Waterloo Student ID Number:

WatIAM/Quest Login UserID:

ECE 750/493 Final Exam - Winter 2024
Exam Duration: 150 Minutes

There should be 10 pages in your exam booklet. You can use the empty pages at the end of the exam booklet.

1. Stochastic Games (15 points). Consider the following two-player, zero-sum, stochastic game (only utility of agent 1 is shown).

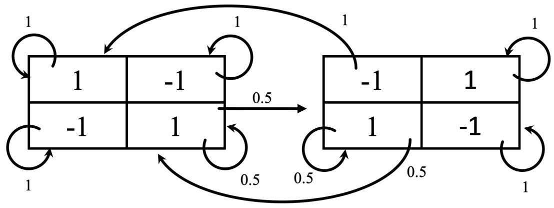

Suppose that $V_0(s_1) = 1$ and $V_0(s_2) = 2$, and $\delta = 0.5$, where s_1 is the stage game on the left, and s_2 is the stage game on the right. Compute the matrix game $G(s_1, V_0)$ (10 points) and the value $V_1(s_1)$ for agent 1 (5 points), according to the Shapley Algorithm.

$G(s_1, V_0)$

$V_1(s_1)$:

2. Commitment (15 points). Consider the following game[1]:

	L	R
U	2, 1	4, 0
D	1, 0	3, 1

In this game, D is strictly dominated by U for the row player. However, if the row player has the ability to *commit* to a pure strategy before the column player chooses their strategy, the row player should commit to D as doing so makes the column player prefer to play R, leading to a utility of 3 for the row player. By contrast, if the row player were to commit to playing U, the column player would prefer to play L, leading to a utility of only 2 for the row player.

If the row player is able to commit to a mixed strategy, they can get an even greater (expected) utility: Suppose that the row player commits to playing D with probability p. If $p > 1/2$, then the column player will prefer to play R, and the row player's expected utility will be $3p + 4(1 - p) = 4 - p \geq 3$. If $p = 1/2$, the column player is indifferent between L and R. In such cases, we will assume that the column player will choose the strategy that maximizes the row player's utility (in this case, R). Hence, the optimal mixed strategy to commit to for the row player is $p = 1/2$. There are a few good reasons for assuming that the column player will choose the strategy that maximizes the row player's utility.

If we were to assume the opposite, then there would not exist an optimal strategy for the row player in our example game: the row player would play D with probability $p = 1/2 + \epsilon$ with $\epsilon > 0$, and the smaller ϵ, the better the utility for the row player. Additionally, with this assumption, it is enough to only consider a pure strategy for the follower. This is because if the follower randomizes between multiple pure strategies in response to the leader's strategy, then all those pure strategies must yield the same utility to the follower and also to the leader (due to the tie-breaking rule). Therefore, we can just focus on one of those pure strategies as all of them produce the same outcome.

Now, consider a more general setting with two players: the leader ℓ and the follower f. The set of pure strategies for ℓ and f are denoted by A_ℓ and A_f, respectively. The leader's utility is given by a function $u_\ell : A_\ell \times A_f \mapsto \mathbb{R}$. Similarly, the follower's utility is given by a function $u_f : A_\ell \times A_f \mapsto \mathbb{R}$. The leader commits to a (mixed) strategy to which the follower best responds. Our goal is to find the equilibrium of this commitment game. A pair of strategies form an equilibrium if they satisfy all the following conditions:

- The leader plays a best-response to the follower's pure strategy;
- The follower plays a best-response to the leader's (mixed) strategy; and
- The follower breaks ties in favor of the leader.

[1] From Conitzer, Vincent, and Tuomas Sandholm. "Computing the optimal strategy to commit to." Proceedings of the 7th ACM conference on Electronic commerce. 2006.

