

15-381 Spring 2007

Assignment 4: Auctions, Probability, and Uncertainty

Questions to Gil (egjones+@cs.cmu.edu)

Spring 2007
Out: March 20
Due: April 3, 1:30pm Tuesday

The written portion of this assignment must be turned in at the beginning of class at 1:30pm on February 20th. Type or write neatly; illegible submissions will not receive credit. Write your name and andrew id clearly at the top of your assignment. If you do not write your andrew id on your assignment, you will lose 5 points.

The code portion of this assignment must be submitted electronically by 1:30pm on April 3rd. To submit your code, please copy all of the necessary files to the following directory:

`/afs/andrew.cmu.edu/course/15/381/hw2_submit_directory/yourandrewid`

replacing yourandrewid with your Andrew ID.

Late Policy. Both your written work and code are due at 1:30pm on 4/3. Submitting your work late will affect its score as follows:

- If you submit it after 1:30pm on 4/3 but before 1:30pm on 4/4, it will receive 90% of its score.
- If you submit it after 1:30pm on 4/4 but before 1:30pm on 4/5, it will receive 50% of its score.
- If you submit it after 1:30pm on 4/5, it will receive no score.

Collaboration Policy.

You are to complete this assignment individually. However, you are encouraged to discuss the general algorithms and ideas in the class in order to help each other answer homework questions. You are also welcome to give each other examples that are not on the assignment in order to demonstrate how to solve problems. But we require you to:

- not explicitly tell each other the answers
- not to copy answers
- not to allow your answers to be copied

In those cases where you work with one or more other people on the general discussion of the assignment and surrounding topics, we ask that you specifically record on the assignment the names of the people you were in discussion with (or “none” if you did not talk with anyone else). This is worth five points: for each problem, your solution should either contain the names of people you talked to about it, or “none.” If you do not give references for each problem, you will lose five points. This will help resolve the situation where a mistake in general discussion led to a replicated weird error among multiple solutions. This policy has been established in order to be fair to everyone in the class. We have a grading policy of watching for cheating and we will follow up if it is detected.

No corporate sponsorship was involved with the writing this homework, despite evidence to the contrary.

Problem 1 (Auctioning a Wii)

Much of this assignment will examine various ways that you go about acquiring a scarce commodity - a Nintendo Wii. In this question you are asked to think about auctions in ways slightly different than what was covered in class - considering things like truthfulness, revenue generation, and susceptibility to collusion. Most of the things you've learned about game theory are true when evaluating auctions, including concepts like Nash equilibria.

Question 1.1

Your first attempt to acquire a Wii involves on-campus auctions held by student entrepreneurs. One auctioneer is holding an English auction: he sets a low price and instructs the participants to stand until they are no longer willing to pay the price. He then slowly raises the price until only one person remains standing, and that person pays the amount at which the last person sat down. The other auctioneer auctions her Wii in a Dutch auction. She sets a high price and instructs the bidders to stand when they are willing to pay that price for the Wii. She then slowly lowers the price. The first person to stand gets the Wii at that price. In the following problems assume that your valuation for the Wii is v .

A dominant strategy for a bidder in an auction is very similar to the dominant strategy concept from game theory: does an agent maximize its payoff (the agent's value for an item less the amount paid for that item) by following a particular strategy no matter what every other agents' strategy is?

1. (3 points) Is it a dominant strategy for the English auction to represent your value truthfully - to stand up until the price of v is reached, and then to sit down? Why or why not?
2. (3 points) Is it a dominant strategy for the Dutch auction to represent your value truthfully - to remain seated until the price of v is reached and then to immediately stand up? Why or why not?
3. (3 points) Suppose that there are people with a set of valuations v_1, v_2, \dots, v_n where $v_1 \geq v_2 \geq \dots \geq v_n$ for the Wiis in each auction, and that the participants in the auction intend to represent their valuations truthfully (as described above). If these people participate in each of the two auctions independently, will one Wii auctioneer make more money than the other?
4. (6 points) Suppose you know that your valuation v is the highest among all auction participants in each of the auctions, that your friend's valuation v_2 is the next highest, and that the third highest valuation, v_3 , is strictly less than either of yours. Furthermore, you know that all bidders with valuations v_3 and below are intending to represent their valuations truthfully. For each of the auctions describe an arrangement you can make with your friend that will increase your payoffs over truthfully representing your bids - for your arrangements show how each of you and the friend will make more by the arrangement than either would make by truthfully revealing your valuations. Also indicate for the given auction whether under what circumstances, if any, given that you both agreed to cheat that there would be an inducement to not perform the agreed upon action - essentially, for each of the auctions does double-crossing pay?

