CPSC 532A, Winter 2006
Homework #2

---

1. **[10 points]** (**Perfect Information Games**)

   Consider the centipede game in Figure 1. It differs from the one appearing in the course reader only in that the payoff pair (4,3) has been changed to (5,3).

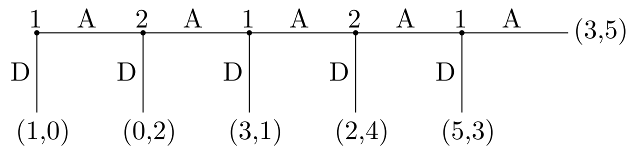

Figure 1: A centipede game.

   Player 2 is one of two types: (i) With probability p player 2 is a rational player who follows the unique subgame perfect equilibrium strategy. (ii) With probability $1 - p$ player 2 is "irrational" and simply flips a fair coin at each of his choice points. For every possible value of p, find a best response (possibly mixed) strategy for player 1 to this player 2. (Obviously there will be ranges of p for which a strategy is always a best response.) Show your work.

2. **[10 points]** (**Imperfect Information Games**)

   Each part of this problem will use the two-player game of imperfect information given in Figure 2. However, the meaning of the numbers at the leaves will differ. In part (a), we consider a common-payoff game. Thus, the value at a leaf defines the payoff of both players. In parts (b) and (c), we switch to a zero-sum game. In that case, the value of a leaf defines the payoff of player 1, and the negative of the payoff of player 2. In each part, briefly justify your answer.

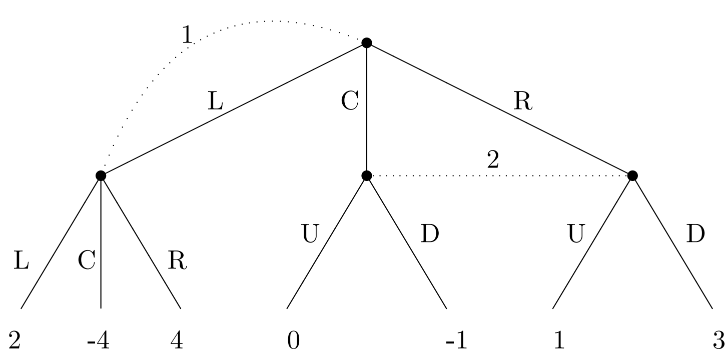

Figure 2: An imperfect information game in which each player has a single information set.

   (a) For the *common-payoff* game defined by Figure 2, list all Nash equilibrium *pure* strategy profiles ("none exists" is a possible answer).

(b) For the *zero-sum* game defined by Figure 2, list all Nash equilibrium *pure* strategy profiles ("none exists" is a possible answer).

(c) Now we will allow mixed strategies. For the *zero-sum* game defined by Figure 2, find all Nash equilibrium (possibly mixed) strategy profiles. Note, listing all Nash equilibria is rather arduous considering that there are an infinite number of equilibria. Instead, you are supposed to characterize the set of all mixed strategy Nash equilibria (for example, using a variable p and giving ranges on p for which the equilibria holds). Two minor hints here: (1) the characterization is simple, and (2) use the fact that the game is zero-sum to limit the space of strategy profiles you have to consider.

3. **[20 points] (Repeated Prisoners' Dilemma)**

Consider the Prisoners dilemma game. Specifically, the following game is going to be played repeatedly:

		Player 2	
		c	d
Player 1	C	-1,-1	-4,0
	D	0,-4	-3,-3

(a) Suppose the game will be played three times.

i. Find (the only) subgame perfect Nash equilibrium. (HINT: make sure you write the equilibrium completely and precisely.)

ii. How does the equilibrium change if the game is repeated n times? (n is common knowledge.)

(b) Now suppose that the game is going to be repeated forever. Suppose the overall payoff is the future discounted rewards from individual games, with a discount factor $0 < \beta < 1$. Construct a set of Nash equilibrium strategies that would lead to both players cooperating every period, and give the values of β for which your solution holds.

(c) Now suppose there is no discounting ($\beta = 1$) but each period there is a probability 9/10 that the game continues and a probability of 1/10 that the game will finish. Are the strategies proposed in part (b) still a Nash equilibrium?

