499 Final Solutions[1]

1. Question 1

TRUE/FALSE

(a) Every two-player zero-sum game has an optimal strategy.

True; this follows from von Neumann's Minimiax Theorem

(b) Suppose I have a polynomial time algorithm that, when given any two-player general sum game defined by two $n \times n$ integer-valued payoff matrices, could output an ESS, or determine that no ESS exists. (This algorithm has a run time that is polynomial in n and in the log of the largest integer in the payoff matrices.) Then P=NP. (Equivalently, it is NP-hard to: find an ESS or decide no ESS exists.)

True. We mentioned in class that deciding whether or not an ESS exists is NP-hard (in fact it is strictly harder than any NP-complete problem).

(c) In the game of chess, it is known that the first player has a winning strategy. That is, the first player can guarantee a win, regardless of what the second player does.

False. Zermelo's Theorem implies that one of three alternatives holds, but it is not known which holds. (The first player has a winning strategy, or the second player does, or both players have a strategy guaranteeing at least a draw.)

(d) Consider the problem of deciding if a general sum game has at least two Nash equilibria. This problem is NP-complete.

True. We mentioned this in class.

(e) Suppose we have a bandit problem with rewards satisfying $0 \leq R_t \leq 1$ for all times $t \geq 0$. Assume that $k \leq H$. Then there is an algorithm for the bandit problem with expected regret at time H bounded by $100\sqrt{kH \log H}$.

True. We proved this in class/ in the notes.

(f) The Condorcet paradox no longer occurs if we consider an election between four candidates. That is, the Condorcet paradox only occurs in Condorcet elections between three candidates.

False. For any number of candidates $\{a_1, \ldots, a_k\}$ with $k > 3$, we can have three voters who rank the candidates $\{a_{k-3}, \ldots, a_k\}$ below $\{a_1, a_2, a_3\}$, then just reproduce the Condorcet paradox by ranking a_1, a_2, a_3 as in the case $k = 3$.

2. Question 2

.

Prove the following. On a standard Hex game board, the first player has a winning strategy. That is, the first player has a strategy that guarantees a win, regardless of what the second player does.

Solution. The game cannot end in a tie. So by Zermelo's Theorem, one player has a (memoryless) winning strategy. We argue by contradiction. Suppose the second player has a winning strategy. Note that the game positions are symmetric with respect to swapping the colors of the game board. So, the first player can use the second player's winning strategy as follows. On the first move, the first player just colors in any hexagon H. On the first player's next move, they just use the winning strategy of the second player (pretending that hexagon H is white). If the first player is ever required to color the hexagon H, they just color any other hexagon instead, if an uncolored hexagon exists. Having this extra hexagon colored can only benefit the first player. So, the first player is guaranteed to win, and both players will win, a contradiction. We conclude the first player has a winning strategy.

3. Question 3

Let $m \geq 1$ be a positive integer.

Show that Δ_m is convex and bounded.

(You need to justify your answer.)

Solution. Let $x, y \in \Delta_n$ and let $t \in [0,1]$. By definition of Δ_n, $\sum_{i=1}^n x_i = 1$ and $x_i \geq 0$ for all $1 \leq i \leq n$, and $\sum_{i=1}^n y_i = 1$ and $y_i \geq 0$ for all $1 \leq i \leq n$. Since $t + (1-t) = 1$, we then have

$$\sum_{i=1}^n tx_i + \sum_{i=1}^n (1-t)y_i = t(1) + (1-t)1 = t + (1-t) = 1.$$

That is, $tx + (1-t)y$ satisfies $\sum_{i=1}^n (tx + (1-t)y)_i = 1$. Since $t, 1-t \geq 0$, and $x_i, y_i \geq 0$ for all $1 \leq i \leq n$, we have $tx_i + (1-t)y_i \geq 0$ for all $1 \leq i \leq n$. Then is $(tx + (1-t)y)_i \geq 0$ for all $1 \leq i \leq n$. In summary, we have shown that $tx + (1-t)y \in \Delta_n$. That is, Δ_n is convex. For boundedness, we use $0 \leq x_i \leq 1$, so that $x_i^2 \leq 1$ for all $1 \leq i \leq n$ to get

$$\|x\| = \sqrt{\sum_{i=1}^n x_i^2} \leq \sqrt{\sum_{i=1}^n 1} = \sqrt{n}.$$

So, we may choose $r := \sqrt{n}$ to see that Δ_n is bounded. We could also use $r = 1$, since $0 \leq x_i \leq 1$ implies $0 \leq x_i^2 \leq x_i$ for all $1 \leq i \leq n$, so

$$\|x\| = \sqrt{\sum_{i=1}^n x_i^2} \leq \sqrt{\sum_{i=1}^n x_i} = \sqrt{1} = 1,$$

using $x \in \Delta_n$ in the penultimate equality.

4. Question 4

Suppose we have a Vickrey auction, i.e. a sealed-bid second-price auction, with $n \geq 2$ bidders. Each bidder submits a bid to the seller (in a sealed envelope), and the winner of the auction is the highest bidder, and they pay the second-highest bid.

