

CSC304 Fall'26

Assignment 1 Solutions

Q1 [20 Points] Nash Equilibria

Compute all (mixed) Nash equilibria in each of the following games. At the end of your solution, explicitly write the set of Nash equilibria using the following notation. A Nash equilibrium can be denoted as $((p_1, p_2, p_3), (q_1, q_2, q_3))$, where p_1 , p_2 , and p_3 are the probabilities of the row player playing T , M , and B respectively, and q_1 , q_2 , and q_3 are the probabilities of the column player playing L , C , and R , respectively.

(a) [5 Points] Game 1:

[Hint: Instead of brute-forcing, apply a trick we learned in class to reduce the size of the game without changing the set of Nash equilibria.]

	L	C	R
T	(5, 5)	(1, 9)	(4, 2)
M	(3, 5)	(9, 6)	(6, 9)
B	(6, 5)	(1, 9)	(5, 4)

(b) [15 Points] Game 2: Note that the only differences compared to Game 1 are in the rewards to the row player in (T,C) and (M,R).

[Hint: It is still possible to reduce the size of Game 2 from 3×3 to 3×2 . Applying the indifference principle to the latter will still result in $(2^3 - 1) \cdot (2^2 - 1) = 21$ cases. Instead of writing each one separately, you can club them into four categories: neither player randomizes, only the row player randomizes, only the column player randomizes, and both players randomize. Provide a brief argument for any categories that do not lead to any Nash equilibria, and show detailed calculations only for the remaining categories.]

	L	C	R
T	(5, 5)	(3, 9)	(4, 2)
M	(3, 5)	(9, 6)	(1, 9)
B	(6, 5)	(1, 9)	(5, 4)

Solution to Q1

(a) We can iteratively eliminate strictly dominated actions without changing the set of equilibria. In this game, we can eliminate L (strictly dominated by C), then T (strictly dominated by M), then B (strictly dominated by M), and then C (strictly dominated by R).

This yields a unique Nash equilibrium (M, R) , or in our representation, $((0, 1, 0), (0, 0, 1))$.

(b) In this game, you can still eliminate L (strictly dominated by C). However, in the resulting game, no actions are strictly dominated. This yields the following game.

	C	R
T	(3, 9)	(4, 2)
M	(9, 6)	(1, 9)
B	(1, 9)	(5, 4)

Now, we apply the indifference principle to solve for the set of Nash equilibria by taking the following cases.

- Both players don't randomize ($|X_1| = |X_2| = 1$). It is easy to check by drawing the best response diagram that there are no pure Nash equilibria in the game. *Note to students: You should draw the diagram here.*
- Only the row player randomizes ($|X_1| > 1, |X_2| = 1$). If $X_2 = \{C\}$, then the rewards for the row player for three actions are different. Hence, by the indifference principle, the row player cannot randomize. The same holds for $X_2 = \{R\}$.
- Only the column player randomizes ($|X_1| = 1, |X_2| > 1$). Again, for each action of the row player, note that the rewards of the column player for her two actions are different. Hence, by the same argument as above, there are no such equilibria.
- Both players randomize ($|X_1|, |X_2| > 1$). This means $X_2 = \{C, R\}$. There are four subcases for X_1 : $\{T, M\}, \{M, B\}, \{T, B\}, \{T, M, B\}$.

Let us consider these remaining four cases. In all cases, we denote column player's strategy by $(0, q, 1 - q)$, where $q \in (0, 1)$. We denote indifference principle by IP in short.

