Math 454, Spring 2026, USC — Instructor: Steven Heilman

Name: ________________ USC ID: ____________ Date: __________

Signature: ________________. Discussion Section: ________________
(By signing here, I certify that I have taken this test while refraining from cheating.)

Final Exam

This exam contains 14 pages (including this cover page) and 8 problems. Enter all requested information on the top of this page.

You may *not* use your books, notes, or any calculator on this exam.

You are required to show your work on each problem on this exam. The following rules apply:

- You have 50 minutes to complete the exam, starting at the beginning of class.
- **Organize your work**, in a reasonably neat and coherent way, in the space provided. Work scattered all over the page without a clear ordering will receive very little credit.
- **Mysterious or unsupported answers will not receive full credit**. A correct answer, unsupported by calculations, explanation, or algebraic work will receive no credit; an incorrect answer supported by substantially correct calculations and explanations might still receive partial credit.
- If you need more space, use the back of the pages; clearly indicate when you have done this. Scratch paper appears at the end of the document.

Do not write in the table to the right. Good luck![a]

Problem	Points	Score
1	15	
2	10	
3	15	
4	10	
5	10	
6	10	
7	10	
8	15	
Total:	95	

Reference sheet. Below are some definitions that may be relevant.

$$\Delta_m := \{x = (x_1, \ldots, x_m) \in \mathbf{R}^m \colon \sum_{i=1}^{m} x_i = 1,\ x_i \geq 0,\ \forall\, 1 \leq i \leq m\}.$$

Let A be an $m \times n$ real payoff matrix defining a zero-sum two-player game. A mixed strategy $\widetilde{x} \in \Delta_m$ is **optimal for player** I if $\min_{y\in\Delta_n} \widetilde{x}^T A y = \max_{x\in\Delta_m} \min_{y\in\Delta_n} x^T A y$. A mixed strategy $\widetilde{y} \in \Delta_n$ is **optimal for player** II if $\max_{x\in\Delta_m} x^T A \widetilde{y} = \min_{y\in\Delta_n} \max_{x\in\Delta_m} x^T A y$. We say the pair $(\widetilde{x}, \widetilde{y})$ are an **optimal strategy** for the payoff matrix A if $\widetilde{x} \in \Delta_m$ is optimal for player I and $\widetilde{y} \in \Delta_n$ is optimal for player II.

Let m, n be positive integers. Suppose we have a two-player general sum game with $m \times n$ payoff matrices. Let A be the payoff matrix for player I and let B be the payoff matrix for player II. A pair of vectors $(\widetilde{x}, \widetilde{y})$ with $\widetilde{x} \in \Delta_m$ and $\widetilde{y} \in \Delta_n$ is a **Nash equilibrium** if

$$\widetilde{x}^T A \widetilde{y} \geq x A \widetilde{y}, \quad \forall\, x \in \Delta_m, \qquad \widetilde{x}^T B \widetilde{y} \geq \widetilde{x} B y, \quad \forall\, y \in \Delta_n.$$

A joint distribution of strategies is an $m \times n$ matrix $z = (z_{ij})_{1\leq i\leq m, 1\leq j\leq n}$ such that $z_{ij} \geq 0$ for all $i \in \{1, \ldots, m\}$, $j \in \{1, \ldots, n\}$, and such that $\sum_{i=1}^{m} \sum_{j=1}^{n} z_{ij} = 1$. We say z is a **correlated equilibrium** if

$$\sum_{j=1}^{n} z_{ij} a_{ij} \geq \sum_{j=1}^{n} z_{ij} a_{kj}, \qquad \forall\, i \in \{1, \ldots, m\}, \forall\, k \in \{1, \ldots, m\}.$$

$$\sum_{i=1}^{m} z_{ij} b_{ij} \geq \sum_{i=1}^{m} z_{ij} b_{ik}, \qquad \forall\, j \in \{1, \ldots, n\}, \forall\, k \in \{1, \ldots, n\}.$$

