Homework Assignment 4

CSC 200Y: Social and Economic Networks

Due: March 30, 2016

Be sure to include your name, student number and tutorial room with your assignment. If your handwriting is *possibly* illegible, be sure to hand in your assignment in some typed form. Make sure that your assignment is *stapled* together.

1. Consider the editor of a restaurant blog who wants to write an article about the more popular of two new restaurants, A and B, in some resort town. She has a group of 21 novice reporters, conveniently named $r_1, r_2, \ldots, r_{21}$, that she plans to send out to help determine which restaurant is more popular. She, and all of her reporters, know that $3/4$ of the townspeople prefer one of the restaurants and $1/4$ prefer the other, but they do not know *which* of A or B is more popular. Their prior beliefs give a 50% chance to each of A and B being the more popular. She sends her reporters to the town, and each reporter asks one random towns-person which of the two restaurants they prefer: each person will answer truthfully, A or B, with their preferred restaurant. After this single interview, each reporter votes *in sequence* by sending a text message A or B to the entire group indicating which he believes to be most popular: if his belief gives one restaurant probability greater than $1/2$ of being most popular, then this is the message he sends. In case of a tie (i.e., both are equally likely to be most popular according to his beliefs), he texts the restaurant reported as most popular by the person he interviews. Each reporter sees the sequence of all votes of anyone who voted before him.

   Because of their large fingers, each of the reporters has a 20% probability of a *reporting error*, that is, of sending an incorrect message (e.g., they will mistakenly text A even if they believe B is more popular). Suppose the reporters vote in order $r_1, r_2, \ldots, r_{21}$.

   For clarity, please use the following random variables to describe your answers:

   - MA: this means the majority of townspeople prefer restaurant A.
   - MB: this means the majority of townspeople prefer restaurant B.
   - I_j: this refers to the response of the jth reporter's interviewee: we write $I_j = A$ if the jth towns-person interviewed says she prefers A; and we write $I_j = B$ if she says she prefers B.
   - T_j: this refers to the *attempted* text message of the jth reporter: we write $T_j = A$ if the jth reporter believes MA with probability greater than 0.5 (i.e., if he *tries* to text A); and we write $T_j = B$ if the jth reporter believes MB with probability greater than 0.5 (i.e., if he *tries* to text B). Tied beliefs (i.e., $\Pr(MA) = \Pr(MB) = 0.5$) are are handled as discussed above.
   - R_j: this refers to the *actual* report/text message of reporter j: we write $R_j = A$ if the jth actually texts (or reports) A, and $R_j = B$ if he actually texts (or reports) B.

   Assume that restaurant A is actually most popular (remember, none of the reporters know this).

The jth reporter r_j's *prior beliefs* refer to his assessment of the probability of MA and MB *before* his interview, namely, $\Pr(MA|R_1, \ldots, R_{j-1})$ and $\Pr(MB|R_1, \ldots, R_{j-1})$ His *posterior beliefs* refer to his assessment *after* his interview, namely, $\Pr(MA|I_j, R_1, \ldots, R_{j-1})$ and $\Pr(MB|I_j, R_1, \ldots, R_{j-1})$.

(a) The prior beliefs of reporter r_1 (i.e., before his interview and before any voting) are simply:

$$Pr(MA) = Pr(MB) = 0.5.$$

What is the probability that $I_1 = A$? that $I_1 = B$?

What are r_1's beliefs *after* he interviews someone who declares B to be her favorite (i.e., after observing $I_1 = B$)? What will r_1 attempt to report after the interview response $I_1 = B$, and with what probability will the actual report R_1 be A or B? In other words, what is $\Pr(R_1 = A|I_1 = B)$ and what is $\Pr(R_1 = B|I_1 = B)$. Explain your answer.

(b) Suppose $R_1 = B$. What are reporter r_2's *prior* beliefs after receiving this message? (Don't ignore the possibility of text messaging mistakes.) Now suppose that r_2's interview results in $I_2 = B$. What are r_2's *posterior* beliefs after this interview, i.e., what are $\Pr(MA|I_2 = B, R_1 = B)$ and $\Pr(MB|I_2 = B, R_1 = B)$? Explain your answer.