- $X_1 = \{T, M\}, X_2 = \{C, R\}$. Suppose row player uses $(p, 1 - p, 0)$ where $p \in (0, 1)$.
 - IP for row player: $E[T] = E[M] \geq E[B]$. So $3q + 4(1 - q) = 9q + (1 - q) \geq q + 5(1 - q)$. From first two, we get $q = 1/3$, and verify that it satisfies all constraints.
 - IP for column player: $E[C] = E[R]$. So $9p + 6(1 - p) = 2p + 9(1 - p)$, so $p = 3/10$. This gives one NE: $((3/10, 7/10, 0), (0, 1/3, 2/3))$.
- $X_1 = \{T, B\}, X_2 = \{C, R\}$. Suppose row player uses $(p, 0, 1 - p)$ where $p \in (0, 1)$.
 - IP for row player: $E[T] = E[B] \geq E[M]$. So $3q + 4(1 - q) = q + 5(1 - q) \geq 9q + (1 - q)$. From first two, we get $q = 1/3$, and verify that it satisfies all constraints.
 - IP for column player: $E[C] = E[R]$. So $9p + 9(1 - p) = 2p + 4(1 - p)$. This doesn't have a solution in $p \in (0, 1)$. So there are no NE in this case.
- $X_1 = \{M, B\}, X_2 = \{C, R\}$. Suppose row player uses $(0, p, 1 - p)$ where $p \in (0, 1)$.
 - IP for row player: $E[M] = E[B] \geq E[T]$. So $9q + (1 - q) = q + 5(1 - q) \geq 3q + 4(1 - q)$. From first two, we get $q = 1/3$, and verify that it satisfies all constraints.
 - IP for column player: $E[C] = E[R]$. So $6p + 9(1 - p) = 9p + 4(1 - p)$, so $p = 5/8$. This gives one NE: $((0, 5/8, 3/8), (0, 1/3, 2/3))$.
- $X_1 = \{T, M, B\}, X_2 = \{C, R\}$. Suppose row player uses $(p, r, 1 - p - r)$ where $p, r, 1 - p - r \in (0, 1)$. (Fun fact: This is equivalent to $p > 0 \wedge r > 0 \wedge p + r < 1$.)

- IP for row player: $E[T] = E[M] = E[B]$. So $3q + 4(1 - q) = 9q + (1 - q) = q + 5(1 - q)$. Check that $q = 1/3$ is the unique value to satisfy this.
- IP for column player: $E[C] = E[R]$. So

$$9p + 6r + 9(1 - p - r) = 2p + 9r + 4(1 - p - r) \Rightarrow 2p + 5 = 8r \Rightarrow r = (5 + 2p)/8.$$

At this point, you can simply say the set of NE in this case are all $((p, r, 1 - p - r), (0, 1/3, 2/3))$, where $p, r, 1 - p - r \in (0, 1)$ and $r = (5 + 2p)/8$, and that will earn full marks. But we can simplify this further, if we want. We want $1 - p - r > 0 \Rightarrow 3 - 10p > 0 \Rightarrow p < 3/10$ and $r < 1 \Rightarrow 5 + 2p < 8$, which is always true (and $r > 0$ is trivially true). Hence, the NE in this case are all $((p, \frac{5+2p}{8}, \frac{3-10p}{8}), (0, 1/3, 2/3))$ where $p \in (0, 3/10)$.

Hence, the set of all NE are:

$$\begin{aligned} & \{((3/10, 7/10, 0), (0, 1/3, 2/3)), ((0, 5/8, 3/8), (0, 1/3, 2/3))\} \\ & \bigcup \{((p, \frac{5+2p}{8}, \frac{3-10p}{8}), (0, 1/3, 2/3)) : p \in (0, 3/10)\}. \end{aligned}$$

(Fun fact: This set is exactly the line segment connecting the two equilibria in the first set.)

Q2 [20 Points] Games Without Pure Nash Equilibria

Consider a two-player game in which each player has two actions: the row player chooses between T and B , and the column player chooses between L and R . Let u_1 and u_2 denote the reward functions of the row player and the column player, respectively. Our goal is to prove the following claim: *If the game has no pure Nash equilibrium, then it has exactly one Nash equilibrium, and in it, both players randomize, i.e., each action of each player is played with a probability strictly between 0 and 1.*

(a) [5 Points] Suppose the row player's action T gives her at least as much reward as her other action B against every action of the column player, i.e., $u_1(T, L) \geq u_1(B, L)$ and $u_1(T, R) \geq u_1(B, R)$. Show that the game has a pure Nash equilibrium. (By symmetry, the same holds if B is at least as good as T against every action of the column player, or if one of the two actions of the column player is at least as good as the other against every action of the row player; you do not need to prove these.)

