

Comments Please see the attached report.

Comp34120 Exam Feedback

General Feedback to Question 1:

- 55 out of 57 students answered this question.
- The average mark for this question is 64% (i.e., 12.8 marks out of 20).
- 26 or 47 % students received a 1st class mark of 70% or better.
- 16 or 34% students received a 2nd class mark between 50%-69% (i.e., between 10 and 13 marks).
- 7 or 13% students received a 3rd class mark between 40%-49% (i.e., between 8 and 9 marks).
- 6 or 11% students received a mark of less than 40% (i.e., 7 marks or fewer).
- Of the above, 4 students received a mark less than 32% (less than 6 marks).

General speaking, the students' performance on this question was very good. I was particularly pleased at how many students could solve the game from part 4.

Detailed Feedback to Question 1:

- i. Most students got full points for this. There were two answers I accepted (most students gave the first answer):
 - a. Player 1 can force a win; player 2 can force a win; either player can force a drawn. (Some people seemed to think that draws are impossible in Kalah, but that is not correct. Draws are possible and happened in the lab.)
 - b. It is known that which player (if either) can force a win, but the optimal strategy is not known; the optimal strategy is known, but only from the standard starting position and only if the other player makes no mistakes; a winning strategy is known for one player even if the other player has made mistakes along the way.
- ii. Most students got both parts of this, solving the first part by maximising the minimum of the rows and minimising the maximum of the columns, and solving the second part using dominance. There were two types of mistakes. First, a few people solved the problem using minimax twice, which I could not give credit for both. Second, a few people forgot that it was a zero sum game so player 2 wanted to minimize. These people tried to maximise player 2's payoff and so got the wrong answer. One or two people turned it into an extensive form game and used minimax search. This does give the right answer, but it is not the same game (think of rock-paper-scissors in extensive form).
- iii. The first part was easy. Just find the value of all the nodes in the tree and this provides the value of the root and the justification. Most students are able to do this. For the second part, there were 5 branches which could be pruned. A lot of people only got some of them or got the wrong ones. Some people wrote out the full alpha-beta values for each node which is fine, but makes the problem harder than I intended. I think you should be able to do this by inspection. For example, once node 1.1 is valued at 10 and node 1.2.1 at 14, you don't need to expand branch 1.2R, because the value of 1.2 will be at least 14 whatever comes from that branch. So, the value of 1 will be 10. One mark was awarded for each correct branch.
- iv. It was very pleasing to see how many students were able to solve this game from scratch. Some people guessed the solution (1/2,1/2) and showed there was not pure strategy which

beat it. Some people actually solved the game by assuming that Robert played a with probability x and b with probability $1-x$, and found that value of x which Cassie could not beat with any pure strategy, and then invoked symmetry to get Cassie's strategy. Among those who did not get this question right, the most common guess was $(2/3, 1/3)$ or $(1/3, 2/3)$. Those who answered this usually could not or did not try to show that it was Nash. This gave less points. Those who answered with a pure strategy got no points.

General Feedback to Question 2:

- 43 out of 57 students answered this question.
- The average mark for this question is 66% (i.e., 13.2 marks out of 20).
- 23 or 53% students received a 1st class mark of 70% or better.
- 13 or 25% students received a 2nd class mark between 50%-69% (i.e., between 10 and 13 marks).
- 5 or 12% students received a 3rd class mark between 40%-49% (i.e., between 8 and 9 marks).
- 2 or 11% students received a mark of less than 40% (i.e., 7 marks or fewer).
- Of the above, 1 students received a mark less than 32% (less than 6 marks).

General speaking, the students' performance on this question was very good. I was particularly pleased at how many students could solve the 2 games in part 4.

Detailed Feedback to Question 2:

- i. Most people at received at least 3 points out of 4. I was looking for 1) zero-sum, two-player, perfect information and no chance; 2) a pure strategy equilibrium exists, or player 1 can force a win, etc. 3) knowledge of the solution of smaller games. Again, some people seemed to think that draws are not possible in Kalah, which is not true.
- ii. This question was ambiguously worded (apologies). I meant for the method to be described, but it could be interpreted as the design of the program to be described. I accepted either interpretation, but I did need to get a sense which method you used (e.g. minimax search, Monte Carlo Tree Search, etc.). In a few cases I could not tell what method was used or the student did not seem to know how their method worked. Most people gave a broad description of the method and how parts of the program fit together, and that was work about 4 marks. Some people gave much more detail and that was worth 5 or 6 marks. A few people seemed to only know about the heuristic they used, but not how it was used.
- iii. It is hard to identify a common set of errors among those who did not get this right, except some people did not really know about methods different than those they used in the lab. A few people could not describe minimax search with alpha-beta pruning. There was also some confusion about what method is a learning method. Some who described MCTS in the previous question, described another learning approach in this question.
- iv. Many people got this or got at least one of the two parts, which was, again, very pleasing to see. Some people solved it by writing out the complete game tree which was fine if they got the correct answer. Breadth-first search works better and only the winning strategy subtree

is needed. One point which I had hoped would make the second part of the problem easier was that once the winning strategy without the pie rule was found, it would be recognized that with the pie rule player 1 is susceptible to player 2 using that strategy if player 1 plays S2. Therefore, player 1 must play S1, and from there it follows that North always wins. I am not sure anyone recognised this.