(c) Suppose that the first two reports are $R_1 = B$ and $R_2 = B$. We won't derive this, but we simply assert that reporters r_3's prior beliefs after these reports (but before his interview) are approximately:

$$\Pr(MA|R_1 = B, R_2 = B) = 0.2248 \qquad \text{and} \qquad \Pr(MB|R_1 = B, R_2 = B) = 0.7752.$$

Using this fact (which you do NOT have to derive), provide a precise quantitative argument regarding r_3's *posterior* beliefs to show that an information cascade has formed after these two reports, i.e., that the outcome of the third interview I_3 will not influence the attempted text T_3 or the actual report R_3. Give a brief *qualitative* argument why the interviews of the remaining reporters, $r_4, \ldots, r_{21}$, will have no impact on their reports either.

(d) Using the above facts, show that the probability of an incorrect cascade forming is *at least* 0.1225. Hint: reason about the probability of the first two reports being "misleading."

(e) Learning about the possibility of such incorrect cascades, the editor has decided to change the process. Instead of reporting in sequence, she has the 21 reporters all send her a personal text message *simultaneously*. Each reporter will send a text as above using their personal beliefs, but these are now based only on their *single* interview. The reports are still noisy: specifically, you may assume that a reporter r_j will *attempt* to text the same restaurant that his interviewee suggested because he has no access to any other reports, i.e., $T_j = I_j$ for all reporters r_j. But there is still a 20% chance that a reporter reports the wrong value, i.e., that $T_j \neq R_j$.

The editor will use a *majority vote* to select the best restaurant—the restaurant for which she receives the greatest number of texts will be selected. Does this increase or decrease the probability that the editor makes the correct choice of restaurant to write about relative to the sequential reporting model? Provide a justification for your response.

Hint: To determine the probability of selecting the correct restaurant in this new simultaneous model, think of it as follows: Each reporter j has a certain chance p of reporting the correct majority restaurant, i.e., a probability p that his report R_j corresponds to the true majority restaurant. Each reporter has the same chance p of a correct report, and each report is independent, like coin flips. You should first compute this probability p (you may have done this as part of an earlier question). Then the probability that the editor decides on the correct restaurant is simply the probability that 11 or more reporters report the correct restaurant. This latter probability is a *binomial experiment* where, given a success probability p, you need to compute the odds of 11 or more successes in 21 trials. You can derive it using a formula, but feel free to use one of many online binomial calculators, e.g.,

`http://stattrek.com/online-calculator/binomial.aspx`.

2. Recall the discussion of self-fulfilling expectations equilibria in a market with positive direct effects in Chapter 17. We considered an individual $x \in [0,1]$ with a reservation price $r(x)$ for some product; and the direct benefit factor $f(z)$ indicates the relative benefit derived if a fraction z of the population also uses the product. The utility x receives from the product, if fraction z also uses the product, is given by the function $u(x,z) = r(x)f(z)$. Hence x will buy the product if and only if the price is at most $u(x,z)$. Suppose $r(x) = 1 - \frac{3}{4}x$ and $f(z) = z$, so overall utility is $u(x,z) = z(1 - \frac{3}{4}x)$.

   Now consider a product (like a newspaper) that is purchased every day and where consumers can make decisions each day that are influenced by the number of people who bought the paper the previous day. *Please justify your answer to each of the following questions.* You may find it helpful to sketch a graph of the function $u(z,z)$ on the interval $[0,1]$.

   (a) What is the maximum number of self-fulfilling expectations equilibrium points $z \in [0,1]$ that can be obtained for any positive price $p^* > 0$.

   (b) Is $z = 1$ a self-fulfilling expectations equilibrium for some price p^*?

   (c) Suppose that some producer (e.g., newspaper publisher) is convinced she can attain an initial market share (say, on the first day of sales) of at least half the consumers. In other words, she believes she can reach an initial market of $z = \frac{1}{2}$. What is the maximum price p^* she can charge so that the *tipping point* does not exceed $z = \frac{1}{2}$?

   (d) Assume the producer adopts the maximum price p^* that sustains the tipping point $z = \frac{1}{2}$, as you computed in part (c). Suppose that initial market z_0 turns out to be $z_0 = \frac{3}{8}$, less than predicted. Given our assumptions regarding who will buy the product the next day based on market share, compute the fraction of people z_1 who will purchase the product the *next* day.