(b) [15 Points] Suppose the game has no pure Nash equilibrium. Using part (a), we already know that $u_1(T, L) - u_1(B, L)$ and $u_1(T, R) - u_1(B, R)$ are both nonzero and have opposite signs, and by symmetry, $u_2(T, L) - u_2(T, R)$ and $u_2(B, L) - u_2(B, R)$ are also both nonzero and have opposite signs; you do not need to prove this. Use this information to prove the claim above.

[Hint: Let x and y be the probabilities with which the row player plays T and the column player plays L , respectively. First, show that in any Nash equilibrium, if one player plays a pure strategy, then so does the other player, which contradicts the assumption that there are no pure Nash equilibria. Next, apply the indifference principle to find the only possible equilibrium, and verify that it is indeed a Nash equilibrium.]

Solution to Q2

(a) Let $t \in \{L, R\}$ be a best response of the column player to T . We claim that (T, t) is a pure Nash equilibrium. The column player has no incentive to deviate because t is her best response to T . The row player has no incentive to deviate because $u_1(T, t) \geq u_1(B, t)$ by assumption.

(b) Let $a = u_1(T, L) - u_1(B, L)$, $b = u_1(T, R) - u_1(B, R)$, $c = u_2(T, L) - u_2(T, R)$, and $d = u_2(B, L) - u_2(B, R)$. We are given that a and b are nonzero with opposite signs, and so are c and d .

Note to students: For completeness, here is why this follows from part (a). If $a \geq 0$ and $b \geq 0$, then T is at least as good as B against every action of the column player, and hence, by part (a), the game has a pure Nash equilibrium, which is a contradiction. Similarly, if $a \leq 0$ and $b \leq 0$, then B is at least as good as T against every action of the column player, and hence, by the symmetric version of part (a), the game has a pure Nash equilibrium, which is again a contradiction. Thus, neither $(a \geq 0 \text{ and } b \geq 0)$ nor $(a \leq 0 \text{ and } b \leq 0)$ holds. The former failing means that $a < 0$ or $b < 0$, and the latter failing means that $a > 0$ or $b > 0$. Since $a < 0$ and $a > 0$ cannot both hold (and the same for b), it must be that $(a < 0 \text{ and } b > 0)$ or $(a > 0 \text{ and } b < 0)$. In both cases, a and b are nonzero and have opposite signs. The proof for c and d is identical.

Let x be the probability with which the row player plays T , and y be the probability with which the column player plays L . If the column player plays L with probability y , then the expected reward of the row player from T minus her expected reward from B is

$$f(y) = y \cdot a + (1 - y) \cdot b.$$

Similarly, if the row player plays T with probability x , then the expected reward of the column player from L minus her expected reward from R is

$$g(x) = x \cdot c + (1 - x) \cdot d.$$

First, we show that in every Nash equilibrium (x, y) , we must have $x, y \in (0, 1)$. Suppose $x \in \{0, 1\}$. Then $g(x) \in \{c, d\}$ is nonzero, which means one of the two actions of the column player is strictly better than the other, and she must play only that action in a Nash equilibrium. That is, $y \in \{0, 1\}$. But then (x, y) is a pure Nash equilibrium, which is a contradiction. The case of $y \in \{0, 1\}$ is identical, using the fact that $f(y) \in \{a, b\}$ is nonzero.

Now, take any Nash equilibrium (x, y) . Because $x, y \in (0, 1)$, both players play both of their actions with a positive probability. By the indifference principle, each player must be indifferent between her two actions. That is, $f(y) = 0$ and $g(x) = 0$, which gives

$$y = \frac{b}{b - a} \quad \text{and} \quad x = \frac{d}{d - c}.$$

(Here, $b \neq a$ because a and b have opposite signs, and similarly $d \neq c$.) Hence, the game has at most one Nash equilibrium.

Finally, we show that the game does have a Nash equilibrium, namely (x, y) given by the above expressions. (Alternatively, existence follows from Nash's theorem.) Because a and b have opposite signs, $|b - a| = |a| + |b| > |b|$, and $b - a$ has the same sign as b . Hence, $y = b/(b - a) \in (0, 1)$.