A common problem seemed to be that some people forgot the rules, and in particular the rules for taking seeds. Without the ability to take seeds obviously a draw can be ensured, but I could not give credit for this answer. If I ever use a question like this again, I will write out the rules. I did not anticipate that anyone would not remember the rules of the game having worked on this game for 5 weeks.

General Feedback to Question 3:

- 50 out of 57 students answered this question.
- The average mark for this question is 63% (i.e., 12.6 marks out of 20). This is a little worse than the last years, where the average mark for this question was 67%.
- 24 or 48 % students received a 1st class mark of 70% or better (i.e., 14 marks or more), in comparing with 22 out of 42 students or 53.6% in the last year.
- 14 or 28% students received a 2nd class mark between 50%-69% (i.e., between 10 and 13 marks), in comparing with 12 out of 42 students or 29% in the last year.
- 4 or 8% students received a 3rd class mark between 40%-49% (i.e., between 8 and 9 marks), in comparing with 3 out of 42 students or 7% in the last year.
- 7 or 14% students received a mark of less than 40% (i.e., 7 marks or fewer), in comparing with 4 out of 42 students or 9.5% in the last year.
- In particular, 2 students did not write even a word in their answer sheets and so got zero mark. This is the main reason behind a lower average mark and a little higher fail rate this year in comparing the last year.

General speaking, the students' performance in this question is good from the teaching outcome point of view and largely similar to the result of the last year. It is disappointed that two students got zero mark due to no effort being made (empty answer sheets).

Same as the previous year, the main reason that almost all students selected this question is that the students have done a project related to this exam question. Therefore they are more familiar to the topic and better prepared.

Detailed Feedback to Question 3:

- Question a). Most students were able to answer this question, but the accuracy and completeness of the answers are different. Some students were able to give a text description of the answer, but could not give the mathematical formulas which are required. Further most students were able to give the correct answer to a-i), but only about half of students gave the correct and complete answers to a-ii). For a-iii), most students are able to list one or two main

advantages of the recursive least square approach but only about one-third students got the full mark by giving three or more advantages. In particular, the common missing or mistake is that the advantage in the computation effectiveness is not mentioned.

- Question b). Except to a few students, most students know how to solve the given problem and answer this question well. The most common mistake is the incorrect calculation. The second common mistake is that, for question b-ii), no justification was given why checking the values of the payoff function at the boundaries 0 and 1 is enough to find the best strategy. On the other hand, the common mistake of the previous years, without checking the boundary strategies when finding the best global strategies, has largely disappeared.

General Feedback to Question 4:

- 19 out of 57 students answered this question and therefore less than half of the students selected this question.
- The average mark for this question was 59% (i.e., 11.8 marks out of 20). This is at the similar level as Q3 and much better than the last year, where the average mark for this question was 53%.
- 6 or 32 % students received a 1st class mark of 70% or better (i.e., 14 marks or more), in comparing with 5 out of 22 students or 22% in the last year.
- 8 or 42% students received a 2nd class mark between 50%-69% (i.e., between 10 and 13 marks), in comparing with 4 out of 22 students or 18% in the last year.
- 3 or 16% students received a 3rd class mark between 40%-49% (i.e., between 8 and 9 marks), in comparing with 4 out of 22 students or 18% in the last year.
- 2 or 11% students received a mark of less than 40% (i.e., 7 marks or fewer), in comparing with 7 out of 22 students or 32% in the last year.

General speaking, the students' performance in this question is good and is much better than the last year from the teaching outcome point of view. Further the similar level of the average mark between Q3 and Q4 shows that these two exam questions are in the similar level of difficulty, which was the goal set up in the exam feedback last year.

Detailed Feedback to Question 4:

- Question a). Most students answered this question well. The common mistake and incompleteness is the fail to mention that 1) the players or agents in the games have self-interest (a key feature for the games considered) and have the private information (a key feature for the mechanism design). 2) The design task or problem is to design the rules of games.
- Question b). There are 4 sub-questions in Question b). Most students were able to answer i), ii) and iv) correctly and completely. However, about half of the students failed to answer sub-question iii) correctly. The correct answer is that the second-price auction does not always generate more or at least the same revenue as the first-price auction. Although the revenue equivalence theorem proves that the second-price and first-price auctions have the same

expected revenue, it is just averagely speaking. In the other words, the first-price auction can generate more revenue than the second-price in some case and vice versa in some other cases.

- Question c). There are 3 sub-questions in Question c). For questions i) and iii), two third students answer correctly. For question ii), less than half of the students answered it correctly or completely. One possible answer to explain why is to give an example to illustrate the bidding with the true valuation is not the best strategy. Another possible answer is to use Myersons Lemma which proves that there is one and only one pricing rule ensuring DSIC. As the unique pricing rule given in Myersons Lemma is different from the extended 2nd pricing rule given in the question, this implies that the extended 2nd pricing rule cannot ensure DSIC.