- Explain why it is an equilibrium when each bidder bids their private value.
- Explain why the total expected revenue of the seller is

$$n \int_{-\infty}^{\infty} F_Y(t) \mathbf{E}(Y \mid Y \leq t) f(t) dt,$$

  where $Y = \max(V_2, \ldots, V_n)$ is the maximum private value of $n-1$ bidders. (Hint: the expected payment of buyer 1 with private value t is $F_Y(t)\mathbf{E}(Y \mid Y \leq t)$ where $F_Y(t) := \mathbf{P}(Y \leq t)$ for all $t \in \mathbf{R}$.)

Solution. A buyer with private value v can make a profit at most $\max(v - m, 0)$ where m is the maximum of all other bids, and this profit is achieved when a buyer bids their private value. Note that this conclusion does not require any assumptions on the buyers, such as independence of private values.

In a sealed-bid second-price auction, at the symmetric equilibrium, the expected payment of buyer 1 with private value t is $F_Y(t)\mathbf{E}(Y|Y \le t)$. So, the expected payment of buyer 1 is $\int_{-\infty}^{\infty} F_Y(t)\mathbf{E}(Y|Y \le t)f(t)dt$. The seller's expected revenue is the sum of payments of all n buyers, hence we multiply this quantity by n.

5. Question 5

Consider a network with four vertices v_1, v_2, v_3, v_4 (cities) and four edges (roads): $(v_1, v_3), (v_3, v_2), (v_1, v_4), (v_4, v_2)$. Each edge has a cost which describes the time it takes for a driver to traverse that road. Suppose the edges have costs $t, 1, 1, t$, respectively.

Suppose there is one unit of traffic, representing a large number of players. Each player wants to go from v_1 to v_2. Each player acts independently of each other player. And each player wants to minimize their travel time. Assume that every player is using the same strategy at equilibrium.

- Under the above assumptions, describe the unique Nash equilibrium for the players and the mean travel time of one player. Justify your answer.

Suppose now we add a short and fast (one way) highway from v_3 to v_4 with zero cost.

- Under the above assumptions, for the new highway system, describe the unique Nash equilibrium for the players and the mean travel time of one player. Justify your answer.
- What is the ratio between your current answer and your previous answer?

Solution. The unique Nash equilibrium occurs when each player chooses the top or bottom path with probability $1/2$ each, for a mean travel time of $1/2 + 1 = 3/2$.

Any driver taking the v_1 to v_3 route will use the new highway to then proceed from v_4 to v_2, since this is strictly better for such a driver, giving them a travel time of 2. In fact, the unique Nash equilibrium occurs when each player chooses this path, since any player starting on the v_1 to v_4 path would rather take the previously mentioned path, as its travel time is shorter. Since we found only one Nash equilibrium, the price of anarchy and price of stability are both equal to $2/(3/2) = 4/3$.

6. Question 6

Let $f\colon \{-1,1\}^n \to \{-1,1\}$.

- Show that the noise stability of f with parameter $0 < \rho < 1$ is at most 1.
- If $\sum_{x\in\{-1,1\}^n} f(x) = 0$, show that the noise stability of f with parameter $0 < \rho < 1$ is at most ρ.

Solution. From Plancherel's Theorem, $\sum_{S\subseteq\{1,\ldots,n\}} |\widehat{f}(S)|^2 = 1$. Since $0 < \rho < 1$, we have $\sum_{S\subseteq\{1,\ldots,n\}} |\rho^{|S|}\widehat{f}(S)|^2 \le \sum_{S\subseteq\{1,\ldots,n\}} |\widehat{f}(S)|^2 = 1$.

The second inequality follows similarly. If $\sum_{x\in\{-1,1\}^n} f(x) = 0$, then $\widehat{f}(\emptyset) = 0$. Since this coefficient is zero, we can use the inequality $\rho^{|S|} \le \rho$ for all S with $|S| \ge 1$ to get $\sum_{S\subseteq\{1,\ldots,n\}} |\rho^{|S|}\widehat{f}(S)|^2 \le \rho \sum_{S\subseteq\{1,\ldots,n\}} |\widehat{f}(S)|^2 = \rho$.

7. Question 7

Suppose we have a two-player symmetric game such with payoffs described by a matrix such as the following.

$$A = \begin{pmatrix} 4 & 3 & 2 & 5 & 6 \\ 3 & 1 & 8 & 9 & 1 \\ 7 & 0 & 7 & 0 & 7 \\ 1 & 3 & 3 & 2 & 1 \\ 0 & 1 & 2 & 9 & 1 \end{pmatrix}$$

Describe in detail an algorithm that outputs a Nash equilibrium of this game.

Make sure to justify why your algorithm outputs a Nash equilibrium.

Also give a bound on the run time of the algorithm.