Similarly, $x = d/(d - c) \in (0, 1)$. Under this strategy profile, $f(y) = 0$ and $g(x) = 0$, i.e., each player is indifferent between her two actions, and hence, cannot gain by deviating. This shows that (x, y) is a Nash equilibrium in which both players randomize.

Q3 [20 Points] Strategic Computing

In a strategic computing game, there is a set of agents $N = \{1, \dots, n\}$, and a set of machines $M = \{1, \dots, m\}$. Each agent i has a job (referred to as job i) that requires p_i time to be processed on any single machine. A pure strategy s_i of agent i is a machine on which job i will be processed (i.e., $s_i \in M$). Given a pure strategy profile $\vec{s} = (s_1, \dots, s_n)$, the load on machine j is the sum of processing times of the jobs assigned to machine j :

$$L_j(\vec{s}) = \sum_{i \in N: s_i = j} p_i.$$

The cost to each agent i is the load on the machine chosen by the agent, i.e., $c_i(\vec{s}) = L_{s_i}(\vec{s})$. The social cost is defined to be the *maximum load on any machine*: $C(\vec{s}) = \max_{j \in M} L_j(\vec{s})$. (Observe that the social cost is not the sum of agent costs, but rather the maximum of agent costs.)

(a) [4 Points] Consider a strategic computing game with $m = 2$ machines and $n = 5$ jobs with processing times $p_1 = 1$, $p_2 = 2$, $p_3 = 3$, $p_4 = 6$, and $p_5 = 6$. Compute the price of stability and the price of anarchy of this game with respect to pure Nash equilibria. Show your calculation.

(b) [6 Points] Prove that the price of stability of any strategic computing game is 1. In other words, you need to show that in every strategic computing game, there is at least one socially optimal pure strategy profile that is also a pure Nash equilibrium.

[Hint: Not every socially optimal pure strategy profile will be a Nash equilibrium. First, observe that in a social optimum, an agent assigned to the machine with the highest load (assuming it is unique) must have no incentive to switch. Then, think about when this would be true for the agents assigned to the other machines.]

(c) [10 Points] Prove that the price of anarchy of any strategic computing game with respect to pure Nash equilibria is at most $2 - 2/(m + 1)$.

[Hint: Let OPT be the optimum social cost. You want to prove that the social cost of any pure Nash equilibrium $\vec{s}$ is at most $(2 - 2/(m + 1)) \cdot OPT$. Let j^* be a machine with the highest load under $\vec{s}$. First, prove the desired result when j^* has a single job under $\vec{s}$. Next, assume j^* has at least two jobs and i^* is a job on j^* with the minimum processing time among all jobs on j^* . To prove the desired result, derive two inequalities: the first relates OPT to the average (not maximum) load across all machines under $\vec{s}$, and the second relates the load on each machine under $\vec{s}$ to the maximum load under $\vec{s}$.]

Solution to Q3

(a) The total processing time is $1 + 2 + 3 + 6 + 6 = 18$. Hence, in every pure strategy profile, one of the two machines has a load of at least $18/2 = 9$, so the optimum social cost is at least 9. Assigning jobs 1, 2, 4 (with processing times 1, 2, 6) to one machine and jobs 3, 5 (with processing

times 3, 6) to the other results in loads (9, 9). Hence, the optimum social cost is 9.

Now, let us find the pure Nash equilibria. In a pure strategy profile with machine loads H and $18 - H$, where $H \geq 9$, a job on the machine with load H can move to the other machine and have a cost of $18 - H + p_i$. A job on the machine with the smaller load can never benefit from moving, since it would face a cost of at least H on the other machine. Hence, the profile is a pure Nash equilibrium if and only if every job i on the machine with load H satisfies $18 - H + p_i \geq H$, i.e., if and only if the smallest job on that machine has processing time $p_{\min} \geq 2H - 18$. Since $p_{\min} \leq 6$, we need $H \leq 12$. Let us consider the possible values of $H \in \{9, 10, 11, 12\}$.