Question 1.2 (10 points)

Your second attempt to acquire a Wii comes in the form of a different kind of auction. One of your friends has tired of his Wii and has offered to auction it to you and one other person by the following method - each of you submits a sealed bid to him, the auctioneer. You each value the object at v , and if your bids are the same you get dual-custody of the Wii, which is worth $\frac{v}{2}$ to each of you. The higher bid wins the Wii, and both of you pay the amount of your bids to him no matter which one of you won. Is there a pure strategy Nash equilibrium in this game? If so, describe it. If not, tell why not.

Problem 2 - Waiting for a Wii

Imagine that you are no longer pursuing the auction route for acquiring a Wii but are instead going straight to the source - the Waterfront Target early on Sunday morning. You still have a valuation v for the Wii - representing your desire for the Wii monetarily. Once you get to Target you know that you can buy a single Wii for the advertised price c , which we assume is strictly less than v (you are always willing to pay more for the Wii than it actually costs). You also know that to get a Wii at that price you must be one of the first m people in line when the store opens at 8 AM, where m is the number of Wiis that Target has in stock. Thus it pays to arrive early, but you also recognize that you don't want to wait too long in the cold on Sunday morning. The desire to wait as little as possible is represented by an opportunity cost o for each minute before 8 AM that you have to wait in line to get the Wii - you could be spending that time sleeping, and depending on what your Saturday night looked like that opportunity cost can be pretty high. Finally, there's a cost k that represents getting to the store, which combines the opportunity cost of the time it takes to get to the Waterfront and the gas and munchies you had to bribe your friend with to drive you. Your goal is to minimize your waiting time and the associated cost while still obtaining a Wii.

Question 2.1 (10 points)

Suppose that you know that m is 20 for a given Sunday, and that there are 100 people that want a Wii. Suppose furthermore that all participants know the values of all other participants, with $v_1 > v_2 > \dots > v_{100} > c$. Also suppose that $k = 0$ (no cost to get to the store) and that $o = 1$ for all participants. Finally, assume that if many people arrive in line simultaneously that the people with the highest valuation get the first spots in line. What is the Nash Equilibrium for this situation?

Notes: A Nash equilibrium for this case follows the same model we've learned about in game theory: given that all other participants are going to arrive at a particular time, when does a single participant want to arrive? At an equilibrium all participants will be satisfied with their strategy given the strategy of all other individuals. Note: you can state your solution and justify it informally.

Question 2.2 (25 points + 5 bonus)

In question 2.1 you found an equilibrium for a simple version of the problem of when to arrive at Target to maximize your payoff. The simplification came from the fact that all valuations were common knowledge. This is not generally the case, the this question examines relaxing that constraint - when should you get to the store given that you don't know anyone's value or the strategy they are using?

This section involves coding. Download the files at

<http://www.cs.cmu.edu/afs/cs.cmu.edu/academic/class/15381-s07/www/hw4/hw4.tar>

These files simulate the line in front of Target on a series of Sunday mornings. You will be changing a single function in a C++, and must write your code in such a way that it will compile in this language. Please come and talk to me (Gil) if this is going to be a significant problem - all that should be required is C style coding - just filling out a single function.

For this problem assume that every Sunday morning each agent that wants a Wii is told both its valuation v for the Wii for that week - the price that it thinks it can get on Ebay for it, say - as well as that week's opportunity cost o . The agents are also told the amount in their bank accounts, and the sorted arrival times of all agents the previous week. Each agent must then use this information to determine if they will go to Target to wait for a Wii, and if so, how many minutes before 8 AM they will arrive in the line.

In this exercise you will write a function that takes your valuation for the week, your opportunity cost for the week, your bank account balance, and the arrival times of all other agents the previous week and returns the number of minutes t before 8 AM that you wish to get in line at the Target. If you choose to go to Target, get in line, and are one of the first m people in line for the week you will buy a Wii and your account will be credited with your profit P as determined by:

$$P = v - c - (t * o) - k$$

Otherwise you will have waited in line until 8 AM in vain and your account will receive a negative profit of

$$P = -(o * t) - k$$

If you choose to stay home your profit is 0.

Your task is to replace the function in `better_ArrivalTime.h` with your own code. See the `readme.txt` file included with the code for more information on the particulars of the coding assignment.

