CPSC 532A, Fall 2006
Homework #1

1. **[10 points] (Utility Theory)**

   (a) In 1738, J. Bernoulli investigated the St. Petersburg paradox, which works as follow. You have the opportunity to play a game in which a fair coin is tossed repeatedly until it comes up heads. If the first heads appears on the nth toss, you win 2^n dollars.

   Show that the expected monetary value of this game is infinite.

   If we assume that utilility increases monotonically with money, is this still consistent with the assertion by Von Neumann and Morgenstern that utility for outcomes must fall within a bounded interval? (Hint: What does this imply about an agent's utility for money?)

   (b) If we remove the Decomposability ("No fun in gambling") axiom, the proof of Von Neumann Morgenstern utility theory (from Appendix C of the textbook) is no longer valid. Demonstrate this by constructing a situation consistent with the other axioms (Completeness, Transitivity, Substitutability and Monotonicity) where the agent's utility for some outcome o is undefined. (Use the same notation as in the textbook.)

2. **[10 points] (Normal Form Games)**

   Consider the following game:

		Player 2 C	D
Player 1	A	-21,0	10,10
	B	6,6	0,-21

   Where the first number in each square is the payoff of player 1 and the second number is player 2's payoff.

   (a) Find all Pareto optimal pure strategy profiles.

   (b) Find the pure strategy Nash Equilibria.

   (c) Find all mixed-strategy Nash Equilibria.

   (d) Which of the above equilibria do you prefer? Suppose player 2 has decided to play according to one of the equilibria that you found in part (b) (but you do not know which.) What would you play as player 1?

3. **[10 points] (More Normal Form Games)**

   Consider the following game:

		Player 2 L	R
Player 1	T	a,e	b,f
	B	c,g	d,h

Where the first number in each square is the payoff of player 1 and the second number is player 2's payoff.

(a) What (in)equalities must hold for the game to have one pure strategy Nash Equilibrium (BL), which is pareto dominated?

(b) What (in)equalities must hold for a Nash Equilibrium to exist where player 1 only plays T but player 2 mixes over L and R (playing L with probability p)?

(c) In general, what must be true of an action A (in terms of dominance) before it can be in the support of a mixed strategy Nash Equilibrium?

(d) In general, what must be true of an action A before it *must* be part of any Nash Equilibrium?

4. **[20 points] (Maxmin and Minmax)**
Consider a game with n players. Denote the maxmin strategy for player i as $\bar{s}_i = \arg\max_{s_i} \min_{s_{-i}} u_i(s_i, s_{-i})$; this is the mixed strategy that i can play to maximize his reward if he believes that his opponent(s) are trying to harm him. Denote the maxmin value of i as $\bar{v}(i) = \max_{s_i} \min_{s_{-i}} u_i(s_i, s_{-i})$. Denote the minmax strategy of some agent $j \neq i$ against i as $\underline{s}_j$, which is defined as j's component of the mixed strategy profile s_{-i} in the expression $\arg\min_{s_{-i}} \max_{s_i} u_i(s_i, s_{-i})$. Intuitively, $\underline{s}_j$ is the mixed strategy profile that player j should play to minimize i's expected payoff. Observe that when there are more than two agents, j is not able to guarantee that i achieves minimal payoff by himself; the computation of the minmax strategy presumes that other agents will also choose their strategies to minimize i's payoff. Denote the minmax value of i as $\underline{v}(i) = \min_{s_{-i}} \max_{s_i} u_i(s_i, s_{-i})$. Here $-i$ denotes the set of players other than i.

(a) Prove that for all games, the maxmin of player i is no greater than the minmax of player i, i.e. $\bar{v}(i) \leq \underline{v}(i)$.

(b) Prove that in all two-player games the maxmin of player i is equal to the minmax of player i, in other words $\bar{v}(i) = \underline{v}(i)$. Hint: you can use the minmax theorem, but note that it only applies to two-player *zero-sum* games.

(c) Now we demonstrate that the result in (b) does not apply to n-player games with $n > 2$, by the following counterexample. Consider the following three-player game; player 1 chooses the column, player 2 chooses the row, and player 3 chooses the matrix.

	L	R
T	1,4,0	1,2,-2
B	5,6,0	5,5,0

U

	L	R
T	3,3,0	2,6,0
B	3,4,-2	5,3,0

D

Compute $\bar{v}(3)$, and show that $\bar{v}(3) < \underline{v}(3)$. (Hint: the minmax value is hard to calculate for this game, but you don't need to compute it exactly in order to show that $\bar{v}(3) < \underline{v}(3)$.)