- $H = 9$. The condition $p_{\min} \geq 0$ always holds. So any profile with loads (9, 9) is a pure Nash equilibrium. Such profiles put $\{1, 2, 6\}$ (one of the two jobs of size 6 together with jobs 1 and 2) on one machine and $\{3, 6\}$ on the other. These are socially optimal.
- $H = 10$. We need $p_{\min} \geq 2$. The only subset of jobs with total processing time 10 is $\{1, 3, 6\}$ (up to the choice of the job of size 6), which contains a job of size 1. So there are no Nash equilibria in this case.
- $H = 11$. We need $p_{\min} \geq 4$, so the machine can only have jobs of size 6, whose total processing time is 6 or 12, never 11. So there are no Nash equilibria in this case.
- $H = 12$. We need $p_{\min} \geq 6$, so the machine must have exactly the two jobs of size 6. The other machine has jobs 1, 2, 3 with load 6. This is indeed a Nash equilibrium: a job of size 6 that moves to the other machine faces a cost of $6 + 6 = 12$, which is not smaller than its current cost of 12.

Hence, there are two types of pure Nash equilibria: one with social cost 9 and one with social cost 12. The price of stability is $9/9 = 1$, and the price of anarchy is $12/9 = 4/3$.

Note to students: An alternative to the above case analysis for showing that no equilibrium is worse than 12 is to use part (c), which shows that the price of anarchy is at most $2 - 2/(2 + 1) = 4/3$ for $m = 2$ machines. Observe that the bound in part (c) is tight for this game.

(b) Let $\vec{s}$ be a pure strategy profile that is a social optimum (minimizes the highest load), among all such profiles minimizes the second-highest load, among all such profiles minimizes the third-highest load, and so forth.

If you want to see why this is mathematically well-defined, let P_1 be the set of socially optimal pure strategy profiles. Iteratively refine this set as follows: for $k \geq 2$, P_k is the subset of profiles in P_{k-1} that minimize the k -th highest load among all profiles in P_{k-1} . Note that P_1 is trivially non-empty, and since we are always optimizing an objective, each refinement P_k is non-empty as well. We are interested in a profile $\vec{s} \in P_m$.

Note that because $\vec{s} \in P_m \subseteq P_1$, it is a social optimum. We argue that it is also a Nash equilibrium. Suppose for contradiction that agent i has an incentive to deviate from machine a to machine b . Then, $L_b(\vec{s}) + p_i < L_a(\vec{s})$. Consider the profile $\vec{s}'$ after agent i deviates to machine b . Then:

$$\begin{aligned} L_a(\vec{s}') &= L_a(\vec{s}) - p_i < L_a(\vec{s}) \\ L_b(\vec{s}') &= L_b(\vec{s}) + p_i < L_a(\vec{s}). \end{aligned}$$

Note that the loads on machines other than a and b aren't affected. Hence, when comparing $\vec{s}'$ and $\vec{s}$, $\vec{s}'$ has exactly the same set of loads that are strictly greater than $L_a(\vec{s})$, but at least one fewer load equal to $L_a(\vec{s})$. In other words, if k is the largest rank at which the load $L_a(\vec{s})$ appears in the sorted order of loads under $\vec{s}$, then $\vec{s}'$ has the same highest, second-highest, $\dots$, $(k-1)$ -st highest load as $\vec{s}$, but a strictly smaller k -th highest load than $\vec{s}$, which contradicts $\vec{s} \in P_m$.

(c) Consider any pure Nash equilibrium NE and any social optimum OPT . We want to prove $C(NE) \leq (2 - 2/(m+1)) \cdot C(OPT)$. Let j^* be its highest load machine under NE .

First, consider the case where j^* has a single job i , so $C(NE) = L_{j^*}(NE) = p_i$. It is easy to see that $C(OPT) \geq p_i$, since i must be assigned to some machine under OPT , and the social cost will be at least the load on that machine. Hence, we get $C(NE) \leq C(OPT)$ (which means it must be that $C(NE) = C(OPT)$).