Question 2.2.1

First implement a strategy that does well against the not-so-smart random players in `random_ArrivalTime.h`. To implement this function you should only need to reason about your value and opportunity cost (and about your bank account balance a little bit as described in the `readme.txt` file).

Comment your code and turn in a brief description of what your function does with the write-up for this question (described below).

Question 2.2.2

For this problem you will implement a strategy to compute arrival times given that your peers will also be implementing strategies of their own. After you've turned in your code we will compile all of your functions together into the `wiiLine` program and run them for 5000 cycles. You must create a function that will maximize your profit given that your classmates will all trying to be doing the same. The vector of `arrivalTimes` (described in the `readme.txt` file) might be useful for this, as it gives information about the trends of arrival times of your peers.

There is no right answer for this problem - we just want you to think about strategies for a competitive real world situation where other agents have similar objectives to your own.

You will get credit for this question by implementing a function and describing your design decisions in a write-up to be submitted with the written portion of your homework.

You can submit different functions for 2.2.1 and 2.2.2 - if you want to do this turn in `better_ArrivalTime.h` with two versions of the function, with the version for 2.2.2 commented out.

Point values and what to turn in:

- (5 points) You turn in your commented code - in your directory please include all code necessary to compile `wiiLine` (almost all of which I've given you), and a compiled `wiiLine` executable which runs your code versus random participants (more info in the `readme.txt`).
- (5 points) Your code performs reasonably against a number of random opponents - we will run this in your submission directory.
- (5 points) Your code performs reasonably against other people's code - we will be extracting your function and running it against other class members'.
- (5 bonus points) Your code performs exceptionally against other peoples' - you finish in the top of the class after a large number of cycles (5000).
- (10 points) You write up and turn it in with the written of your homework (not a README) an analysis of how you made the design decisions you made in your arrival time for both questions 2.2.1 and 2.2.2.

Problem 3 - Probability (8 points)

Your Doritos are gone, and you have two apartment mates as suspects - Marc and Steve.

You know the following things:

1. In previous chip-swiping incidents Marc was implicated 85% of the time, and Steve only 15% of the time.
2. Your across-the-hall neighbor believes she saw Marc eating your Doritos last night, but her eyesight is notoriously poor. You estimate that her rate of Marc/Steve differentiation is only 80% (i.e. 80% of the time she thinks she sees Marc it's actually Marc, and 20% of the time it's actually Steve. And if she thinks she sees Steve she's correct 80% of the time).

Question 3.1 (4 points)

Given the above data what is the probability that Marc purloined the Doritos? Show your work.

Question 3.2 (4 points)

What would the probability be that Marc stole the chips if your neighbor thought she saw Steve? Show your work.

Problem 4 - Bayes Nets (22 points)

Your loyal dog Clyde has been howling for the last three hours and you want to decide whether or not to take him to the vet or just to put in ear plugs and go back to sleep. You know that Clyde often howls when there's a full moon, when he's genuinely sick, or occasionally when a particular neighborhood dog starts howling. That neighborhood dog sometimes howls at the full moon and sometimes howls when her owner isn't home, but is not affected by Clyde's howls. If Clyde's really sick he probably won't have eaten very much and should have a bunch of food left in his bowl - but he sometimes just isn't very hungry despite not being sick.

Question 4.1 (10 points)

Create a Bayesian network for the scenario described above - use single letter names for each boolean variable (explaining what they mean of course). You can make up the exact numbers in the conditional probability tables (CPTs), but both the CPT values and the causal topology of the network should be reasonable. Briefly explain why you are setting up the topology as you are.

Question 4.2 (4 points)

Identify two conditional independencies in your Bayesian network - express them using the notation $A \perp B|C$. Explain why it makes sense that these variables are conditionally independent given the structure of the problem.

Question 4.3 (4 points each)

For each of the following, write the expression you would need to compute the probabilities using your Bayes Net from 4.1 for inference. You don't need to compute the actual probabilities using your CPT values - just give the expression. Use the notation $p(A|B)$ to indicate the probability that A is true given that B is true - use $\neg A$ to indicate when A is false. Be sure to exploit the structure in your graphs to make the computation easier.

1. What is the probability that Clyde is sick given that there's no full moon, the neighbor's dog is howling, and that his food bowl is full?
2. What is the probability that Clyde is sick given that there's a full moon and you know your neighbor is away, but you have no other information?