5. **[15 points] (Rationalizability, Correlated Equilibria)** Consider the following two-player game:

		Player 2		
		D	E	F
Player 1	A	9,10	3,5	5,4
	B	1,6	17,9	8,5
	C	0,5	2,6	6,13

Where the first number in each square is the payoff of player 1 and the second number is player 2's payoff.

(a) Find a pure strategy a_2 for player 2 and prove that it is not rationalizable.

(b) Find a different pure strategy a'_1 of player 1 and prove that it *is* rationalizable.

(c) Find a correlated equilibrium of the game where player 1 achieves an expected payoff of 14. As randomizing devices, you have three publicly observable, fair coins: a nickel, a dime and a quarter. You may not use any other randomizing devices.

6. **[25 points] (Grading and Peer Review)**

Students' grades in CPSC 523A will be determined mainly by the instructor; however, they will also depend on peer-review evaluations performed by other students. Because this course focuses on systems in which multiple self-interested agents take strategic action to maximize their rewards, it seems sensible to ask whether such peer-review grading will work. Specifically, what will happen if self-interested students are willing to strategically manipulate their peer reviews to maximize their own grades?

(a) We first introduce a formal model of the peer-review grading scenario. Let $S = \{0, \ldots, N\}$ be the set of participants in CPSC 523A: let 0 denote the instructor, and let $1, \ldots, N$ denote each of the N students in the class. Let α be the fraction of a student's final grade which is determined by the instructor. Let $g : S \times S \setminus \{0\} \mapsto [0, 1]$ be the grading function, where $g(i, j)$ denotes the grade given by participant i to student j. For all $1 \leq i \leq N$, let $g(i, i) = 0$. Student j's unadjusted final grade is:

$$f_j = \alpha g(0, j) + \sum_{i=1}^{N} \frac{1 - \alpha}{N - 1} g(i, j)$$

Argue that student j cannot affect f_j by changing $g(j, \cdot)$.

(b) **Grading on a curve:** Let μ and σ denote the mean and standard deviation of final grades. Assume that the instructor wants to curve grades so that the mean is μ' and the standard deviation is σ'. He could do this by giving student j the adjusted final grade:

$$\frac{\sigma'(f_j - \mu)}{\sigma} + \mu'$$

However, let's keep things simple in this section and assume that the professor doesn't want to change the standard deviation. He can thus assign adjusted final grades as follows:

$$f'_j = f_j + (\mu' - \mu)$$

i. Argue that j *can* affect f'_j by strategically changing $g(j, \cdot)$.

ii. How should j select values $g(j, \cdot)$ in order to maximize f'_j?

iii. Show that the strategy shown as the answer to the previous question is a *strongly dominant strategy*: i.e., each student is strictly better off following this strategy regardless of the peer-review strategies employed by other students.

(c) **Incentive-compatible grading:** Define

$$f_{i\sim j} = \begin{cases} \alpha g(0,i) + \left(\sum_{k=1}^{N} \frac{1-\alpha}{N-2} g(k,i)\right) - \frac{1-\alpha}{N-2} g(j,i) & i \neq j; \\ \alpha g(0,i) + \left(\sum_{k=1}^{N} \frac{1-\alpha}{N-1} g(k,i)\right) & i = j. \end{cases}.$$

Define $\mu_{\sim j}$ and $\sigma_{\sim j}$ as the mean and standard deviation of $f_{\cdot\sim j}$. To try to prevent the manipulation of peer-review grades, the instructor calculates curved grades using these values:

$$f_j^* = \frac{\sigma'(f_j - \mu_{\sim j})}{\sigma_{\sim j}} + \mu'$$

Note that in this case we're allowing the instructor to vary the standard deviation, because it doesn't make things any more complicated :-)

i. Show that student j cannot affect f_j^* by strategically changing $g(j,\cdot)$.

ii. Note that when each student j receives the grade f_j^* the mean and standard deviation of the grades are not exactly μ' and σ'. Explain why there is no way of choosing f_j^* which simultaneously satisfies the following properties:

A. the mean and standard deviation are exactly μ' and σ';

B. no student j has incentive to strategically change $g(j,\cdot)$;

C. f_j^* is strictly increasing in $g(i,j)$ for all $i \neq j$.

Academic Honesty Form

For this assignment, it is acceptable to collaborate with other students provided that you write up your solutions independently. The only reference materials that you can use are the course notes and textbook, and the reference textbooks listed on the course web page. In particular, getting help from students or course materials from previous years is not acceptable.

List any people you collaborated with:

List any non-course materials you refered to:

Signature:

Fill in this page and include it with your assignment submission.