University of Illinois **Fall 2018**

ECE 586BH: Problem Set 4: Problems and Solutions
Extensive form (aka sequential) games

Due: Tuesday, October 30, at beginning of class
Reading: Course notes, Sections 4.1 and 4.2

1. **[Variations of pirate game]**
Five pirates, A, B, C, D, E, in a boat at sea must decide how to allocate 100 gold coins among themselves. The pirates are lettered in order of seniority, with A being the most senior. The pirate protocol is that the most senior pirate proposes an allocation of coins to all pirates, with a nonnegative integer number of coins going to each pirate. Then all pirates, including the proposer, vote on the proposal. If a majority (meaning strictly more than half the number voting) approve the proposal, then the proposal is implemented and the game ends. Otherwise, the proposer is thrown overboard and votes no more, and the remaining pirates use the same protocol again, with the most senior remaining pirate being the next proposer, and so on. Pirates base decisions on four factors. First, each pirate prefers to not be thrown overboard, even if the pirate receives no coins. Second, given a pirate is not thrown overboard, the pirate prefers to have more coins. Third, if in any round the pirate would do as well to vote either way, the pirate will vote no in order to increase the chance the proposer is thrown overboard. Finally, pirates follow the protocol, but otherwise don't trust each other and so can't enter into binding side agreements with each other. (This differs from the version on wikipedia https://en.wikipedia.org/wiki/Pirate_game because we assume that in case of a tie vote the proposer is thrown overboard.)

   (a) What is the payoff vector for a subgame perfect strategy? Is such payoff vector unique?
   Solution: The subgames have payoff vectors described as follows. If only pirate E remains, the payoff vector is $(-,-,-,-,100)$. If only pirates D and E remain, pirate D will be thrown overboard, so the payoff vector is the same as if only pirate E remains. If pirates C,D,E remain, pirate C proposes $(-,-,100,0,0)$, which is accepted by C and D, and the game ends. If pirates B,C,D,E remain, pirate B proposes $(-,98,0,1,1)$, which is accepted by B, D and E, and the game ends. So starting at the beginning of the entire game, when all pirates are present, pirate A proposes either (97, 0, 1, 2, 0) which is accepted by A,C, and D, or (97, 0, 1, 0, 2), which is accepted by A,C, and E. There are thus two possible payoff vectors for subgame perfect equilibria.

   (b) Repeat part (a), but now assume that a proposal is accepted if and only if the number of pirates voting to accept the proposal is 2/3 or more of the number of pirates voting (including the proposer).
   Solution: Answer is the same as in part (a) for up to four pirates, because the number of yes votes in those cases is at least 2/3. Starting at the beginning of the entire game, when all pirates are present, pirate A proposes (95, 0, 1, 2, 2), which is accepted by A,C,D, and E, and the game ends. This is the unique payoff vector for subgame perfect equilibria.

2. **[Market power in territory control game, continued]**
Recall the territory control game for a graph $G = (V, E)$ and $n \geq 2$ players, defined in problem

set 1 and revisited in problem set 2. As in problem set 2, suppose that players 2 and 3 are grouped together into a single player that selects two vertices. For this problem, we refer to the combination of players 2 and 3 as a single player, namely player 2. The new aspect for this problem is the order that the selections are made, making it an extensive form game. First, player 2 declares a vertex at which the player will place a restaurant. Then player 1 declares a vertex for placement of a restaurant. Then player 2 declares a vertex for placement of the other restaurant of player 2. Consider the special case of a line graph $V = \{1, 2, \ldots, m\}$. Keep in mind that one vertex can be selected multiple times.

(a) Draw the game tree for the special case $m = 2$ and describe the set of all subgame perfect equilibria in pure strategies (i.e. those strategies that can be produced by backwards induction (aka dynamic programming).

