Networks and Markets

Homework 4

Due: 4/18/2024, 11:59pm

Submit solutions as a PDF file to Gradescope. On Gradescope, match pages with the corresponding problem (we will make a one-point deduction per problem if pages are not matched). Show work throughout, with legible handwriting. Clearly mark your answer by putting a box around it.

As always, you are allowed to use computational tools to perform algebra/calculus, as long as you clearly show the algebra setup and explain how you calculated it. Similarly, you must disclose the use of LLM/Copilot tools that are used for any problem.

1 Efficiency of scale: Why bigger can be better (12 pts + 4pts bonus)

Consider the below circular geographic area, which is going to represent a ride-hailing marketplace where the driver must drive to the passenger to pick them up before the ride starts. Long pickup times are undesirable, because they are "unproductive" – they represent a cost (to the rider, driver, and the system) without benefitting anyone. We are going to use this system to illustrate the "efficiency of scale," why bigger systems can be more efficient if designed well.

Note: the basic principle behind this homework problem is present in many types of markets: electricity transmission, phone lines, ride-hailing marketplaces, etc, when it is more efficient for one provider to service the entire area so that you can benefit from increased density of customers. Of course, this often yields to competitive concerns, and so government responds by giving one company the sole right to operate in an area but then regulating them. These are called "regulated monopolies." For example in NYC, Con Edison is the sole distributor for gas and electricity.

For this problem, we're going to model passengers and drivers as being on the circumference of a circle with radius 1.

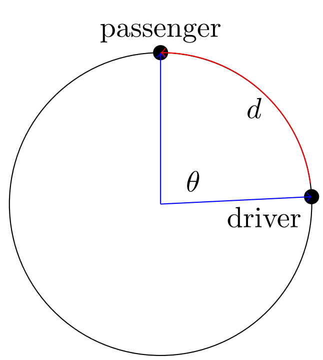

Part (a) (3 pts) As a warmup, let's suppose that there is one passenger and one driver. Each of the driver and passenger are located somewhere on the circle's circumference, each uniformly at random independently. (In other words, each of the driver and passenger can be located anywhere on the circle with equal probability). After being matched, suppose that the driver takes the shortest route on the circle to the passenger (in other words, they'll never need to travel more than half the circle to reach the passenger). In the above image, for example, the driver would travel the red portion.

What is the expected distance between the single driver and the single passenger? *Hint: for each angle θ, what is the distance d? Then, the distribution of θ is uniform, from what value to what value?*

Part (b) (3 pts) Now, still suppose there is one passenger, but now there are two drivers, each distributed uniformly at random on the circle. What is the expected distance between the passenger and the *closest* driver to them? *Hint: this might require a double integral, over the (random) positions of each of the drivers.*

Part (c) (2 pts) Now, suppose that there are 2 passengers and 2 drivers. To make things simpler, we will assume that the two passengers are on opposite ends of the circle (i.e., the angle between the two passengers is 1 radian, or 180 degrees).

First, we assume that the platform doesn't do any smart matching – it matches the drivers to the riders uniformly at random. What is the expected distance between each rider and their assigned driver?

Part (d) (4 pts) Same 2 drivers and passenger set up as in part (c).

However, now, the platform does smart matching – each driver is still matched to a different passenger, but now the system matches drivers to passengers such that the overall distance that the drivers must travel is minimized (i.e., if the distances are d_1 and d_2, the choice of matching minimizes $d_1 + d_2$).

What is the expected distance between each rider and their assigned driver?

Is this more or less than part (a)?

In one sentence, say what the results imply about two competing ride-hailing companies (e.g., Uber and Lyft), where one is bigger than the other.

Part (e) (4 pts) BONUS Same setup as part (d), but now there are N drivers and N passengers, and each is distributed uniformly at random on the circle. The platform matches drivers to riders such that the overall driving distance is minimized. Write code to simulate this setup. Plot the average distance traveled by each driver for $N = 1 \ldots 100$. *Hint: For each N, randomly place the drivers and passengers on the circle, calculate the minimum distance matching, and calculate the average distance. Do this many times for each N to smooth out randomness. You might want to use a "max weight matching" function for bipartite graphs.* Include your code and plot in your solutions upload.

2 Herding: Wisdom and foolishness of crowds (13 pts)

We will now explore how crowds can either *wise* (making the correct decisions collectively even if each person makes error-prone decisions) or *foolish* (collecting making the wrong decision).

Suppose there are two restaurants, where one restaurant is clearly better. There are N people, who are each deciding which restaurant they want to visit. Suppose that N is odd.

First, suppose each person decides for themselves which restaurant is better, and that they each make the correct decision with probability $p > 1/2$.

Part (a) (1 pts) What is the expected fraction of people who choose the better restaurant?

