CS 1360 Spring 2025
Midterm Exam (Practice)

Problem 1: Game Theory

1. **[10 pts]** Provide a definition (using mathematical notation) of *best response* in a normal-form game.

   Solution: x_i is a best response to $\mathbf{x}_{-i}$ if for all $x_i' \in \Delta(S)$, $u_i(x_i, \mathbf{x}_{-i}) \geq u_i(x_i', \mathbf{x}_{-i})$.

2. **[15 pts]** Consider a 2-player game in normal form, and denote the strategy set of each player by S. Let $B_1(x_2)$ denote the set of (possibly mixed) best response strategies of player 1 to the (possibly mixed) strategy x_2 of player 2. For convenience, let us fix some mixed strategy $x_2^\star$ for player 2, and denote $\alpha = u_1(x_1, x_2^\star)$ for all $x_1 \in B_1(x_2^\star)$, that is, α is the maximum utility player 1 can achieve against $x_2^\star$.

   Show that $x_1 \in B_1(x_2^\star)$ *if and only if* every pure strategy $s \in S$ in the support of x_1 (i.e., every pure strategy s such that $x_1(s) > 0$) is itself in $B_1(x_2^\star)$ (i.e., $u_1(s, x_2^\star) = \alpha$).

   Note: Do not forget to show both directions.

   Solution: ($\implies$) Assume that $x_1 \in B_1(x_2^\star)$. Then, let $s \in S$ be any strategy such that $s \notin B_1(x_2^\star)$. Assume for the sake of contradiction that $x_1(s) > 0$. Then, since α is the maximum utility player 1 can achieve against $x_2^\star$, it must be true that $u_1(s, x_2^\star) < \alpha$ and also that

   $$u_1(x_1, x_2^\star) = \sum_{s \in S \,:\, x_1(s) > 0} x_1(s) u_1(s, x_2^\star) < \sum_{s \in S \,:\, x_1(s) > 0} x_1(s) \alpha = \alpha$$

   and so $x_1 \notin B_1(x_2^\star)$. However, this is a contradiction, and so it must be true that $x_1(s) = 0$. Therefore, every pure strategy that is in the support of x_1 must itself be in $B_1(x_2^\star)$.

   ($\impliedby$) Assume that x_1 is a possible mixed strategy such that every pure strategy that is in the support of x_1 is in $B_1(x_2^\star)$. Then, for every $s \in S$ such that $x_1(s) > 0$, we know that $u_1(s, x_2^\star) = \alpha$ and so

   $$u_1(x_1, x_2^\star) = \sum_{s \in S \,:\, x_1(s) > 0} x_1(s) u_1(s, x_2^\star) = \sum_{s \in S \,:\, x_1(s) > 0} x_1(s) \alpha = \alpha$$

and so $x_1 \in B_1(x_2^\star)$ by definition.

Problem 2: Equilibrium Computation

1. **[10 pts]** Define (using mathematical notation) the notion of *strong Stackelberg equilibrium (SSE)*.

   Solution: In a strong Stackelberg equilibrium (SSE) the leader plays a mixed strategy in
   $$\text{argmax}_{x_1 \in \Delta(S)} \max_{s_2 \in B_2(x_1)} u_1(x_1, s_2)$$
   where $\Delta(S)$ is the set of all mixed strategies. The follower plays the corresponding strategy s_2 in the best response set (breaking ties in favor of the leader).

2. **[15 pts]** Describe a polynomial-time algorithm that computes an SSE in a given Stackelberg game.

   Note: This was done in class.

   Solution: The leader's mixed strategy is defined by variables $x(s_1)$, which give the probability of playing each strategy $s_1 \in S$. For each follower strategy s_2^*, we compute a strategy x for the leader such that playing s_2^* is a best response for the follower, and under this constraint, x is optimal. This computation is done via the following LP:
   $$\begin{array}{ll} \max & \sum_{s_1 \in S} x(s_1) u_1(s_1, s_2^*) \\ \text{s.t.} & \forall s_2 \in S, \sum_{s_1 \in S} x(s_1) u_2(s_1, s_2^*) \geq \sum_{s_1 \in S} x(s_1) u_2(s_1, s_2) \\ & \sum_{s_1 \in S} x(s_1) = 1 \end{array}$$
   Finally, we take the x resulting from the "best" s_2^*.

Problem 3: The Price of Anarchy

1. **[10 pts]** Describe in words the concept of *price of anarchy* of a class of games.

   Solution: Fixing an objective function and an equilibrium concept for the class of games, the price of anarchy is the worst-case ratio between the worst objective function value of an equilibrium of the game, and that of the optimal solution.