Suppose we have a two-player symmetric game (so that the payoff matrix for player I is A, the payoff matrix for player II is B, and with $A = B^T$). Assume that A, B are $n \times n$ matrices. A mixed strategy $x \in \Delta_n$ is said to be an **evolutionarily stable strategy (ESS)** if, for any pure strategy w, we have

$$w^T A x \leq x^T A x, \qquad \text{and} \qquad \text{If } w^T A x = x^T A x, \text{ then } w^T A w < x^T A w.$$

Suppose we have a game with n players together with a characteristic function $v \colon 2^{\{1,\ldots,n\}} \to \mathbf{R}$. For each $i \in \{1, \ldots, n\}$, we define the **Shapley value** $\phi_i(v) \in \mathbf{R}$ to be any set of real numbers satisfying the following four axioms:

(i) (Symmetry) If for some $i, j \in \{1, \ldots, n\}$ we have $v(S \cup \{i\}) = v(S \cup \{j\})$ for all $S \subseteq \{1, \ldots, n\}$ with $i, j \notin S$, then $\phi_i(v) = \phi_j(v)$.

(ii) (No power/ no value) If for some $i \in \{1, \ldots, n\}$ we have $v(S \cup \{i\}) = v(S)$ for all $S \subseteq \{1, \ldots, n\}$, then $\phi_i(v) = 0$.

(iii) (Additivity) If u is any other characteristic function, then $\phi_i(v + u) = \phi_i(v) + \phi_i(u)$, for all $i \in \{1, \ldots, n\}$.

(iv) (Efficiency) $\sum_{i=1}^{n} \phi_i(v) = v(\{1, \ldots, n\})$.

A set $K \subseteq \mathbf{R}^n$ is called **convex** if, for any $x, y \in K$ and for any $0 \leq t \leq 1$, we have $tx + (1-t)y \in K$. A set $K \subseteq \mathbf{R}^n$ is called **bounded** if there exists $r > 0$ such that $||x|| \leq r$ for all $x \in K$.

We define a **symmetric auction**. A single object is for sale at an auction. The seller is willing to sell the object at any nonnegative price. There are n buyers, which we identify with the set $\{1, 2, \ldots, n\}$. All buyers have some set of **private values** in $[0, \infty)$. We denote the private value of buyer $i \in \{1, \ldots, n\}$ by V_i, so that V_i is a random variable that takes nonnegative real values. We assume that all of the random variables $V_1, \ldots, V_n$ are independent. We also assume that $V_1, \ldots, V_n$ are identically distributed, with a continuous density function. That is, there exists some continuous function $f\colon \mathbf{R} \to [0, \infty)$ with $\int_{-\infty}^{\infty} f(x)dx = 1$ such that: for each $i \in \{1, \ldots, n\}$, for each $t \in \mathbf{R}$, the probability that $V_i \leq t$ is equal to $\int_{-\infty}^{t} f(x)dx$. We also assume that all buyers are **risk-neutral**, so that each buyer seeks to maximize their expected profits.

Finally, we assume that all of the above assumptions are **common knowledge**. That is, every player knows the above assumptions; every player knows that every player knows the above assumptions; every player knows that every player knows that every player knows the above assumptions; etc.

Under the above assumptions, a **pure strategy** for Player $i \in \{1, \ldots, n\}$ is a function $\beta_i\colon [0, 1] \to [0, \infty)$. So, if Player i has a private value of V_i, he will make a bid of $\beta_i(V_i)$ in the auction. (We will not discuss mixed strategies in auctions.)

Given the strategies $\beta = (\beta_1, \ldots, \beta_n)$, and given any $v \in [0, 1]$, Player i has expected profit $P_i(\beta, v)$, if her private value is v. (If buyer i wins the auction, and if buyer i has private value v and bid b, then the profit of buyer i is $v - b$.) We say that a strategy β is an **equilibrium** if, given any $v \in [0, 1]$, any $b \geq 0$, and any $i \in \{1, \ldots, n\}$,

$$P_i(\beta, v) \geq P_i((\beta_1, \ldots, \beta_{i-1}, b, \beta_{i+1}, \ldots, \beta_n), v).$$