   (e) Suppose somehow (and we do not care how this has come to happen) that on a certain day $t > 1$, the market share is $z_t = \frac{5}{8}$. Given our assumptions regarding who will buy the product the next day based on market share, compute the fraction of people z_{t+1} who will purchase the product the *next* day.

3. Suppose the Province wants to encourage more competition in sales of beer and now will allow certain supermarkets to sell a limited number of brands : Molson, Keith, Stella, Corona, and Heineken. The store's marketing policy is to sell only one brand to each customer. The store manager has decided to install a fancy digital board in the store that shows how many bottles of each brand is sold during a day. Whenever, a customer buys her beer, the digital board immediately gets updated. After installation of this board, Jack noticed that the behaviour of some customers is changed and their purchases are influenced by the volume sold for each brand.

   (a) The manager observes that the demand distribution follows a power law. He also knows that a p portion of customers select their beer brand without looking at the board and select uniformly at random but the remaining $1 - p$ portion are influenced by the board's statistics. Propose two models (i.e. what does the i^{th} customer choose) that might explain the power law distribution for sales that the manager is observing. Briefly indicate why you think each model might explain the observed power law behaviour. NOTE: This is just a "thought question" and any "plausible" answer and justification will be given full credit.

   (b) As is sometimes the case, some people see an opportunity to exploit the new system. Knowing how people make their decisions on selecting the brands, the manager wants to improve the profit margin. As different brands have various profit margins, he would like to manipulate the initialization of the digital board so as to sell more of high profit beer brands. Also, he doesn't want the manipulation to be easily detected, so he decides to put some constraints on his morning dishonest initializations: all brands have at least the value of 1 and the sum of all initial values on the board is not more than 15. According to your two models of behaviour, which morning initialization do you think he is going to use and why?

   (c) A smart customer understands some manipulation is taking place and just when the store opens the next day, she takes a snapshot of the digital board. Then, she writes a review of the store on her blog informing the reader regarding the dishonestly in reporting statistics. The manager is moved to another location and the new manager immediately stops the dishonest manipulation. Suggest an honest business plan for for the new manager which helps her still make increased profits by selling more high profit brands? (You don't need to come up with specifics of the plan just a high-level plan is enough).

4. Consider the following social network of undergraduate students at UofT. Each undergrad has to enrol in one computer science course and they are now about to enrol. Each node in the following graph corresponds to an undergrad. Undirected edges in the graph represent friendships.

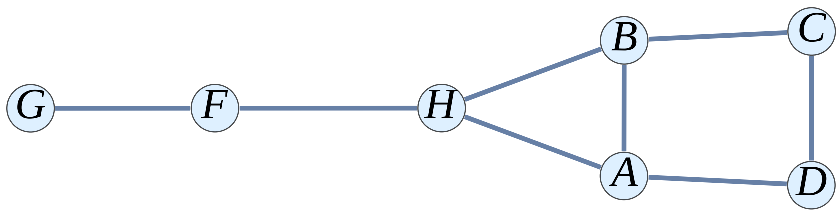

(a) Suppose that each undergrad decides to enrol in CSC200 if at least a θ fraction of his/her friends enrol. Assume that undergrads A, C, and G decided to enrol, regardless of their friends' decision (i.e., they are our "early adopters"). What is the maximum value of θ for which *everyone* will decide to enrol after the diffusion process takes place? Provide the value, along with a brief justification.

(b) Suppose that the threshold is set to $\theta = \frac{2}{3}$. Let the set function $f(S)$ be the total number of students influenced when S is chosen as the initial set of adopters. Show that $f()$ is *not* a submodular function.

(c) Even though the influence process may not be submodular, one can still execute the "greedy by marginal gain" algorithm. Suppose that the threshold is $\theta = \frac{2}{3}$ and as an instructor I can quietly promise great grades to at most two students and they will then became early adopters. Execute the algorithm to determine the two students who will be chosen. (You should break any ties lexicographically.)