2. **[15 pts]** Assignment 2 introduced scheduling games on *related machines*. Here we are interested in scheduling games on *unrelated machines*. The players $N = \{1, \ldots, n\}$ are associated with tasks and there is a set M of m machines. Each player chooses a machine to place their task on, that is, the strategy space of each player is M. The weights of players (or tasks) are now machine-dependent: player i has weight $w_{i\mu}$ on machine μ. A strategy profile induces an assignment $A : N \to M$ of players (or tasks) to machines. The total load on machine μ is $\ell_\mu = \sum_{i \in N:\ A(i)=\mu} w_{i\mu}$. The *cost* of player i is $\ell_{A(i)}$. Our objective function is the *makespan*, which is the maximum load on any machine: $\text{cost}(A) = \max_{\mu \in M} \ell_\mu$.

Show that for any number of players $n \geq 2$, scheduling games on unrelated machines have an unbounded (i.e., arbitrarily high) price of anarchy.

Solution: Note that there are many correct solutions to this problem. Consider a case where $m = n$ and

$$w_{i\mu} = \begin{cases} 1 & \text{if } i = \mu \\ x & \text{if } i \neq \mu \end{cases}$$

for some $x \geq 2$. Then, it is clear that the strategy profile $A(i) = i$ for all $i = 1, \dots, n$ results in the optimal (lowest) cost as $\text{cost}(A) = 1$ since $\ell_\mu = 1$ for all $\mu = 1, \dots, m$. However, consider the strategy profile B where

$$B(i) = \begin{cases} i+1 & \text{if } i \leq n-1 \\ 1 & \text{if } i = n \end{cases}$$

Here, we have that $\text{cost}(B) = x$ because $\ell_\mu = x$ for all $\mu = 1, \dots, m$. Further, note that B is a Nash equilibrium, because if any player unilaterally deviates to a different machine, their cost will increase from x to $x+1$ (if they deviate to machine i) or $2x$ (if they deviate to machine $j \neq i$). Thus, based on this example of an equilibrium for this game

$$\text{PoA} \geq \frac{x}{1} = x$$

and since we can set x to be arbitrarily large, it follows that the price of anarchy for these scheduling games on unrelated machines is unbounded.

Problem 4: The Epistemic Approach to Voting

1. **[10 pts]** Define (using mathematical notation) the *Mallows noise model*.

   Solution: The Mallows model is parameterized by $\phi \in (0, 1]$. The probability of a voter having a ranking σ given the true ranking π is:

   $$\Pr[\sigma \mid \pi] = \frac{\phi^{d_{KT}(\sigma,\pi)}}{\sum_\tau \phi^{d_{KT}(\tau,\pi)}}.$$

   Here, d_{KT} is the Kendall tau distance, and this is defined as:

   $$d_{KT}(\sigma, \sigma') = |\{(a, b) : a \succ_\sigma b \text{ and } b \succ_{\sigma'} a\}|$$

   i.e., the number of pairwise disagreements between the two rankings.

2. **[15 pts]** Construct a preference profile such that, under the Mallows model with any value of the parameter $\phi \in (0, 1)$, the ranking given by Borda count (i.e., ranking the

alternatives by Borda score, breaking ties as you wish) is *not* a maximum likelihood estimator for the ground-truth ranking.

Hint: What would Kemeny do?

Solution: First, note that under the Mallows model, we have shown in class that the MLE of the ground-truth ranking is the ranking given by the Kemeny Rule, i.e. the ranking π that minimizes the sum of the KT distances with the preference profile:

$$\text{argmin}_\pi \sum_{i \in N} d_{KT}(\sigma_i, \pi).$$

Thus, we must come up with a strategy profile such that the ranking given by Borda count disagrees with the ranking given by the Kemeny Rule. There are many such examples. Consider the following setup with 3 voters on alternatives $\{A, B, C, D\}$. Consider the following preference profile:

2 voters	1 voter
B	A
A	C
C	D
D	B

Calculating Borda scores for each alternative, we get that A gets 7 points, B gets 6 points, and C gets 4 points, and D gets 1 point. Thus, the Borda score returns a ranking $A \succ B \succ C \succ D$. Now, note that since the majority of the voters in this scenario have the ranking $B \succ A \succ C \succ D$, this is actually the output of the Kemeny rule. This is because the majority of the voters agree with each of the head-to-head comparisons made in this ranking, and so any ranking that disagrees with this ranking on a head-to-head comparison will disagree with the majority, thus increasing that ranking's total KT distance with the preference profile. Thus $B \succ A \succ C \succ D$ minimizes the KT distance with the preference profile and is therefore the MLE for the ground-truth ranking under the Mallows model, but the Borda count gives a different ranking!