibrium (in this game, there is a unique, pure strategy Nash equilibrium).

Solution. A choice of q_1^* and q_2^* is a Nash equilibrium if and only if the firms are best-responding to each other. When the firms are best-responding to each other, we should have that

$$q_1* = \frac{1}{2}(a - q_2^* - c)$$

and

$$q_2^* = \frac{1}{2}(a - q_1^* - c)$$

Let' plug the expression for q_2^* into Firm 1's best-response function.

$$q_1^* = \frac{1}{2}(a - \left(\frac{1}{2}(a - q_1^* - c)\right) - c)$$

. This is an equation in one unknown, so we can solve it for q_1^*. Doing this, we get

$$q_1^* = \frac{a-c}{3}$$

Plugging this in for q_1^* in Firm 2's best-response function (or just arguing by symmetry), we get

$$q_2^* = \frac{a-c}{3}$$

Part d What is the market price at the Nash equilibrium?

Solution. The firms together are producing $\frac{2a-2c}{3}$ units of the good. Since the price is a minus the total quantity produced, we get

$$p = a - \frac{2a-2c}{3} = \frac{a+2c}{3}$$

Part e How much profit does each firm make?

Solution. At this price and quantity, we get

$$\pi_1 = \left(\frac{a+2c}{3}\right)\left(\frac{a-c}{3}\right) - c\frac{a-c}{3} = \left(\frac{a-c}{3}\right)^2$$

And symmetrically for Firm 2,

$$\pi_2 = \left(\frac{a-c}{3}\right)^2$$

Part f Suppose the firms instead colluded and each produced $q_i = \frac{a-c}{4}$. What would the market price be? What is the profit of each firm?

Solution. The price is

$$p = a - (\frac{a-c}{2}) = \frac{a+c}{2}$$

. Each firm sees profit

$$\pi_i = \left(\frac{a+c}{2}\right)\left(\frac{a-c}{4}\right) - c\left(\frac{a-c}{4}\right) = \frac{1}{2}\left(\frac{a-c}{2}\right)^2$$

Part g In this collusive case, we see that the profit for each firm is higher than in the Nash equilibrium. Explain why this collusive agreement is unsustainable. That is, why is this collusive agreement not a Nash equilibrium?

Solution. Suppose Firm 1 followed through with the collusive agreement and produced $q_1 = \frac{a-c}{4}$. Firm 2 can maximize its profit by best-responding to this choice of q_1, so Firm 2 picks

$$q_2 = \frac{1}{2}\left(a - \frac{a-c}{4} - c\right) = \frac{3(a-c)}{8}$$

When Firm 2 deviates like this, the price of the good is $p = a - \frac{5(a-c)}{8}$. Firm 1, which followed the agreement, gets profit

$$\pi_1 = \left(a - \frac{5(a-c)}{8}\right)\left(\frac{a-c}{4}\right) - c\frac{a-c}{4} = \frac{3}{8}\left(\frac{a-c}{2}\right)^2$$

Observe that this is *less* than what Firm 1 would have gotten had Firm 2 stuck by the agreement.

Firm 2, which deviated, sees profit

$$\pi_2 = \left(a - \frac{5(a-c)}{8}\right)\left(\frac{3(a-c)}{8}\right) - c\frac{3(a-c)}{8} = \frac{9}{16}\left(\frac{a-c}{2}\right)^2$$

Observe that this is *more* than what Firm 2 would have gotten by sticking with the agreement, so it is better off deviating from the agreement. Since Firm 2 is better off by deviating, the collusive agreement can't be a Nash equilibrium.