(d) Suppose we run an instance of the independent cascade influence diffusion process on the above network: initially everyone not enrolled. A set of initial adopters $\{A, B, C\}$ decides to enrol. Immediately after enrolling, each student makes a single attempt to convince his friends to enrol and succeeds with probability p. What is the minimum value of p for which both students D and H enrol with probability at least $\frac{9}{16}$?

5. We imagine a community where people are arranged on a grid, and interact almost exclusively with their local (vertical and horizontal) neighbors. However, a few people "travel" and develop (symmetric) friendships with others that are more distant on the grid. The following graph represents people as nodes, where nodes 13, 25, 36, 75, and 78 are the "travellers." Friendships are represented as edges on the grid. Homophilous friendships between local neighbors are shown using thin edges, while the friendships (or weak ties) involving connections between the travellers and non-local contacts are shown using bold edges. For example, node 25 (a traveller) as developed relationships with non-local nodes 28 and 54. Notice that this forms "small world" graph based on a variant of the Watts-Strogatz model.

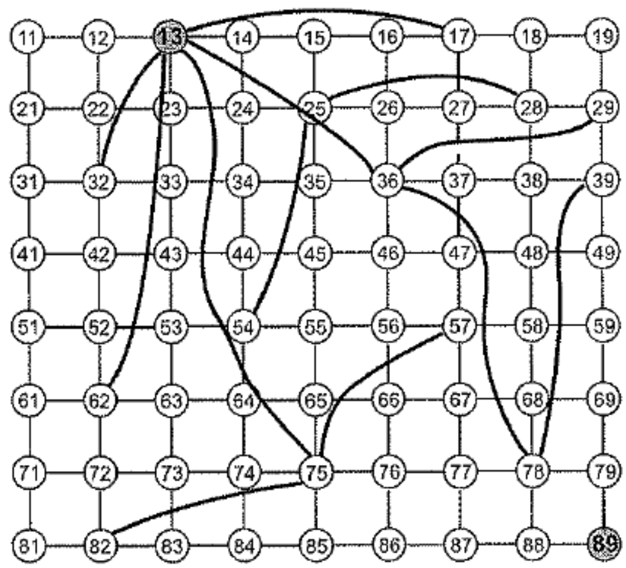

We're interested in the process of *decentralized search* on this graph involving node 13 trying to communicate a message to node 89 (the two shaded nodes). In decentralized search, if a node n is asked to forward a message so that it will reach a target node t quickly, it must forward the message to one of its friends f (who will then continue the process). Node n will forward the message to the friend f that is "closest" to target node t, where closeness is measured by grid distance (or city block distance). The grid distance is simply the length of (smallest) path between f and t using only local edges (thin edges in the picture). If there are several friends f that are equally close to the target, n can choose any such node.

(a) Node 13 is trying to get a message to node 89 using the decentralized search process. What path will the message take? (There may be more than one acceptable answer). How many hops (links) will the message need to traverse?

(b) What is a *shortest* path that the message from 13 to 89 could take (not using decentralized search)? (There may be several). How long is it?

6. Suppose that we have a set of 4 candidates $C = \{a, b, c, d\}$, and a population of 11 voters, with the following distribution of preferences over the candidates:

- 4 *voters :* $a \succ b \succ c \succ d$
- 1 *voter :* $b \succ a \succ c \succ d$
- 3 *voters :* $c \succ a \succ d \succ b$
- 3 *voters :* $d \succ c \succ a \succ b$

(a) Who is the plurality winner for this preference profile?

(b) Who is the Borda winner for this preference profile?

(c) Who is the STV (single transferable vote) winner for this proference profile? In case of a tie between two candidates (say some x and y) in any round, eliminate the candidate y who lost the most pairwise comparisons. If this still doesn't resolve the tie or if there are more than two candidates who are tied, then flip a fair coin and report back the probability of winning for each possible winner.

(d) Suppose that candidates b and d are willing to drop out of the election at a modest fee (that is, they can be bribed). Who should candidate c bribe in order to win the election under the **plurality** voting rule? Give a *minimal size* set of candidate(s). Briefly justify your answer.

(e) Now assume that we run a **Borda** election. Who should candidate c bribe now in order to win the election? Give a *minimal size* set of candidate(s). Briefly justify your answer.