

CS285 Spring 2009 – Homework 2 (due April 14, in class)

Professor David Parkes

*Points will be awarded for clarity, correctness and completeness of the answers. Typed answers will get a 5% bonus, but typing is not required. Submissions should be electronic or printed in class. Each student will be allotted 3 late days throughout the semester. Delays exceeding 3 days will cause penalties of 5% each to the overall homework grade. Submissions more than 4 days late on any one homework will not be accepted. Students may work **in pairs** and submit only one solution, but they **both** have to understand it.*

Total points: 310 – extra credit 50 points

1. (45 points) Big monkey (M) and little monkey (m) live happily in the jungle, but they are often hungry. Both monkeys eat the fruit of a particular tree, which grows high above the ground. They both have two actions: “bang on tree” (B) and “climb to obtain the fruit” (C). Being more senior, big monkey gets to act first. Banging on the tree by big monkey yields 0.6 probability for the fruit to come down, but only 0.2 probability when performed by little monkey. However, big monkey requires 10 Kcal to bang on the tree, while little monkey need spend only 4 Kcal. Climbing up on the tree is certain to deliver the fruit, but costs 60 and 50 Kcal for the two monkeys, respectively. Also, if a fruit drops on the ground from banging on the tree, big monkey eats 60% of the fruit and little monkey eats 40%, but, when obtained through climbing, the monkey that got it eats 90% of it. If both monkeys choose to bang on the tree and the fruit does not come down, big monkey gets angry and can choose to punish little monkey for being lazy, which induces a cost of 10 Kcal to the former and 100 Kcal to the latter. Each fruit gives 80 Kcal to both monkeys and there is just one fruit on the tree.
 - (a) Draw the extensive form game tree. Is this a game of perfect or imperfect information?
 - (b) Find a Nash equilibrium.

- (c) How would this equilibrium change if big monkey's punishment gave it pleasure (worth 10 Kcal), instead of a 10 Kcal deficit?
- (d) What should be the probability of a fruit falling when little monkey plays B (currently 0.2), such that little monkey is indifferent between B and C ?
- (e) Represent this game in its sequence form.
2. (40 points) Alice and Bob want to share c ounces of ice-cream. They both think that bargaining is the fairest and most efficient way to share it. In the beginning (round number 1) Alice makes an offer $x_1 \in [0, c]$ to Bob. If Bob accepts, he gets x_1 ounces and Alice gets to keep the rest. If he rejects, he can make a counter-offer of x_2 ounces to Alice in the next round (2), which Alice can accept or reject, etc. However, since the day is sunny, after every rejection the available quantity of ice-cream is halved, i.e., $x_2 \in [0, \frac{c}{2}]$. (**Note:** This is an infinitely repeated game with discount rate $\frac{1}{2}$.)
- (a) Formulate this as an extensive form game and find the unique subgame perfect Nash equilibrium.
- (b) Find another Nash equilibrium (one that is not subgame perfect). Can you represent the strategies of the two players in this Nash equilibrium as finite-state automata?
3. (30 points) Consider Alice and Bob, two producers in a Cournot duopoly market. They both manufacture the same product at zero cost and compete in quantities. (This is similar to question 3(a) of the previous homework set, with $c = 0$.) There are two price forecasting agencies in town, offering two different predictions: Agency X predicts the price will be $p = 1 - 2(q_a + q_b)$, whereas agency Y predicts $p = 1 - (q_a + q_b)$ (here q_i is the quantity produced by firm i and p is the price-per-item, hence for example Alice's profit equals $\pi_a = q_a p$). We assume that it is common knowledge that Bob believes and uses X 's prediction. However, he doesn't know which agency Alice trusts and uses. His prior is that, with probability α she also trusts agency X , while with probability $1 - \alpha$ she trusts Y . (Assume for simplicity that α is common knowledge.)
- (a) Formulate the above as a Bayesian game. Define (i) the agents' types, (ii) their common prior over types, (iii) their utility functions.
- (b) Find a Bayes-Nash equilibrium (q_a^*, q_b^*) for the above game. Find the payoffs for both agents, for the case that Alice trusts agency Y .
4. (45 points) Consider the routing network on the next page (Fig. 1).