Solution:

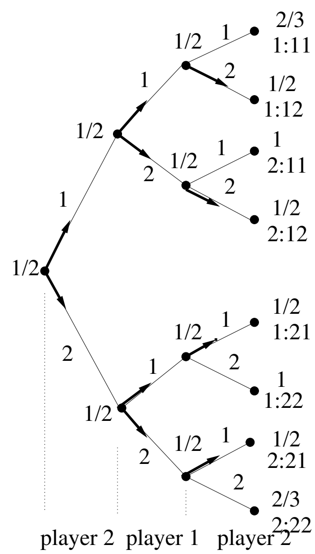

Leaf vertices are labeled with payoff to player 1 and final state ($v_1 : v_2 v_3$). The set of all subgame perfect equilibria can be described as follows. Player 2 can initially select either vertex. Then player 1 can select either vertex. Then player 2 selects the vertex player 2 did not select the first time. The payoff vector is $(1/2, 3/2)$ for these equilibria.

(b) Let $F^*_{1,m}$ denote the value of the game for player 1, which is well defined because the game has two players and the sum of payoffs is constant. Identify $\lim_{m\to\infty} F^*_{1,m}/m$, which is the fraction of customers attracted to the restaurant of player 1. (Hint: The game can be solved numerically on a computer for small m, but you can also think of the case m is large and explore backwards from the end of the game. To set notation, let v_2 be the first choice for player 2, let v_1 be the choice of player 1, and let v_3 be the second choice for player 2. Assume without loss of generality that $v_2 \geq (n+1)/2$, in which case it is easy to see that a best response v_1 for player 1 will satisfy $v_1 \leq v_2$. Then suppose $v_1 \approx xm$ and $v_2 \approx ym$ with $0 \leq x \leq y$ and $y \geq 1/2$. Don't forget the possibility $v_1 = v_2$.)

Solution: Suppose v_2 and v_1 are given. Let's identify the optimal choice of v_3 by player 2. Let $x = v_1/m$ and $y = v_2/m$ and use F to denote the payoff of player 1, which varies depending on the state of the game.

Case 1: $v_1 = v_2$ An optimal choice of v_3 for player 2 is $v_3 = v_2 - 1$, yielding $F =$

$(m - v_2)/2 \approx m\left(\frac{1-y}{2}\right)$.
Case 2: $v_1 < v_2$ At least one of the following three choices of v_3 is optimal: $v_1 - 1, v_1$, or $v_1 + 1$ leading to payoff to player 1 of approximately $(v_2 - v_1)/2$, $(v_1 + v_2)/4$ or v_1. Player 2 would select whichever of these gives the smallest payoff to player 1. Since the second of these payoffs is the average of the first and third, it can be assumed without loss of generality that the second choice is not used. So the payoff of player 1 in case 2 is $\min\{(v_2 - v_1)/2, v_1\} \approx m \min\left\{\frac{y-x}{2}, x\right\}$.
Case 3: $v_2 > v_1$ Since we have assumed $v_2 > (m + 1)/2$, such choice of v_1 would lead to smaller payoffs for player 1 so we shall not consider it further.

Next, let's consider the optimal choice of v_1 given v_2, or equivalently for large m, the choice of x for given y. Once y is selected with $1/2 \le y \le 1$, player 2 can receive reward $m\left(\frac{1-y}{2}\right)$ or $m \min\left\{\frac{y-x}{2}, x\right\}$. The reward for the second case is maximized by selecting x to make the two terms equal (by the equalizer principle). Solving $\frac{y-x}{2} = x$ for x yields $y/3$. Thus, given y, player 1 can get reward $m\left(\frac{1-y}{2}\right)$ (by selecting $v_1 = v_2$) or $y/3$ (by selecting $v_1 \approx my/3$ and will select whichever gives larger reward. So, given y, the payoff to player 1 will be $m \max\left\{\frac{1-y}{2}, \frac{y}{3}\right\}$.
Finally, let's consider the optimal choice of v_2, or equivalently, for large m, the choice of y. Player 2 will select y to minimize $\max\left\{\frac{1-y}{2}, \frac{y}{3}\right\}$, which by the equalizer principle gives $y = 3/5$.
The value of the game to player 1 thus satisfies $\frac{F^*_{1,m}}{m} \to 1/5$ as $m \to \infty$. This results from $v_2 \sim (3/5)m$, $v_1 \sim (1/5)m$, and $v_3 \sim v_1$.

3. **[Call my bluff card game]**
Consider the call my bluff card game in the class notes (Examples 4.5 and 4.8).

