Assignment 2

Due on Tuesday March. 3^{rd}, 23:55 EST

- Please read the rules for assignments on the course webpage: (`https://ece.uwaterloo.ca/~smzahedi/crs/ece752/`).
- Use Piazza or directly contact the instructor (smzahedi@uwaterloo.ca) with any questions.
- For all questions, you must show your work. Final answers alone will not receive full credit.

1. (20 points) Consider a market with exactly two firms producing identical products. Each firm chooses a quantity $q_i \in [0, 5]$. The market price is determined by total output according to

$$P(q_1, q_2) = \max\{6 - (q_1 + q_2), 0\}.$$

Both firms have zero production costs and therefore seek to maximize profit:

$$u_i(q_i, q_{-i}) = q_i P(q_i + q_{-i}).$$

a. (5 points) Suppose the firms choose their quantities simultaneously. Find the unique Nash equilibrium.

b. (5 points) Now suppose the game is sequential: firm 1 first chooses q_1, which is observed by firm 2, and then firm 2 chooses q_2. Is the Nash equilibrium outcome from part (a) still a Nash equilibrium outcome of this extensive-form game? Is it unique? Explain your reasoning.

c. (10 points) Find the subgame perfect equilibrium (SPE) of this sequential game. How does it compare to the equilibrium outcome(s) identified in part (b)? Briefly explain any differences.

2. (10 points) Consider the following game and the probability distribution π over its outcomes.

	A	B	C
A	$1, 1$ 35%	$-1, -1$ 0%	$0, 0$ 0%
B	$-1, -1$ 0%	$1, 1$ 35%	$0, 0$ 0%
C	$0, 0$ 0%	$0, 0$ 0%	$-1.5, -1.5$ 30%

a. (5 points) Is π a coarse correlated equilibrium (CCE) of the game? Justify your answer.

b. (5 points) Is π a correlated equilibrium (CE) of the game? Justify your answer.

3. (10 points) There are three agents bargaining over the division of one dollar. The bargaining protocol is as follows. At the beginning of round 1, agent i is selected as the proposer with probability p_i, where $p_1 + p_2 + p_3 = 1$ and $p_1 < p_2 < p_3$. The proposer makes an offer (x_1, x_2, x_3), where $x_j \geq 0$ for all j, and $x_1 + x_2 + x_3 \leq 1$. The agents then vote on the proposal: if it receives **three** votes in favor (assume the proposer always votes in favor), the proposal is accepted and the split is implemented. Otherwise, the game proceeds to round 2.

Round 2 follows the same procedure as round 1, with proposer selection again determined by the same probabilities (independently of round 1). However, round 2 is the final round–if the proposal fails to receive unanimous approval, all agents receive a payoff of 0. Each agent discounts payoffs in round 2 proportionally by a common discount factor $\delta \in [0, 1]$.

a. (5 Points) What are the expected utilities of each agent in the unique equilibrium of the second round?

b. (5 Points) What are the expected utilities of each agent in the unique SPE of the game?

3. (10 points). Consider the following game:

	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 can *commit* to a pure strategy before the column player moves, it is optimal to commit to D. Doing so induces the column player to choose R, yielding a payoff of 3 to the row player. By contrast, committing to U induces the column player to choose L, yielding a payoff of only 2.

If the row player can instead commit to a mixed strategy, an even higher expected payoff can be achieved. Suppose the row player commits to playing D with probability p. The column player's expected payoff from L is

$$(1 - p) \cdot 1 + p \cdot 0 = 1 - p,$$

while from R it is

$$(1 - p) \cdot 0 + p \cdot 1 = p.$$

Thus, the column player prefers R whenever $p > 1/2$, prefers L whenever $p < 1/2$, and is indifferent when $p = 1/2$.

If $p > 1/2$, the row player's expected payoff (inducing R) is

$$(1 - p) \cdot 4 + p \cdot 3 = 4 - p.$$

This expression is decreasing in p, so the optimal commitment is $p = 1/2$, yielding payoff $4 - 1/2 = 3.5$.

When $p = 1/2$, the column player is indifferent between L and R. We assume that ties are broken in favor of the leader (row player), so the column player selects the action maximizing the leader's payoff (here, R).

This tie-breaking assumption ensures existence of an optimal commitment strategy. Without it, no optimal strategy would exist in this example: the leader would choose $p = 1/2 + \varepsilon$ with $\varepsilon > 0$, and smaller ε would strictly improve the payoff. Moreover, under this assumption, it suffices to consider pure best responses by the follower. If the follower randomizes among multiple best responses, all such actions must yield the same payoff to both players (by indifference and the tie-breaking rule), so it is without loss of generality to select any one of them.

We now consider the general setting with two players: a leader ℓ and a follower f. Let A_ℓ and A_f denote their respective sets of pure strategies. The utility functions are

$$u_\ell : A_\ell \times A_f \to \mathbb{R}, \qquad u_f : A_\ell \times A_f \to \mathbb{R}.$$

The leader commits to a (mixed) strategy, and the follower best responds after observing it.

An equilibrium of this commitment game consists of a mixed strategy for the leader and a pure strategy for the follower such that:

- The follower plays a best response to the leader's mixed strategy;
- The follower breaks ties in favor of the leader; and
- The leader's strategy maximizes u_ℓ given the follower's induced response.

To compute such an equilibrium, we formulate the problem as a mixed-integer linear program.

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

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.