Next, consider the case where j^* has at least two jobs. Let i^* be the smallest job on j^* . Hence, there is at least one other job on j^* that has processing time at least that of i^* . So:

$$p_{i^*} \leq (1/2) \cdot L_{j^*}(NE).$$

Further, in NE , i^* must have no incentive to deviate from j^* to any machine $j \neq j^*$. Thus:

$$\forall j \in M \setminus \{j^*\}, L_j(NE) \geq L_{j^*}(NE) - p_{i^*} \geq \frac{1}{2} L_{j^*}(NE),$$

where the final transition uses the above equation.

Adding these for all $j \neq j^*$, and then adding $L_{j^*}(NE)$ once on both sides, we get

$$\sum_j L_j(NE) = \sum_i p_i \geq \left(\frac{m-1}{2} + 1 \right) \cdot L_{j^*}(NE) = \frac{m+1}{2} \cdot C(NE).$$

On the other hand, under OPT , the maximum load across all machines must be at least as much as the average load across all machines, so:

$$C(OPT) \geq \frac{1}{m} \cdot \sum_j L_j(OPT) = \frac{1}{m} \cdot \sum_i p_i.$$

Combining these two equations, we get

$$C(OPT) \geq \frac{m+1}{2m} \cdot C(NE) \Leftrightarrow C(NE) \leq \left(2 - \frac{2}{m+1} \right) C(OPT).$$

Q4 [20 Points] Potential Functions and Cost-Sharing Games

Recall that in the cost-sharing game, we defined the potential function $\Phi(\vec{P}) = \sum_{e \in E(\vec{P})} \sum_{k=1}^{n_e(\vec{P})} \frac{c_e}{k}$, where $E(\vec{P})$ is the set of edges taken by at least one player in $\vec{P}$, and $n_e(\vec{P})$ is the number of players taking edge e in $\vec{P}$.

Consider a cost-sharing game with n players on a network with m edges, where the cost of each edge is 1 (i.e., $c_e = 1$ for each e). Suppose we start from an arbitrary pure strategy profile $\vec{P}^1$ and

follow an iterative process. The pure strategy profile at the beginning of iteration t is denoted $\vec{P}^t$. In iteration t , we find a player i who can change her strategy from P_i^t to P'_i to strictly decrease her cost (i.e., $c_i(P'_i, \vec{P}_{-i}^t) < c_i(\vec{P}^t)$) and set $\vec{P}^{t+1} = (P'_i, \vec{P}_{-i}^t)$. If multiple such players exist, we choose an arbitrary one, and if no such players exist, the process ends.

Our goal is to prove that this process will necessarily converge to a pure Nash equilibrium in finitely many steps, and prove an upper bound on the number of steps required.

In parts (b), (c), and (d) below, all bounds must be in terms of only n and m . Better bounds can fetch higher marks. The best possible bound will receive full marks, but, as in real life, you will not know what this bound is. So, spend a reasonable amount of time trying to improve your bound, and when you cannot find a better one, submit the one you have.

(a) [2 Points] Show that if this process stops, the final pure strategy profile must be a pure Nash equilibrium.

(b) [5 Points] Find an upper bound on the initial potential value, i.e., $\Phi(\vec{P}^1)$.

(c) [10 Points] Find a strictly positive lower bound on the amount by which the potential function must decrease in every iteration, i.e., on $\Phi(\vec{P}^t) - \Phi(\vec{P}^{t+1})$.

(d) [3 Points] Using parts (b) and (c), prove an upper bound on the number of steps required for the process to converge to a pure Nash equilibrium. Your bound must hold regardless of the choices made during the process (i.e., wherever we said “arbitrary” in the process description).

Solution to Q4

(a) If the process stops, then by the way the process is defined, no player i can reduce their cost by unilaterally deviating. This is the definition of a pure Nash equilibrium.