(a) Find the expected payoff for player 1 at Nash equilibrium. Also, in order to get some practice with definitions, identify the behavioral strategy profile σ and the belief vector μ such that the assessment (σ, μ) is the sequential equilibrium.
Solution: Examining the normal form of the game, we see there are no pure strategy Nash equilibria. Strategy Cr is weakly dominated by strategy Rr and strategy Cc is dominated by strategy Rr. We then find $(1/3)[Rr] + (2/3)[Rc], (2/3[M] + (1/3)[P])$ is the unique Nash equilibrium, with value 1/3 for player 1. Also, the unique sequential equilibrium is the assessment (σ, μ) such that $\sigma = ((\sigma(1.a), \sigma(1.b)), \sigma(2.c)) = ((1, 0), (1/3, 2/3)), (2/3, 1/3))$ and $\mu = ((\mu(1.a), \mu(1.b)), \mu(2.c)) = (((1), (1)), (3/4, 1/4))$. In particular, by Bayes rule, given information state 2.c is reached, the game is in state x with probability $(1 \times 0.5)/(1 \times 0.5 + (1/3) * 0.5) = 3/4$.

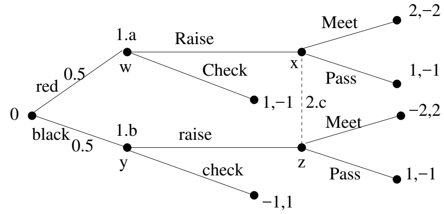

(b) Consider the following variation of the game. Suppose player 2 is given a third possible action in case player 1 raises. Namely, if player 1 raises, player 2 can (re)double the stakes. In other words, if player 1 raises (proposal to double the stakes) and player 2 doubles, then the payoff vector is (4,-4) if the card is red and (-4,4) if the card is black. Suppose if player 2 doubles, player 1 must accept it and the game ends. As before, if player 1 checks then player 2's action is irrelevant. Give both the extensive and normal form version of the game and determine the expected value of the game to player one.

Solution: The new normal form game is shown.

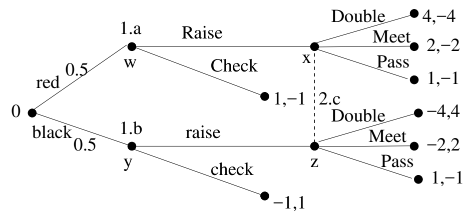

	D	M	P
Rr	$0,0$	$0,0$	$1,-1$
Rc	$1.5,-1.5$	$.5,-.5$	$0,0$
Cr	$-1.5,1.5$	$-.5,.5$	$1,-1$
Cc	$0,0$	$0,0$	$0,0$

Again strategies Cr and Cc are weakly dominated and can be eliminated, and then strategy D becomes weakly dominated by M. So by iterated elimination of weakly dominated strategies, we arrive at the same normal form game as in part (a). The Nash equilibrium found in part (a) continues to be the unique Nash equilibrium and the value to player 1 is 1/3 as before.

4. **[Entry deterrence game with random cost]**
Consider the following two variations of the entry deterrence game discussed in the notes. The node with label 0.b is controlled by nature which selects each possible outcome with probability 0.5. For version (a) player 2 knows the random outcome of nature when it needs to select an action, and for version (b) player 2 doesn't know it.

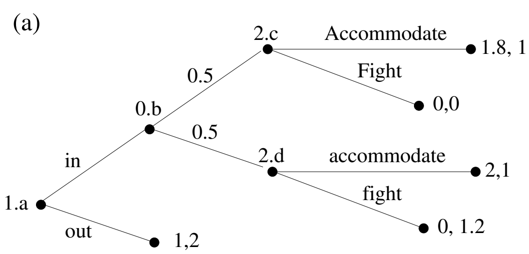

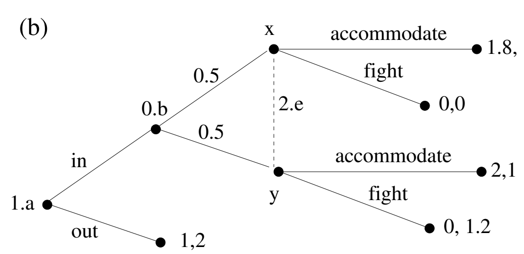

(a) Find the subgame perfect equilibria and expected payoff vector for version (a).