Let n be a positive integer. Let $x = (x_1, \ldots, x_n) \in \{-1, 1\}^n$. Let $f, g\colon \{-1, 1\}^n \to \mathbf{R}$. For any subset $S \subseteq \{1, \ldots, n\}$, define a function $W_S\colon \{-1, 1\}^n \to \mathbf{R}$ by $W_S(x) := \prod_{i \in S} x_i$. Define also the inner product $\langle f, g \rangle := 2^{-n} \sum_{x \in \{-1,1\}^n} f(x)g(x)$. Any $f\colon \{-1, 1\}^n \to \mathbf{R}$ can be expressed as $f(x) = \sum_{S \subseteq \{1,\ldots,n\}} \langle f, W_S \rangle W_S(x)$. For any $S \subseteq \{1, \ldots, n\}$, if we denote $\widehat{f}(S) := \langle f, W_S \rangle = 2^{-n} \sum_{y \in \{-1,1\}^n} f(y)W_S(y)$, then we have $f(x) = \sum_{S \subseteq \{1,\ldots,n\}} \widehat{f}(S)W_S(x)$.

The **noise stability** of f with parameter $\rho \in (-1, 1)$ is defined to be $\sum_{S \subseteq \{1,\ldots,n\}} \rho^{|S|} |\widehat{f}(S)|^2$.

Let $\mathcal{A} = \{1, \ldots, k\}$. Let $(P_a)_{a \in \mathcal{A}}$ be a set of probability distributions. Let H be a positive integer or ∞. A **bandit** problem is a two-player game with incomplete information played over H rounds. For each round $0 \leq t \leq H$ of the game, the learner chooses an action $A_t \in \mathcal{A}$. The learner then obtains a reward R_t, where R_t is sampled from P_{A_t}.

Since the game occurs over many rounds, A_t is a function of the history $A_0, \ldots, A_{t-1}$, $R_0, \ldots, R_{t-1}$ at time t. A **policy** π is a function whose input is the history (at any time t) and whose output is an action in $\mathcal{A}$.

For any $a \in \mathcal{A}$, let μ_a be the expected value of a random variable with distribution P_a. The **expected regret** r_H at time H is:

$$H \max_{a \in \mathcal{A}} \mu_a - \sum_{t=0}^{H-1} \mathbf{E} R_t.$$

1. Label the following statements as TRUE or FALSE. If the statement is true, **EXPLAIN YOUR REASONING**. If the statement is false, **PROVIDE A COUNTEREXAMPLE AND/OR EXPLAIN YOUR REASONING**.

   (a) (3 points) In the game of Nim, the first player always has a winning strategy.

   TRUE FALSE (circle one)

   (b) (3 points) Consider the problem of deciding if a general sum game has at least two Nash equilibria. This problem is NP-complete.

   TRUE FALSE (circle one)

   (c) (3 points) There exists a symmetric two-person general-sum game such that all of its Nash equilibria are not symmetric. (The $n \times n$ payoff matrices A, B satisfy $A = B^T$, and all equilibria $(x, y) \in \Delta_n \times \Delta_n$ satisfy $x \neq y$.)

   TRUE FALSE (circle one)

(d) (3 points) Let $f\colon \{-1,1\}^n \to \{-1,1\}$ be a function. Suppose we do a Condorcet election with f (so that if we just look at the votes between any pair of two candidates, the aggregate preference is decided using the function f). Suppose a Condorcet winner always exists. Then there exists $i \in \{1,\ldots,n\}$ such that $f(x_1,\ldots,x_n) = x_i$ for all $(x_1,\ldots,x_n) \in \{-1,1\}^n$.

TRUE FALSE (circle one)

(e) (3 points) Let $f, g\colon \{-1,1\}^n \to \{-1,1\}$ be functions. Assume that n is odd and f is the majority function ($f(x_1,\ldots,x_n) = \text{sign}(x_1+\cdots+x_n)$). Assume $f \neq g$. Then for any $\rho \in (0,1)$, the noise stability of f exceeds the noise stability of g.

TRUE FALSE (circle one)

2. (10 points) Recall the prisoner's dilemma, which has the following payoffs.

Prisoner I \ Prisoner II	silent	confess
silent	$(-1, -1)$	$(-10, 0)$
confess	$(0, -10)$	$(-8, -8)$

Find all Nash equilibria for this game.