(b)

$$\begin{aligned}\Phi(\vec{P}^1) &= \sum_{e \in E(\vec{P}^1)} \sum_{k=1}^{n_e(\vec{P}^1)} \frac{1}{k} \\ &\leq \sum_e \sum_{k=1}^n \frac{1}{k} = mH_n \leq m(\ln n + 1).\end{aligned}$$

(c) Consider one step of the process that changes the strategy profile from $\vec{P}^t$ to $\vec{P}^{t+1}$ due to a deviation by player i , who changes the path from P_i to P'_i . Then, by the definition of the potential function, we have

$$\Phi(\vec{P}^t) - \Phi(\vec{P}^{t+1}) = c_i(\vec{P}^t) - c_i(\vec{P}^{t+1}) = \sum_{e \in P_i \setminus P'_i} \frac{1}{n_e(\vec{P}^t)} - \sum_{e \in P'_i \setminus P_i} \frac{1}{n_e(\vec{P}^{t+1})}.$$

We want a lower bound on it. Let us represent the expression as a fraction a/b in its reduced form. We know that this difference is positive, so we know that $a > 0$. In fact, this implies $a \geq 1$. Thus, our goal is to now provide an upper bound on b . We provide two incomparable upper bounds.

1. We note that there are at most m terms in the expression above because each edge can appear in at most one of the two sets $P_i \setminus P'_i$ and $P'_i \setminus P_i$, and even in that set, at most once. Hence, at most, b is the product of all denominators. Since each denominator is at most n , we get $b \leq n^m$.
2. Alternatively, since each term is a unit fraction with denominator in $\{1, \dots, n\}$, the expression can be written with the common denominator $\text{lcm}(1, \dots, n)$, and hence $b \leq \text{lcm}(1, \dots, n) = 2^{\Theta(n)}$ (the latter equality follows from the prime number theorem).

Given this, the desired lower bound is $\max(1/n^m, 1/2^{\Theta(n)})$. We can take $\max$ because both lower bounds are valid, and thus so is the best of the two.

(d) The potential function is always non-negative. Since we start from a potential value of at most $m(\ln n + 1)$ and chip away at least $\max(1/n^m, 1/2^{\Theta(n)})$ in each step, we get that the number of steps is at most the ratio between the two, which is $m \cdot \min(n^m, 2^{\Theta(n)}) \cdot (\ln n + 1)$.

Note to students: This is not necessarily the best possible bound since I am not deriving a lower bound and proving tightness. This or any better bounds will receive full marks. Any worse bound will receive appropriate partial marks.

Q5 [20 Points] Simultaneous Zero-Sum Games

There are three players $P1$, $P2$, and $P3$. $P1$ is simultaneously playing two zero-sum games: one with $P2$ in which $P1$'s reward matrix is A_{12} , and one with $P3$ in which $P1$'s reward matrix is A_{13} . The total reward for $P1$ is the sum of his rewards from the two games. Suppose all three players have the same set of actions S . The twist is that $P1$ is required to play the same (possibly mixed) strategy in both games.

- (a) [5 Points] Argue that this can be modeled as a single three-player zero-sum game.
- (b) [15 Points] Using Nash's theorem for the game derived in part (a), argue that an extension of the minimax theorem holds in this situation. Specifically, let x_1 , x_2 , and x_3 denote mixed strategies of $P1$, $P2$, and $P3$, respectively. Then show that

$$\max_{x_1} \min_{x_2, x_3} (x_1)^T (A_{12}x_2 + A_{13}x_3) = \min_{x_2, x_3} \max_{x_1} (x_1)^T (A_{12}x_2 + A_{13}x_3).$$

[Hint: Proceed along the line of the proof of minimax theorem we did in the class. To prove the two quantities are equal, we are going to show an inequality in each direction. First, find the "obvious inequality". Then, take a Nash equilibrium $(\tilde{x}_1, \tilde{x}_2, \tilde{x}_3)$. Write the three "best response" equations that describe each $\tilde{x}_i$ as the best response of player i given the strategies of the other players. Combine the three equations to get the inequality in the reverse direction.]