Solution: Player 2 selects action A at the node labeled 2.c because it is optimal for the subgame beginning there. Similarly player 2 selects action f at the node labeled 2.d. Thus, the payoff vectors at nodes labeled 2.c and 2.d are (1.8, 1) and (0, 1.2), respectively. The expected payoff vector for the game beginning at 0.b for subgame perfect strategy profiles is thus (0.9, 0.6). Hence, player 1 selects action out at state 1.a. In summary, the unique subgame perfect equilibria is (out, Af), with payoff vector (1,2).

(b) Find the sequential equilibrium and associated expected payoff vector for version (b). In order to get some practice with definitions, identify the behavioral strategy profile σ and the belief vector μ such that the assessment (σ, μ) is the sequential equilibrium.

Solution: Player 2 knows that if information state 2.e is reached, then the two nodes in it are equally likely. So the expected payoff to player 2 is 1 for action a and 0.6 for action f. Thus, the sequentially rational choice for player 2 is accommodate. That makes the payoff vector at the nature node (1.9, 1), and the rational action of player 1 is "in." Thus, the unique sequential equilibrium is the assessment (σ, μ) such that $\sigma = (\sigma(1.a), \sigma(2.e)) = (in, accommodate)$ and $\mu = (\mu(1.a), \mu(2.e)) = ((1), (0.5, 0.5))$, with payoff vector (1.9,1). (Player 2 does better in (a) where player 2 has more information.)

5. **[Trembling hand perfect strategies and perfect information games]**

   (a) Consider a perfect information sequential game such that any player controls at most one node along any instance of the game (i.e. along any path from root to a leaf of the game). Show that a trembling hand perfect behavioral strategy profile σ is subgame perfect. You may start from scratch or use a result from the notes.

   Solution: Let σ be THP. By definition, there exists a sequence $(\sigma^k)_{k\geq 1}$ of fully mixed strategies such that for any player i, $\sigma_i \in B_i(\sigma^k_{-i})$. Let x denote an arbitrary decision node of the game and let i be the player controlling x. For any $k \geq 1$, σ^k is fully mixed, so the game reaches node x with strictly positive probability under σ^k, so the strategy player i uses at node x, denoted by $\sigma(i, x)$, (it is a probability distribution on the actions at node x) must be a best response in subgame(x) when σ^k is used at all the other decision nodes in the subgame, for any $k \geq 1$. Since the best response sets are upper semicontinuous, $\sigma(i, x)$, is also a best response in subgame(x) when the limit, σ, is used at all the other decision nodes. So σ can be obtained by the backwards induction algorithm, so by Proposition 4.3 it is subgame perfect.

   ALTERNATIVELY, the assumption of the first sentence implies equivalence between the normal form and multiagent form of the game. Also, since the game has perfect information, each information set has only one node in it, so there is only one belief vector μ, where μ assigns probability one to the single node within each information set. Thus, by Proposition 4.18 of the notes, if σ is THP, then (σ, μ) is a sequential equilibrium. It means that for any node controlled by any player, if the decision probabilities of all other nodes are given, the decision probabilities at the node maximize the expected payoff of the player. Therefore the strategy profile can be obained by the backwards induction algorithms, so by Proposition 4.3 of the notes it is subgame perfect.

   (b) Give an example of a game satisfying the conditions of the first sentence in (a) showing the converse to (a) is not true: there is a subgame perfect equilibrium that is not THP.

   Solution: Consider the game shown.

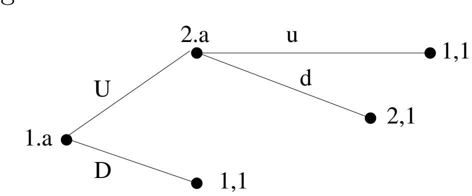

The strategy profile ((0.5, 0,5), (1,0)) is subgame perfect. However, if σ' is any profile of fully mixed strategies, then player 2 must use action d with strictly positive probability, so action U is strictly suboptimal for player 1. That is, $\sigma \notin B(\sigma')$. Therefore, σ is not THP.