Part (b) (3 pts) What is the probability that a *majority* of people make the correct decision, as a function of N and p?

For $N = 3$, $p = \frac{3}{5}$, what is this probability numerically?

What is this probability as $N \to \infty$?

Now, for the problems below, suppose that people also depend on the behavior of others, for example people look online to see which restaurant is more popular. In our model, people arrive sequentially. The first person makes the correct decision with probability p.

Each person after does the following: with probability q, makes their own decision (i.e., correct with probability p). With probability $1 - q$, they instead follow the crowd: they make the same decision as the majority decision before them (or 50/50 if the majority before them is evenly split).

Part (c) (3 pts) Suppose $N = 3$. For each of $k = 1, 2, 3$, what is the probability that the kth person makes the correct decision, as a function of p and q? *Hint: Draw a tree, where each level of the tree is a person and the edges represent decisions, with some probability. For example, the first person can make 2 decisions. For each decision they make, the 2nd person can make 2 decisions, each with some probability. For each of the 4 cases (of first and second persons' decisions), the 3rd person's decisions have respective probabilities. You can refer to edges that you need to multiply and add in order to write these expressions simply.*

Part (d) (3 pts) In the same setup as Part (c), what is the probability that a majority (at least 2) people make the right decision, as a function of p and q? As before, you can refer to edges that you need to multiply and add in order to write these expressions simply.

For $N = 3$, $p = \frac{3}{5}$, $q = \frac{1}{5}$, what is this probability numerically? Compare to part (b).

Part (e) (3 pts) Same setup as part (c), with generic N. Write code to simulate this setup.

Simulate the setup 1000 times, for $p = \frac{3}{5}$ and $N = 51$, for each of $q = 0, \frac{1}{2}$, and 1.

For each value of q:

(i) What is the average number of people who make the correct decision?

(ii) What is the fraction of times that the majority (at least 26) makes the correct decision?

3 Diversity in Recommendations (15 points)

A classic challenge in recommender systems is ensuring that recommendations are sufficiently diverse. There are many reasons one may want diverse recommendations: for example, users have been shown empirically to prefer diversity; when they open Netflix, for example, they don't want to see just 10 action movies, preferring a mix of action and comedy. However, maximizing *accuracy*—the standard objective of ML-based systems—tends to produce homogeneous recommendations. This might be surprising: why would accuracy-maximizing recommendations differ from the recommendations users prefer? In this problem, we will provide a (partial) answer by analyzing a simple model of recommendations.

Suppose that Alice prefers action movies with probability 0.7, and otherwise prefers comedy movies (i.e., with probability 0.3). Now suppose that when Alice prefers genre $g \in \{\text{Action}, \text{Comedy}\}$, she likes any movie of that genre independently with probability 0.4. If the movie is not in that genre, Alice does not like it. (Therefore, Alice likes a given action movie with probability $0.7 \cdot 0.4 = 0.28$, a given comedy movie with probability $.3 \cdot .4$, but never likes both a comedy and action movie.)

A recommender is tasked with choosing n movies to recommend Alice. Define the *accuracy* of a set of recommendations to be the proportion of recommended items in the set that Alice likes. So if Alice is recommended 5 movies, and she likes 2 of them, then that set has accuracy 0.4.

Part (a) (2 points) What is the expected accuracy of a set of recommendations containing m action movies and $n - m$ comedy movies?

Part (b) (2 points) What value of m maximizes the expected accuracy?

Now consider the following observation: most people, when presented with a list of movie recommendations, often only choose one of the movies to watch. So Alice's utility is perhaps better captured by the probability that the recommendations contains *at least one* movie she likes (rather than the total proportion of recommended items she likes).

Part (c) (2 points) What is the probability that Alice likes at least one movie among a set of recommendations containing m action movies and $n - m$ comedy movies? Write your answer as a function of m and n.

Part (d) (2 points) What value of m maximizes the expression you found in part (c)? (As a function of n).

Part (e) (2 points) Let $m(n)$ denote the value you calculated in part (d). What is $\lim_{n\to\infty} m(n)/n$? Explain, very briefly, how this illustrates the importance of diversity in recommendations.

We now consider a slightly more sophisticated version of the model. In this version, when Alice prefers action movies, she likes a given action movie with probability 0.6; when Alice prefers comedy movies, she likes a given comedy movie with only probability 0.3. Again, if Alice does not prefer a genre, then she does not like any movie from that genre.

Part (f) (2 points) What is the probability that Alice likes a given action movie? A given comedy movie?

Part (g) (3 points) Let $m(n)$ be the value that maximizes the probability that Alice likes at least one movie in a set of $m(n)$ action movies and $n - m(n)$ comedy movies. What is $\lim_{n\to\infty} m(n)/n$? Briefly (in at most 2 sentences) interpret the result, noting why it might be surprising.