Figure 1: The routing network

- (a) Find a Nash equilibrium of this routing game with until flow from X_1 to X_5 . (**Hint:** In a Nash equilibrium, all the different routes that receive non-zero traffic must have the same cost. You may use Matlab or Mathematica for equation solving.)
- (b) What is the price of anarchy in this routing game under the Nash equilibrium you found? (**Hint:** Find the optimal network flow.)
5. (50 points) A MAID is a directed acyclic graph with chance nodes (ovals), decision nodes (rectangles) and utility nodes (diamonds). Chance nodes contain conditional probability distributions over their domain values, given the values of their parents in the graph. Each decision node belongs to a particular agent and its domain is the set of actions available to that agent. Finally, utility nodes are deterministic functions of their parents' values. We define the following concepts for a directed acyclic graph $G = (V, E)$.

- **Blocking:** A path $\pi = (v_i, \dots, v_j)$ from node v_i to node v_j is blocked at node $v \in \pi$ by a set of nodes $Z \subset V$ if either (i) the path π does not have “converging arrows” at v (see Fig. 2a) and $v \in Z$, or (ii) the path has “converging arrows” at v (see Fig. 2b) and neither v nor any of its descendants are in Z .
- **Active path:** A path π from node v_i to node v_j is active given subset $Z \subset V$ if it's not blocked in any of its nodes by Z .
- **d-separation:** Two nodes v_i, v_j are d-separated given $Z \subset V$ if there is no active path from v_i to v_j given Z .
- **s-reachable:** A decision D_1 (belonging to player 1) is s-reachable from decision D_2 (belonging to player 2) in the graph of a MAID if, when we

add a new parent p to D_1 , then p is not d-separated from at least one of the utility nodes of player 2, given the subset Z containing node D_2 and all its parents.

- *relevance graph*: For every MAID, its relevance graph consists of a node for each decision node in the MAID, and there is an edge from D_2 to D_1 if and only if D_1 is s-reachable from D_2 in the MAID graph.




Figure 2: (a) Non-converging and (b) Converging arrows

Now consider the following game: There is a room with treasure, worth 1,000. However, the room might also contain a tiger with probability 0.4. There are two agents in this world, Alice and Bob. Neither Alice nor Bob know whether a tiger is in the room. Bob has to decide whether he's going to enter the room and claim the treasure. If he does, he gets 800 and gives the remaining 200 to Alice. If there is a tiger, however, he also incurs a cost of 2,000. Alice has the option, before Bob makes his decision, to open a hole on the room door so that Bob can see the contents of the room. Either action ('open door,' 'not open door') is costless to Alice.

- Draw the MAID for this game. Present the conditional probability tables in each of the chance and utility nodes, as well as the action domain of each of the decision nodes.
 - Draw the game's relevance graph.
 - Draw the relevance graph for the same game, where now Bob observes the content of the room through a second, already-opened hole. Which of these two versions of the game is computationally easier to solve?
 - Why are compact representations of games critical for progress in algorithmic game theory? (A few sentences are fine.)
6. (30 points) Consider a population where there is 80% of A-type species, and 20% of B-type species. The payoffs for the pairwise interaction of the two agents are given in Table 1 (see next page).