(d) An alternative commonly used way to calculate payoffs in an infinitely repeated game is the *limit of the means* reward, sometimes also known as the average reward.

For the pair of strategies you used in part (b), is the payoff in the infinitely repeated game under the limit of means criterion well defined? If so, what is it? Do your strategies still constitute a Nash equilibrium?

4. **[15 points]** Consider a problem where two students must simultaneously decide between working on their research in their (separate) offices and going to Koerner's. Each student has a preference for one of the two choices. The students don't know each other's preferences, but know that they are drawn from a commonly known joint distribution. This distribution is described in Table 1. Starting from a baseline utility of zero, a student gains 2 unit of utility if she goes to the place that she prefers. However, the students are working on a course project together, and so both students lose 3 units of utility if they

both attend the bar and reveal to each other that they were slacking off (independent of whether they gained 1 unit of utility based on their preference). Thus, for example, if they both prefer bar, and they both go to the bar, they each get a utility of $0 + 2 - 3 = -1$.

Student 1	Student 2	Probability
b_1	b_2	0.1
b_1	$\neg b_2$	0.5
$\neg b_1$	b_2	0.1
$\neg b_1$	$\neg b_2$	0.3

Table 1: The common prior joint distribution on student preferences. b_i means that student i prefers to go to the bar, $\neg b_i$ means she prefers to work in the lab.

(a) Model the setting as a Bayesian game. Recall that you need a set of agents N, a set of actions A, a set of types for each agent Θ_i, a probability function mapping from one agent's type to a distribution over the types of the other agent(s) $p_i : \Theta_i \to \Delta(\Theta_{-i})$, and a payoff function for each agent mapping from the agents' joint actions and types to a real number $u_i : A \times \Theta \to \mathbb{R}$. Denote by B and L the actions of going to the bar and staying in the lab respectively. types, and $I = \{I_1, I_2\}$ are the partitions over games for the two agents. Your entire answer can be a figure similar to Figure 3.21, which shows the games, the common prior, and the partitions of the agents. Denote by B and L the actions of going to the bar and staying home, respectively.

(b) Find all Bayes-Nash equilibria of this game.

(c) Draw the payoff matrix of the induced normal form of the game and justify why your equilibrium/equilibria hold(s). Explicitly state the meaning of an action in the induced normal form game; please write the actions in alphabetical order.

5. **[10 points]** Folk Theorem

Show that any payoff profile that can be achieved in a correlated equilibrium can also be achieved in a Nash Equilibrium of the infinitely repeated game (for average rewards).

6. **[25 points]** Consider the following voting scheme:

Each agent submits a total preference ordering, along with additional information of that agent's utility for each outcome. (Obviously, the preference ordering reflects the ordering of the outcomes by utility, with the highest utility outcome being most preferred.) Let the utility for each outcome be an element of $[0, 100]$. The social welfare function orders each outcome by the sum of the utilities of that outcome for each agent. In the case of ties, the outcome with the earlier lexiographic ordering is preferred.

(a) Is this Pareto optimal? Justify.

(b) Is this independent of irrelevant utilities? (In this setting, define IIU as the proposition that the social ordering of o and o' does not change as long as agents do not change their utilities for o and o'.) Justify.

(c) Is this non dictatorial? Justify.

(d) Does this voting scheme contradict Arrow's theorem? If yes, explain why; if not, demonstrate that Arrow's theorem is not violated.

(e) Is it incentive compatible? Why or why not?

7. **[20 points]** Consider a social choice problem with outcome set $\mathcal{O}$, $|\mathcal{O}| \geq 3$, where agents may have any possible ordering over preferences. Let $a \succ_i b$ denote the proposition that i prefers a to b, for $a, b \in \mathcal{O}$. Let θ denote the set of declared preferences, and let a^* denote an arbitrary (but commonly-known) element of $\mathcal{O}$. Prove from first principles (i.e., without referring to an impossibility theorem) that the following social choice function cannot be implemented in a dominant-strategy equilibrium by any direct or indirect mechanism.

$$f(\theta) = \begin{cases} a & \text{if } \forall i, \forall b \neq a,\ a \succ_i b \\ a^* & \text{otherwise} \end{cases}$$