CS109 Midterm Exam

This is a closed calculator/computer exam. You are, however, allowed to use notes in the exam. The last page of the exam is a Standard Normal Table, in case you need it. You have 2 hours (120 minutes) to take the exam. The exam is 120 points, meant to roughly correspond to one point per minute of the exam. You may want to use the point allocation for each problem as an indicator for pacing yourself on the exam.

In the event of an incorrect answer, any explanation you provide of how you obtained your answer can potentially allow us to give you partial credit for a problem. For example, describe the distributions and parameter values you used, where appropriate. It is fine for your answers to include summations, products, factorials, exponentials, and combinations, unless the question specifically asks for a numeric quantity or closed form. Where numeric answers are required, the use of fractions is fine.

I acknowledge and accept the letter and spirit of the honor code. I pledge to write more neatly than I have in my entire life:

Signature: _____

Family Name (print): _____

Given Name (print): _____

Email (preferably your gradescope email): _____

1 Enigma Machine [20 points]

One of the very first computers was built to break the Nazi "enigma" codes in WW2. It was a hard problem because the "enigma" machine, used to make secret codes, had so many unique configurations. Let's count!

a. (4 points) The enigma machine has **three** rotors. Each rotor can be set to one of 26 different positions. How many unique configurations are there of the three rotors?

b. What's more! The machine has a plug board with single plug for each letter in the alphabet. On the plug board, wires can connect any pair of letters to produce a new configuration.

  i. (4 points) How many ways are there to place exactly **one** wire that connects two letters? A wire from 'K' to 'L' is not considered distinct from a wire from 'L' to 'K'. A wire can't connect a letter to itself.

ii. (8 points) How many ways are there to place exactly **two** wires that each connect two letters? Wires are not considered distinct. Each letter can have at most one wire connected to it, thus you couldn't have a wire connect 'K' to 'L' and another one connect 'L' to 'X'.

iii. (4 points) How many ways are there to place any number of wires? Warning: This part is worth less points per expected time than any other part of the exam. Don't get stuck :-).

2 Daycare.ai [34 points]

Providing affordable daycare would have a tremendously positive effect on society. We are building an app for daycare centers and we want to charge our customers as little as possible while paying our staff a living wage. California mandates that the ratio of babies to staff must be ≤ 4.

We have a challenge: just because a baby is **enrolled**, doesn't mean they will **show up**. At a particular location, 6 babies are enrolled. We estimate that the probability an enrolled child actually shows up on a given day is $5/6$. Assume that babies show up independent of one another.

a (4 points) What is the probability that either 5 or 6 babies show up?

b (4 points) If we charge $50 per baby that shows up, what is our expected revenue?

c (6 points) If 0 to 4 babies show up we will hire one staff member. If 5 or 6 babies show up we will hire two staff. We pay each staff member $200 a day. What are our expected staff costs? You may express you answer in terms of a, the answer to part (a).

d (8 points) What is the lowest value $k that we can charge per child in order to have an expected profit of $0? Assume that our only costs are staff. Recall that Profit = Revenue - Cost. You may express your answer in terms of a, b or c, the answers to part (a), (b) and (c) respectively.

e (12 points) Each family is unique. With our advanced analytics we were able to estimate an individual show-up probability for each of the six enrolled babies: $p_1, p_2, \ldots, p_6$ where p_i is the probability that baby i shows up. Write a new expression for the probability that 5 or 6 babies show up. You may still assume that babies show up independent of one another.

3 μ Girls [26 points]

You are on your way to buy tickets to see the new hit movie, "μ Girls" with your friends! At the movie theater, you have to make the usual decision - which line should I wait in? The lines are long so you and your friends briefly observe the rate of people served per minute in each line. How should we balance the rate at which people are served and the length of a line? **The movie is starting soon. You will make it on time if it takes less than 10 mins to buy your tickets.**

| Line	Rate (people / min)	Num People in line
Line A	$^1/_2$	4
Line B	$^1/_3$	3
Line C	$^1/_4$	2

a (8 points) It turns out that line B (which already has three people in it) is the riskiest line to wait in. If you wait in line B, what is the probability that you and your friends are on time for the movie?

b (9 points) If you choose a line uniformly at random to wait in, what is the probability that you and your friends are on time for the movie?

c (9 points) Your genius friend says, "why don't we just each take a line and if *anyone* is able to buy tickets in time, we can make the movie." Under her plan, what is the probability that you and your friends are on time for the movie? Assume that each line acts independently.