- (a) Calculate the average payoff for each type of species.
 - (b) Calculate the rate of change of their fractions in the population according to the replicator dynamic in the initial state.
 - (c) Find the three stable points of the replicator dynamic. Which of these are robust to minor perturbations? (**Hint:** Find the population fractions for which the rate of change of A-type is zero; plotting this as a figure might help.)
7. (20 points) Consider an absent-minded driver who, in order to get home, has to take the highway and get off at the second exit. Turning at the first exit leads into a bad neighborhood (payoff 0). Turning at the second exit yields the highest reward (payoff 4). If he continues beyond the second exit, he will have to go a very long way before he can turn back home (payoff 1). The driver is absent-minded and is aware of this fact. When reaching an intersection, he cannot tell whether he is at the first or the second intersection; that is, he cannot remember how many intersections he has passed.
- (a) Formulate the situation as a one player game.
 - (b) Show that the best behavioral strategy is better than the best mixed strategy.
8. (15 points) Suppose an agent plays a repeated game in which the payoff in each period is π , the agent's discount rate is ρ , and the probability that the agent dies at the end of each round is $\sigma > 0$. Prove that the present value of the game to the agent is the same as the present value of the game if the probability of death were zero and the discount rate were $\rho + \sigma$.
9. (20 points) It is observed experimentally that humans in the ultimatum game often propose more than some minimal amount, and decline positive offers.
- (a) What theory does this behavior seem to violate, i.e. for which simple formal model is this not a Nash equilibrium?
 - (b) Provide three (3) plausible explanations for this apparent deviation from the theory of Nash equilibrium. No more than a couple of sentences are required for each explanation.
10. (15 points) Describe and motivate three different "learning agendas" for multi-agent systems and provide a concrete instance of prior work that fits within each one. Which one do you find most satisfying, and why? (A couple of sentences for each is fine.)

	A	B
A	(8, 8)	(1, 2)
B	(2, 1)	(5, 5)

Table 1: The payoffs in the game of question 6

11. (50 points extra credit) In games of imperfect information, some agents have information sets consisting of more than one node in the game tree. In those information sets, agents need to formulate *beliefs* about how likely each of these (otherwise indistinguishable) nodes are. Below are some definitions:

- *Prior*: Given a strategy profile $s = (s_1, \dots, s_n)$, the prior probability of each node of the tree is computed as follows: The root of the tree has prior 1. Each of the other nodes d has a prior equal to the prior of its parent, times the probability that the agent (or nature) playing at the parent node chooses the arc leading to d . The prior thus captures the probability that node d will be reached in a game where all players follow strategy profile s . Let $\pi[d|s]$ stand for the prior probability of node d under s .
- *Posterior*: During an actual “run” of the game, some information sets will be reached, while others might not. If information set Y , consisting of nodes y_1, \dots, y_m , is reached, the probability that we are actually at node y_i , given that Y was reached under s , is called the posterior of y_i given Y and s . If $\text{pr}[Y|s] = \sum_{j=1}^m \pi[y_j|s] > 0$, then we can use Bayes’ rule to compute the posterior: $p[y_i|Y, s] = \pi[y_i|s]/\text{pr}[Y|s]$. If Y has zero probability under strategy profile s , then the posterior at $y_i \in Y$ is not well-defined.
- *Sequential equilibrium*: A strategy profile s is a sequential equilibrium if there are beliefs μ for all information sets such that: (i) whenever an information set Y has non-zero probability, μ assigns probabilities to each of Y ’s nodes that are equal to their posterior given Y and s , as computed by Bayes’ rule, (ii) whenever an information set Y has zero probability there is a sequence of fully-mixed¹ strategy profiles s^0, s^1, \dots , which converge to s at the limit ($\lim_{k \rightarrow \infty} s^k = s$), and μ assigns probabilities equal to the limit of the corresponding posteriors in this sequence ($\mu[d|s] = \lim_{k \rightarrow \infty} p^{(k)}[d|Y, s]$), and (iii) at each information set where agent j plays, s_j is a best response given μ and s_{-j} ².

Now consider the following game (Fig. 3).

¹A strategy profile is fully mixed if every arc in the game tree has non-zero probability.

²implied: in the subgame originating at that information set.




Figure 3: The game tree

- Verify that the strategy profile (GX, KN) is a Nash equilibrium of the game.
- Verify that the above strategy profile is *not* a sequential equilibrium. Which of the above conditions does it break?
- Find a sequential equilibrium. Give the sequence of fully-mixed strategies that converges to this equilibrium and present what agent 1's beliefs should be in his second information set (the one containing two nodes). (**Hint:** Assume that player 1 plays H with probability ϵ and take a sequence where $\epsilon \rightarrow 0$.)

Good luck!