Solution to Q5

- (a) We construct a three-player game in which each of $P1$, $P2$, and $P3$ chooses an action from S , and the rewards are given by

$$\forall i, j, k \in S, \quad A_{123}(i, j, k) = (A_{12}(i, j) + A_{13}(i, k), -A_{12}(i, j), -A_{13}(i, k)),$$

where the three entries are the rewards to $P1$, $P2$, and $P3$, respectively, when they play i , j , and k , respectively. Here, when $P1$ plays i (in both of his games), $P2$ plays j , and $P3$ plays k , $P1$ is rewarded the sum of his two individual games, while $P2$ and $P3$ are rewarded the reward from their respective individual games (each of which is the negative of the corresponding reward of $P1$ since each individual game is zero-sum). The three rewards sum to zero for every triple (i, j, k) , so this is a three-player zero-sum game. Because $P1$ chooses a single (possibly mixed) strategy in this game, which affects his rewards against both opponents, this captures the twist that he plays the same strategy in both games. Under (independent) mixed strategies x_1, x_2, x_3 , the expected rewards to $P1$, $P2$, and $P3$ are $(x_1)^T(A_{12}x_2 + A_{13}x_3)$, $-(x_1)^T A_{12}x_2$, and $-(x_1)^T A_{13}x_3$, respectively.

(b) For brevity, let $f(x_1, x_2, x_3) = (x_1)^T(A_{12}x_2 + A_{13}x_3)$. All maximums and minimums below exist because f is continuous and the sets of mixed strategies are compact. We will prove

$$\max_{x_1} \min_{x_2, x_3} f(x_1, x_2, x_3) \leq \min_{x_2, x_3} \max_{x_1} f(x_1, x_2, x_3)$$

and the reverse inequality.

First, we prove the “obvious inequality” above. This is known as the max-min inequality and holds in general for real-valued functions. For completeness, here is a proof. For any x'_1, x'_2, x'_3 , we have

$$\min_{x_2, x_3} f(x'_1, x_2, x_3) \leq f(x'_1, x'_2, x'_3) \leq \max_{x_1} f(x_1, x'_2, x'_3).$$

Taking the maximum over x'_1 on the left-hand side and the minimum over (x'_2, x'_3) on the right-hand side gives the desired inequality.

Now we prove the reverse inequality. Using Nash’s theorem, we get that the game from part (a) has a (mixed) Nash equilibrium $(\tilde{x}_1, \tilde{x}_2, \tilde{x}_3)$ with associated rewards $(\tilde{v}_2 + \tilde{v}_3, -\tilde{v}_2, -\tilde{v}_3)$, where $\tilde{v}_2 = (\tilde{x}_1)^T A_{12} \tilde{x}_2$ and $\tilde{v}_3 = (\tilde{x}_1)^T A_{13} \tilde{x}_3$. Using the fact that in an equilibrium each player is playing a best response, we get

$$\begin{aligned} \tilde{v}_2 + \tilde{v}_3 &= \max_{x_1} f(x_1, \tilde{x}_2, \tilde{x}_3) \geq \min_{x_2, x_3} \max_{x_1} f(x_1, x_2, x_3), \\ -\tilde{v}_2 &= \max_{x_2} (-(\tilde{x}_1)^T A_{12} x_2) \Rightarrow \tilde{v}_2 = \min_{x_2} (\tilde{x}_1)^T A_{12} x_2, \\ -\tilde{v}_3 &= \max_{x_3} (-(\tilde{x}_1)^T A_{13} x_3) \Rightarrow \tilde{v}_3 = \min_{x_3} (\tilde{x}_1)^T A_{13} x_3. \end{aligned}$$

Combining these three formulas, we get

$$\begin{aligned} \min_{x_2, x_3} \max_{x_1} f(x_1, x_2, x_3) &\leq \tilde{v}_2 + \tilde{v}_3 \\ &= \min_{x_2} (\tilde{x}_1)^T A_{12} x_2 + \min_{x_3} (\tilde{x}_1)^T A_{13} x_3 \\ &= \min_{x_2, x_3} f(\tilde{x}_1, x_2, x_3) \\ &\leq \max_{x_1} \min_{x_2, x_3} f(x_1, x_2, x_3), \end{aligned}$$

where the third line follows because x_2 and x_3 can be chosen independently, so the minimum of the sum is the sum of the minimums. This completes the proof.