

Algorithmic Game Theory: Problem Set 2

Due online via GradeScope before the start of class on Tuesday, February 18

Aaron Roth

Remember you can work together on problem sets, but list everyone you worked with, and everyone turn in their own assignment. Ask questions on Slack.

Problem 1) Strategic Monty Hall (15 points)

Alice is a contestant on a game show hosted by Monty, who gives Alice a choice between three doors. One door has a car behind it, and the other two doors have nothing behind them. Alice does not know which door the car is behind, but Monty knows. Consider the following sequence:

1. Alice chooses one of the three doors, uniformly at random.
2. Out of the two doors that Alice didn't choose, Monty opens one of the doors that Alice didn't choose, and opens it, always revealing it to not have the car behind it. There are two remaining closed doors; Monty asks Alice if she'd like to switch doors.
3. Alice chooses either to remain at her initial door or switch to the other closed door.
4. After Alice makes her decision, her selected door is opened.

Alice receives a payoff of 1 if her final door has a car behind it, and 0 otherwise.

Part 1 (5 pts) What is Alice's expected payoff for remaining at her initial door? What is her expected payoff for switching? Which one maximizes her expected payoff?

We now modify the setting above so that in step 2, Monty can *choose whether* to open a door and reveal that there is nothing behind it; before, Monty was required to reveal an empty door. If Monty chooses to reveal an empty door, Alice has the opportunity to switch just as before; if he does not, then Alice just gets whatever is behind her initial choice. Alice wants to win the car — and Monty doesn't want her to. This new interaction can be modeled as a zero sum game, where Monty knows the location of the car, and Alice doesn't:

- Monty's payoff is -1 if Alice's door has a car behind it, and 0 otherwise.
- Monty can either Reveal (R) or Don't Reveal (D) an empty door, and can take different actions depending on whether Alice's initial door has the car (C) or has nothing (N). As such, Monty's set of actions is $A_M = \{R^C R^N, D^C R^N, R^C D^N, D^C D^N\}$. For example, $D^C R^N$ means that Monty doesn't reveal if Alice's initial door has a car, but reveals if Alice's initial door has nothing.

- Alice’s set of actions is $A_A = \{\text{Remain}, \text{Switch}\}$. We assume that when Alice plays ”Switch”, she only actually switches if Monty reveals a door. I.e., if Monty doesn’t reveal, ”Switch” behaves the same way that ”Remain” does.

Note that the game above takes place after Alice selects one of the three doors randomly.

Part 2 (5 pts) Write the expected payoffs of this zero sum game.

	R ^C R ^N	D ^C R ^N	R ^C D ^N	D ^C D ^N
Remain				
Switch				

Table 1: Game matrix

Part 3 (5 pts) Using the game matrix from part 2, find all the pure strategy Nash equilibria of the game and provide a brief explanation. What is Alice’s expected payoff in equilibrium?

Problem 2) Maximum-Weight Best Response Dynamics (20 points)

Consider the load balancing game on identical machines that we studied in Lecture 4. We proved that best response dynamics converges for this game. In this problem, we prove that a modification of best response dynamics converges quickly.

Maximum-Weight Best Response Dynamics

Let $a = (a_1, \dots, a_n)$ be an arbitrary action profile.

while a is not a pure strategy Nash equilibrium **do**

Among all players i who are not playing best responses in a , let i be the index of the player with largest weight w_i , and let a'_i be a best response to a_{-i} .

Update the action profile to (a'_i, a_{-i}) .

end while

Output a .

1. Prove that $\min_j \ell_j(a)$, the minimum load among all the machines, is non-decreasing as Maximum-Weight Best Response Dynamics is run (5 pts)
2. Call a player i “active” if she is currently not playing a best response, and inactive otherwise. Show that a player never goes from being “inactive” to being “active” unless another player moves onto her machine. (5 pts)
3. Let i denote the active player of maximum weight at some intermediate step of Maximum-Weight Best Response Dynamics. Show that in the round after i makes a best response move, any other players i' who become active must have $w'_i < w_i$. Conclude that no player ever makes a best response move more than once, and hence that Maximum-Weight Best Response Dynamics converges after at most n steps, where n is the number of players. (10 pts)

Problem 3) Bandwidth Sharing Game (20 points)

As you saw on the last homework, even games with infinite action sets can have pure strategy Nash equilibria; here you will show another example of such a game using the potential function method.

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). For a profile of actions $x \in A$, player i has utility $u_i(x_i, x_{-i}) = x_i(1 - \sum_{j=1}^n x_j)$.

1. (10 points) Show that this game is an exact potential game, and conclude that it has a pure strategy Nash equilibrium. *Hint: Check out this function — $\phi(x_1, \dots, x_n) = \sum_{k=1}^n (x_k - x_k^2) - \sum_{k \neq j} x_k x_j$. Note also that our argument from class that a potential function implies convergence of best response dynamics relies on the game being finite, which this game is not, so the last part of this question requires some thought.*
2. (5 points) 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.)
3. (5 points) 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.)

Problem 4) Characterizing Exact Potential Functions (10 points)

Given that best-response dynamics converge if and only if used in a potential game, we would like to characterize in which games do exact potential functions, in particular, exist. Consider an n -player game G . We say that a game is *utility separable* if there exist functions $U : A \rightarrow \mathbb{R}_{\geq 0}$ and $V_i : A_{-i} \rightarrow \mathbb{R}_{\geq 0}$ for all $i \in [n]$ such that the following holds for all players $i \in [n]$:

$$u_i(a_i, a_{-i}) = U(a) + V_i(a_{-i})$$

In other words, each player's utility function can be decomposed into the sum of two functions; a common value function U that all players' actions affect, and a personal value function V_i that only depends on the actions of the player's opponents.

1. (5 pts) Prove that if G is utility separable, then U is an exact potential function for G .
2. (5 pts) Prove that if G has an exact potential function U , then G is utility separable.