4 Name2Age [18 points]

You are working on a team that uses probability and CS to help a rockband better understand their audience. The band recently played a New Years day concert. After selling the tickets online you know the **first name** of all attendees and would like to update your belief about their ages. Names can be quite indicative of age:
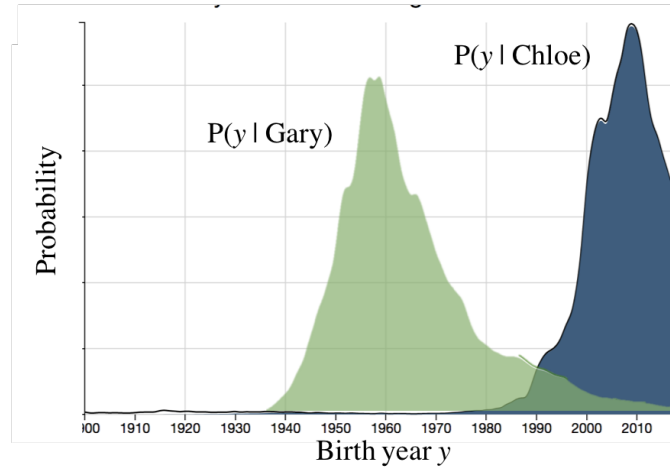
You have access to a digitalized census which tells you how many US residents were born with a given name in a given year. You also had a previous belief about the age of concert attendees (and you assume all attendees are US residents). Specifically you have functions:

Function	Description
`count(year, name)`	The number of residents, born in a given year with a given name.
`prior(year)`	The prior belief than a resident was born in a given year, given they attend the concert

Express all of your answers in terms of the `count` and `prior` functions. All people are born after 1900. Let $N_1, N_2, \ldots, N_k$ be all possible names.

a (8 points) Use census data to estimate the probability that a resident in America is named "Gary" given that they were born in 1950.

b (10 points) What is the updated belief that a resident was born in 1950, given that they attended the concert **and** their name was Gary? *Make the reasonable assumption that a person's name is **conditionally independent** of whether or not they attend the concert, given their age.*

5   Grades are not Normal [22 points]

Sometimes you just feel like squashing normals:

Logit Normal

The logit normal is the continuous distribution that results from applying a special "squashing" function to a Normally distributed random variable. The squashing function maps all values the normal could take on onto the range 0 to 1. If $X \sim \text{LogitNormal}(\mu, \sigma^2)$ it has:

PDF:
$$f_X(x) = \begin{cases} \frac{1}{\sigma(\sqrt{2\pi})x(1-x)} e^{-\frac{(\text{logit}(x)-\mu)^2}{2\sigma^2}} & \text{if } 0 < x < 1 \\ 0 & \text{otherwise} \end{cases}$$

CDF:
$$F_X(x) = \Phi\left(\frac{\text{logit}(x) - \mu}{\sigma}\right)$$

Where:
$$\text{logit}(x) = \log\left(\frac{x}{1-x}\right)$$

A new theory shows that the Logit Normal better fits exam score distributions than the traditionally used Normal. Let's test it out! We have some set of exam scores for a test with min possible score 0 and max possible score 1, and we are trying to decide between two hypotheses:

H_1: our grade scores are distributed according to $X \sim \text{Normal}(\mu = 0.7, \sigma^2 = 0.2^2)$.
H_2: our grade scores are distributed according to $X \sim \text{LogitNormal}(\mu = 1.0, \sigma^2 = 0.9^2)$.
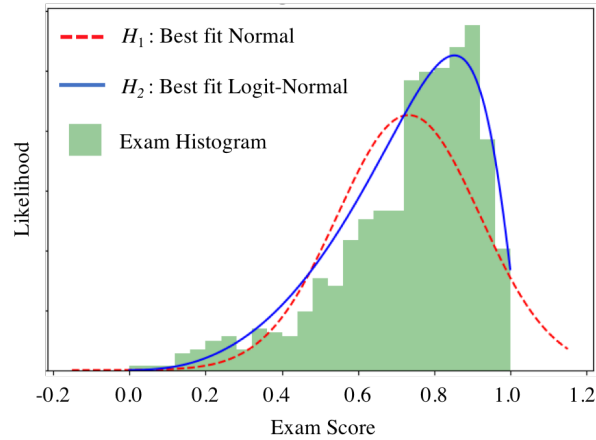
a. (5 points) Under the normal assumption, H_1, what is $P(0.9 < X < 1.0)$? Provide a numerical answer to two decimal places.

b. (5 points) Under the logit-normal assumption, H_2, what is $P(0.9 < X < 1.0)$?

c. (2 points) Under the normal assumption, H_1, what is the maximum value that X can take on?

d. (10 points) Before observing any test scores, you assume that (a) one of your two hypotheses is correct and (b) that initially, each hypothesis is equally likely to be correct, $P(H_1) = P(H_2) = 1/2$. You then observe a single test score, $X = 0.9$. What is your updated probability that the Logit-Normal hypothesis is correct?

Standard Normal Table

An entry in the table is the area under the curve to the left of z, $P(Z \leq z) = \Phi(z)$.

