Algorithmic Game Theory and Applications: Homework 1

This homework is due at 3:00pm, on Wednesday, February 28th. (This is a firm deadline. Please hand it in by that time to the ITO. Do not collaborate with other students on the coursework. Your solutions must be your own.) Each question counts for 25 points, for a total of 100 points.

1. Consider the following 2-player strategic game, G:

$$\begin{bmatrix} (7,4) & (7,6) & (4,4) & (4,3) \\ (9,5) & (5,3) & (4,6) & (8,4) \\ (9,4) & (5,3) & (5,8) & (5,4) \\ (6,8) & (5,9) & (4,8) & (9,8) \end{bmatrix}$$

   This is a "bimatrix", to be read as follows: Player 1 is the row player, and Player 2 is the column player. If the content of the bimatrix at row i and column j is the pair (a,b), then $u_1(i,j) = a$ and $u_2(i,j) = b$.

   Find (i.e., compute) *all* Nash equilibria for this game G.

   For any profile x that you claim is an NE of G, prove that x is indeed an NE of G.

   Argue why there are no other (pure or mixed) NEs, other than the profiles you claim are NEs.

2. Consider the 2-player zero-sum game given by the following payoff matrix for player 1 (the row player):

$$\begin{bmatrix} 4 & 5 & 2 & 3 & 8 \\ 7 & 3 & 7 & 4 & 1 \\ 2 & 4 & 6 & 5 & 2 \\ 9 & 1 & 2 & 9 & 1 \\ 6 & 5 & 4 & 4 & 2 \end{bmatrix}$$

   Set up a linear program associated with this game, and use some linear program solver to compute both the minimax value for this game, as well as a minimax profile, i.e., "optimal" (i.e., minmaximizer and maxminimizer) strategies for players 1 and 2, respectively. Specify the linear program(s) that you used to solve the game. Also, specify the *dual* of the linear program, and explain how to interpret the variables of the dual program.

(You can, for example, use the linear programming solver package `linprog` in MATLAB,available on DICE machines. To run MATLAB, type "matlab" at the shell command prompt. For guidance on using the `linprog` package, see:
`http://uk.mathworks.com/help/optim/ug/linprog.html`.)

3. Consider the simple *2-player zero-sum* game of *Rock-Paper-Scissors*, where payoff table for Player 1 (the row player) is given by:

$$\begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix}$$

   Here, the rows of player 1 (from top to bottom) correspond to Rock (R), Paper (P), and Scissors (S), respectively, and likewise, the columns of player 2 (from left to right) correspond to Rock (R), Paper (P) , and Scissors (S), respectively.

   (a) (5 points) First, a very easy question: what is the unique Nash equilibrium, or equivalently the unique minimax profile of mixed strategies for the two players, in this game? (Justify your answer.)

   (b) (20 points) Now, suppose that the two players play this game against each other over and over again, for ever, and suppose that both of them use the following method in order to update their own (mixed) strategy after each round of the game.

      i. At the beginning, in the first round, each player chooses any of the pure strategies, R, P, or S, arbitrarily, and plays that.

      ii. After each round, each player i accumulates statistics on how its opponents has played until now, meaning how many times the opponent has played R, how many times it has played P, and how many times it has played S. Suppose these numbers are r, p, and s, respectively.
      Then player i uses these statistics to "guess" its opponents *"statistical mixed strategy"* as follows. It assumes that its opponent will next play a mixed strategy σ, which plays R with probability $r/(r+p+s)$ and plays P with probability $p/(r+p+s)$, and plays S with probability $s/(r+p+s)$.
      Under the assumption that its opponent is playing the *"statistical mixed strategy"* σ, in the next round player i plays a pure strategy (R, P, or S) that is a pure *best response* to σ.

If the statistical mixed strategy σ has more than one pure best response, then you are allowed to use **any tie breaking rule you wish** in order to determine the pure strategy best response to σ played in the next round by player i.

iii. They repeat playing like this forever.

Do one of the following two things (preferrably the first):

i. Prove that, regardless how the two players start playing the game in the first round, the "statistical mixed strategies" of both players in this method of repeatedly playing the Rock-Paper-Scissors game will, in the long run, as the number of rounds goes to infinity, converge to their mixed strategies in the unique Nash equilibrium of the game.
You are allowed to show that this holds using any specific tie breaking rule that you want. Please specify the precise tie breaking rule you have used. (It turns out that it hold true for any tie breaking rule. So it would be better if you actually prove that any tie breaking rule works.)

ii. Alternatively, instead of proving that it works, you can "show" this **experimentally** by writing a simple program that plays this strategy update method for both players in repeated Rock-Paper-Scissors, and show experimentally that, for all possible start strategies of both players, the "statistical mixed strategies" of the two players *looks like* it is converging to their NE strategies. (You will need to provide your program code, as well as the experimental output which shows that convergence "looks like" it is happening.)
Note that experimentally you can only "show" that the "statistical mixed strategies" *look like* they are converging to the NE, by repeating the game some finite number of times, but you can not be sure that they do actually converge to the NE, without proving this. This is why a mathematical proof is preferable.

(c) (0 points. This is a really hard question that I want you to think about; but you do not need to submit a solution for it, unless you want to impress us, since you get zero marks for it.)

Does this same method of updating strategies always converge to *statistical mixed strategies* that yield a Nash Equilibrium for *any* finite 2-player normal form game? If so, explain why it does. If not, give an example of a 2-player finite game where it doesn't work, and argue why it doesn't work.

4. One variant of the Farkas Lemma says the following:

   Farkas Lemma A linear system of inequalities $Ax \leq b$ has a solution x if and only if there is no vector y satisfying $y \geq 0$ and $y^T A = 0$ (i.e., 0 in every coordinate) and such that $y^T b < 0$.

   Prove this Farkas Lemma with the aid of Fourier-Motzkin elimination. (*Hint:* One direction of the "if and only if" is easy. For the other direction, use induction on the number of columns of A, using the fact that Fourier-Motzkin elimination "works". Note basically that each round of Fourier-Motzkin elimination can "eliminate one variable" by pre-multiplying a given system of linear inequalities by a *non-negative* matrix.)