Solution. We are given the claim: if $(\widetilde{x}, \widetilde{y}) \in \Delta_m \times \Delta_n$ is a Nash equilibrium and if

$$I := \{1 \leq i \leq m : \widetilde{x}_i > 0\}, \qquad J := \{1 \leq j \leq n : \widetilde{y}_j > 0\}. \qquad (*)$$

then with payoff matrices A, B respectively, we have

$$\max_{i=1,\ldots,m} (A\widetilde{y})_i = (A\widetilde{y})_i, \quad \forall i \in I. \qquad \text{and} \qquad \max_{j=1,\ldots,n} (\widetilde{x}^T B)_j = (\widetilde{x}^T B)_j, \forall j \in J.$$

It then follows from the claim that

$$\widetilde{x}^T A \widetilde{y} = (A\widetilde{y})_i, \quad \forall\, i \in I. \qquad \widetilde{x}^T B \widetilde{y} = (\widetilde{x}^T B)_j, \quad \forall\, j \in J.$$

Algorithm 1 Input: $m \times n$ matrices A, B, $I \subseteq \{1, \ldots, m\}$, $J \subseteq \{1, \ldots, n\}$ Assumption: there exists a Nash equilibrium $\widetilde{x} \in \Delta_m, \widetilde{y} \in \Delta_n$ with I, J satisfying $(*)$. Output: $\widetilde{x}, \widetilde{y}$.

- Let $\widetilde{x}_i, i \in I$ and let $\widetilde{y}_j, j \in J$ be the $|I| + |J|$ unknown variables.
- Consider $|I| - 1$ equations $(A\widetilde{y})_i = (A\widetilde{y})_{i'} \quad \forall i, i' \in I$.
- Consider $|J| - 1$ equations $(\widetilde{x}^T B)_j = (\widetilde{x}^T B)_{j'} \quad \forall j, j' \in J$.
- Together with the equations $\sum_{i \in I} \widetilde{x}_i = 1$, $\sum_{j \in J} \widetilde{y}_j = 1$, we have $|I| + |J|$ linear equations in $|I| + |J|$ equations. By assumption, there exists at least one positive solution to these equations. Output such a solution, if one can be found where $\widetilde{x}, \widetilde{y}$ have nonnegative values and they also maximize the expected payoffs of the players (i.e. if $\widetilde{x}^T A \widetilde{y} \geq (A\widetilde{y})_i$ for all $1 \leq i \leq m$ and $\widetilde{x}^T A \widetilde{y} \geq (\widetilde{x}^T B)_j$ for all $1 \leq j \leq n$.)

Algorithm 2 (Finding a Nash Equilibrium). Input: $m \times n$ matrix A, B. Output: $\widetilde{x}, \widetilde{y}$ a Nash equilibrium

- Run Algorithm 1 for every possible pair of subsets $I \subseteq \{1, \ldots, m\}, J \subseteq \{1, \ldots, n\}$.
- From Nash's Theorem, there exists at least one pair I, J where Algorithm 1 succeeds.

From the above discussion, Algorithm 1 finds a Nash equilibrium, if it exists with a given support as in $(*)$. So, Algorithm 2 finds at least one Nash equilibrium, since it searches over all possible supports I, J defined in $(*)$. Moreover, Algorithm 2 has run time at least 2^{m+n}, since we might need to check this many pairs of I, J. If each solution of the linear equations takes, say $10(m+n)^3$ time, (say using Gaussian elimination on at most $m+n$ equations.) then the total run time is at most $10(m+n)^3 2^{m+n}$.

8. Question 8

Let $X_1, \ldots, X_n$ be real-valued i.i.d. (independent identically distributed) random variables. Assume that

$$\mathbf{E}e^{\alpha X_1} \leq e^{\alpha^2/2}, \qquad \forall\, \alpha \in \mathbf{R}.$$

Assume also that $\mathbf{E}X_1 = 0$. Prove that

$$\mathbf{P}\Big(\frac{1}{n}\sum_{i=1}^{n} X_i > t\Big) \leq e^{-nt^2/2}, \qquad \forall\, t > 0.$$

Solution. Let $\alpha > 0$. Using Markov's inequality, and $\alpha t \geq 0$,

$$\mathbf{P}(\frac{1}{n}\sum_{i=1}^{n} X_i \geq t) = \mathbf{P}(e^{\alpha \frac{1}{n}\sum_{i=1}^{n} X_i} \geq e^{\alpha t})$$

$$\leq e^{-\alpha t}\mathbf{E}e^{\alpha \frac{1}{n}\sum_{i=1}^{n} X_i} = e^{-\alpha t}\prod_{i=1}^{n}\mathbf{E}e^{\alpha X_i/n}.$$

By assumption, $\mathbf{E}e^{\alpha X_i/n} \leq e^{\alpha^2/(2n^2)}$, so for any $t \geq 0$

$$\mathbf{P}(\frac{1}{n}\sum_{i=1}^{n} X_i \geq t) \leq e^{-\alpha t}e^{\alpha^2/(2n)}.$$

Since $\alpha > 0$ is arbitrary, we choose α to minimize the right side. This minimum occurs when $\alpha = tn$, so that $-\alpha t + \alpha^2/(2n) = -t^2 n + t^2 n/2 = -t^2 n/2$, giving the first desired bound.