Z	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09
0.0	0.5000	0.5040	0.5080	0.5120	0.5160	0.5199	0.5239	0.5279	0.5319	0.5359
0.1	0.5398	0.5438	0.5478	0.5517	0.5557	0.5596	0.5636	0.5675	0.5714	0.5753
0.2	0.5793	0.5832	0.5871	0.5910	0.5948	0.5987	0.6026	0.6064	0.6103	0.6141
0.3	0.6179	0.6217	0.6255	0.6293	0.6331	0.6368	0.6406	0.6443	0.6480	0.6517
0.4	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772	0.6808	0.6844	0.6879
0.5	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123	0.7157	0.7190	0.7224
0.6	0.7257	0.7291	0.7324	0.7357	0.7389	0.7422	0.7454	0.7486	0.7517	0.7549
0.7	0.7580	0.7611	0.7642	0.7673	0.7703	0.7734	0.7764	0.7793	0.7823	0.7852
0.8	0.7881	0.7910	0.7939	0.7967	0.7995	0.8023	0.8051	0.8078	0.8106	0.8133
0.9	0.8159	0.8186	0.8212	0.8238	0.8264	0.8289	0.8315	0.8340	0.8365	0.8389
1.0	0.8413	0.8438	0.8461	0.8485	0.8508	0.8531	0.8554	0.8577	0.8599	0.8621
1.1	0.8643	0.8665	0.8686	0.8708	0.8729	0.8749	0.8770	0.8790	0.8810	0.8830
1.2	0.8849	0.8869	0.8888	0.8906	0.8925	0.8943	0.8962	0.8980	0.8997	0.9015
1.3	0.9032	0.9049	0.9066	0.9082	0.9099	0.9115	0.9131	0.9147	0.9162	0.9177
1.4	0.9192	0.9207	0.9222	0.9236	0.9251	0.9265	0.9279	0.9292	0.9306	0.9319
1.5	0.9332	0.9345	0.9357	0.9370	0.9382	0.9394	0.9406	0.9418	0.9429	0.9441
1.6	0.9452	0.9463	0.9474	0.9484	0.9495	0.9505	0.9515	0.9525	0.9535	0.9545
1.7	0.9554	0.9564	0.9573	0.9582	0.9591	0.9599	0.9608	0.9616	0.9625	0.9633
1.8	0.9641	0.9649	0.9656	0.9664	0.9671	0.9678	0.9686	0.9693	0.9699	0.9706
1.9	0.9713	0.9719	0.9726	0.9732	0.9738	0.9744	0.9750	0.9756	0.9761	0.9767
2.0	0.9772	0.9778	0.9783	0.9788	0.9793	0.9798	0.9803	0.9808	0.9812	0.9817
2.1	0.9821	0.9826	0.9830	0.9834	0.9838	0.9842	0.9846	0.9850	0.9854	0.9857
2.2	0.9861	0.9864	0.9868	0.9871	0.9875	0.9878	0.9881	0.9884	0.9887	0.9890
2.3	0.9893	0.9896	0.9898	0.9901	0.9904	0.9906	0.9909	0.9911	0.9913	0.9916
2.4	0.9918	0.9920	0.9922	0.9925	0.9927	0.9929	0.9931	0.9932	0.9934	0.9936
2.5	0.9938	0.9940	0.9941	0.9943	0.9945	0.9946	0.9948	0.9949	0.9951	0.9952
2.6	0.9953	0.9955	0.9956	0.9957	0.9959	0.9960	0.9961	0.9962	0.9963	0.9964
2.7	0.9965	0.9966	0.9967	0.9968	0.9969	0.9970	0.9971	0.9972	0.9973	0.9974
2.8	0.9974	0.9975	0.9976	0.9977	0.9977	0.9978	0.9979	0.9979	0.9980	0.9981
2.9	0.9981	0.9982	0.9982	0.9983	0.9984	0.9984	0.9985	0.9985	0.9986	0.9986
3.0	0.9987	0.9987	0.9987	0.9988	0.9988	0.9989	0.9989	0.9989	0.9990	0.9990
3.1	0.9990	0.9991	0.9991	0.9991	0.9992	0.9992	0.9992	0.9992	0.9993	0.9993
3.2	0.9993	0.9993	0.9994	0.9994	0.9994	0.9994	0.9994	0.9995	0.9995	0.9995
3.3	0.9995	0.9995	0.9995	0.9996	0.9996	0.9996	0.9996	0.9996	0.9996	0.9997
3.4	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9997	0.9998
3.5	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998	0.9998

That's all folks! We hope you had fun. The counting you did was for the real enigma machine. Folks who program the internet think deeply about queuing theory, and like movies too. The last question is an introduction to how we chose between different probabilistic models.