3. For all questions below, JUSTIFY YOUR ANSWER.

   (a) (5 points) Give an example of a closed and convex subset K of Euclidean space, and give an example of a continuous function $f\colon K \to K$ such that f has no fixed point.

   (b) (5 points) Give an example of a bounded and closed subset K of Euclidean space, and give an example of a continuous function $f\colon K \to K$ such that f has no fixed point.

   (c) (5 points) Give an example of a function $f\colon \mathbf{R} \to \mathbf{R}$ such that the only fixed point of f is the point $x = 1$.

4. (10 points) Ten people are standing together in a room. They are presented with the following problem: each person chooses a real number between (and including) 0 and 100. (So, someone could guess: 20, 51.5, π, $\sqrt{2}$, 99.999, etc.) The person who chooses the number closest to two-thirds of the average of all of the numbers wins. That is, if

$$0 \leq a_1, \ldots, a_{10} \leq 100,$$

each person wants to choose a number closest to $(2/3)(1/10)\sum_{i=1}^{10} a_i$. The people do not communicate with each other in any way. It is common knowledge that every person wants to win the game, and every person is rational. Explain what number each person will choose.

5. (10 points) Explain in detail the Condorcet voting paradox. You should probably use an example of three voters ranking three candidates in order to explain the paradox.

6. (10 points) Let n be a positive integer. Let $v\colon 2^{\{1,\ldots,n\}} \to \mathbf{R}$ be a characteristic function (so that $v(\emptyset) = 0$). For any $1 \leq i \leq n$, define

$$B_i(v) := 2^{-n} \sum_{S\subseteq\{1,\ldots,n\}\colon\, i\notin S} |v(S) - v(S \smallsetminus \{i\})|.$$

(Here $S \smallsetminus \{i\} = \{s \in S\colon s \neq i\}$.)

- Which of Shapley's four axioms do the values $(B_1, \ldots, B_n)$ satisfy? JUSTIFY YOUR ANSWER.
- Is B_i equal to the i^{th} Shapley value for each $1 \leq i \leq n$? JUSTIFY YOUR ANSWER.

7. (10 points)

- Describe in detail the Upper Confidence Bound Algorithm.

  The input for this algorithm is a bandit problem with k arms. The output for this algorithm is a policy (strategy) for the bandit problem.

- Describe in detail a bound for the expected regret r_H of this algorithm, when the rewards satisfy $0 \leq R_t \leq 1$ for a k-arm bandit problem with horizon H. (You do NOT have to prove this regret bound holds, just state what the bound is.)

Hint: let $N_t(a)$ be the number of times action $a \in \mathcal{A}$ has been played at time t, let $M_{t,a}$ be the sample mean of the rewards obtained at time t from action a being played. For any $t \geq 0$ where $N_t(a) \geq 1$, define

$$UCB_t(a) := M_{t,a} + \sqrt{\frac{2 \log H}{N_t(a)}}.$$

8. (15 points) Suppose we have a two-player symmetric game with payoffs described by a matrix such as the following.

$$A = \begin{pmatrix} 4 & 3 & 2 & 5 & 6 \\ 3 & 1 & 8 & 9 & 1 \\ 7 & 0 & 7 & 0 & 7 \\ 1 & 3 & 3 & 2 & 1 \\ 0 & 1 & 2 & 9 & 1 \end{pmatrix}$$

Describe in detail an algorithm that outputs a Nash equilibrium of this game.

Make sure to justify why your algorithm outputs a Nash equilibrium.

Also give a bound on the run time of the algorithm.

Hint: you can freely use that, if $(\widetilde{x}, \widetilde{y}) \in \Delta_m \times \Delta_n$ is a Nash equilibrium and if

$$I := \{1 \leq i \leq m : \widetilde{x}_i > 0\}, \qquad J := \{1 \leq j \leq n : \widetilde{y}_j > 0\}. \qquad (*)$$

then with payoff matrices A, B respectively, we have

$$\max_{i=1,\ldots,m} (A\widetilde{y})_i = (A\widetilde{y})_i, \quad \forall\, i \in I. \qquad \text{and} \qquad \max_{j=1,\ldots,n} (\widetilde{x}^T B)_j = (\widetilde{x}^T B)_j, \quad \forall\, j \in J.$$

(Scratch paper)