To find such an equilibrium, we use mixed-integer linear programming (MILP).

$$
\begin{array}{lll}
\text{Max.} & U_\ell, & \\
\text{s.t.} & p_{a_\ell} \in [0,1] & \forall a_\ell \in A_\ell, \\
& \displaystyle\sum_{a_\ell \in A_\ell} p_{a_\ell} = 1, & \\
& q_{a_f} \in \{0,1\} & \forall a_f \in A_f, \\
& \displaystyle\sum_{a_f \in A_f} q_{a_f} = 1, & \\
& U_\ell - \displaystyle\sum_{a_\ell \in A_\ell} u_\ell(a_\ell, a_f) p_{a_\ell} \le L(1 - q_{a_f}) & \forall a_f \in A_f, \\
& (\cdots) & \forall(\cdots),
\end{array}
$$

where L is a constant number with an arbitrary large value. In this formulation, the variables p_{a_ℓ} determine the mixed strategy of the leader. The first and second constraints ensure that these variables define a probability distribution. Similarly, the variables q_{a_f} determine the pure strategy of the follower. The third and fourth constraints ensure that only a single pure strategy is selected.

In addition to p_{a_ℓ}'s and q_{a_f}'s, we have two other variables: U_ℓ and U_f to capture the expected utility of the leader and follower in the equilibrium, respectively. For the leader, this is enforced by the fifth constraint and the objective function: when $q_{a_f} = 1$, U_ℓ becomes equal to $\sum_{a_\ell \in A_\ell} u_\ell(a_\ell, a_f) p_{a_\ell}$.

What should be the last set of constraints to make sure that the follower best responds to the leader. In other words, what should be the last set of constraints so that U_f is equal to the expected utility of the follower in the equilibrium. Write down your answer in the box below **and explain it**.

Note that your answer could encapsulate multiple inequalities in one formula in the form of $xx \le yy \le zz$. If you need to add more than one formula, then add another set of boxes and write your answer in those.

3. Minmax and maxmin (15 points). Consider the following game:

		Agent 2	
		L	R
Agent 1	U	$20, -20$	$-10, 5$
	D	$-5, 15$	$25, -10$

a. What is the maxmin and minmax values of Agent 1 (each 7.5 points)? Show your work.

Maxmin: $\boxed{}$ Minmax: $\boxed{}$

4. All-pay Auction (15 points). Suppose that there are N bidders and a single object for sale. Bidder i has value v_i for the object and bids b_i. The object is allocated to the bidder with the highest bid. ALL bidders pay their bids regardless of who wins. Suppose v's are independent and identically distributed random variables from $[0, 1]$ with a commonly known cumulative distribution function (CDF) $F(v) = v$.

a. What is the unique, symmetric, increasing, and differentiable equilibrium strategy $\beta(v)$? Show your work.

5. Correlated equilibrium (20 points) Consider the following game and the distribution π on its outcomes.

	A	B	C
A	$1, 1$ 30%	$-1, -1$ 0%	$0.5, 0$ 5%
B	$-1, -1$ 0%	$1, 1$ 30%	$0, 0$ 0%
C	$0, 0.5$ 5%	$0, 0$ 0%	$-1.5, -1.5$ 30%

a. Is π a coarse correlated equilibrium of the game? Why (10 points)?

b. Is π a correlated equilibrium of the game? Why (10 points)?

6. BNE (20 points). Consider the following game:

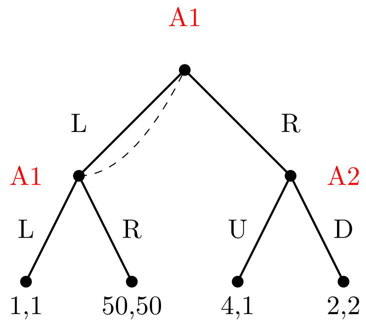

a. What is the expected utility of A1 in the mixed-strategy NE of this game (5 points)? Show your work.

b. What is the expected utility of A1 in the behavioral NE of this game (10 points)? Show your work.