

Data 102, Spring 2024

Midterm 2

- You have **110 minutes** to complete this exam. There are **5 questions**, totaling **50 points**.
- You may use **two** 8.5×11 sheets of handwritten notes (front and back), and the provided reference sheet. No other notes or resources are allowed.
- You should write your solutions inside this exam sheet.
- You should write your Student ID on every sheet (in the provided blanks).
- Make sure to write clearly. We can't give you credit if we can't read your solutions.
- Even if you are unsure about your answer, it is better to write down something so we can give you partial credit.
- We have provided a blank page of scratch paper at the **end** of the exam. No work on this page will be graded.
- You may, without proof, use theorems and facts given in the discussions or lectures, **but please cite them**.
- We don't answer questions individually. If you believe something is unclear, bring your question to us and if we find your question valid we will make a note to the whole class.
- Unless otherwise stated, no work or explanations will be graded for multiple-choice questions.
- Unless otherwise stated, you must show your work for free-response questions in order to receive credit.

Last name	
First name	
Student ID (SID) number	
Berkeley email	
Name of person to your left	
Name of person to your right	
Seat number	

Honor Code [1 pt]:

As a member of the UC Berkeley community, I act with honesty, integrity, and respect for others. I am the person whose name is on the exam, and I completed this exam in accordance with the Honor Code.

Signature: _____

1 CaffeinatEd GLMs (10 points)

A professor wants to understand the effect of caffeine on student activity, so she gathers the following data for each of the 250 students in her class:

- `gpa`: the student's GPA from previous semesters
- `ed`: the number of Ed posts the student made over the semester (including follow-ups)
- `hw`: the student's average score on weekly homework assignments, each out of 10 points
- `coffee`: whether the student drank coffee every day (1) or not (0)
- `tea`: whether the student drank tea every day (1) or not (0)

The professor uses negative binomial regression to predict `ed` from the other four variables. Using the data from her 250 students, she fits both a frequentist and a Bayesian GLM, and then for each one, she makes predictions on 50 students from last semester.

She thinks that caffeine (in coffee and tea) will cause students to make more posts, so her Bayesian GLM uses a prior that assigns higher probability to larger coefficients for `coffee` and `tea`.

- (a) [4 Pts] The professor uses each model to predict the number of Ed posts for one of the students from the previous semester. She obtains credible and confidence intervals for their predicted number of posts:

- A 95% confidence interval for the predicted number of posts: (1.1, 2.4)
- A 95% credible interval for the predicted number of posts: (0.8, 1.7)

For each of the following statements, determine whether it applies to the confidence interval, the credible interval, both, or neither, by filling in the circle next to the correct answer.

- (i) According to the model, the number of posts has a 95% probability of being in the interval given the data from her current students.
- Confidence interval Credible interval Both Neither
- (ii) The interval does not account for any uncertainty in the estimated coefficients, only in the prediction itself.
- Confidence interval Credible interval Both Neither
- (iii) The interval takes into account her prior belief that drinking coffee and/or tea is associated with making more Ed posts.
- Confidence interval Credible interval Both Neither
- (iv) A 99% interval would be wider than the 95% interval given (i.e., would have a larger difference between the lower and upper bounds).
- Confidence interval Credible interval Both Neither

For the remainder of the question, she uses the same dataset to predict whether a student drinks coffee from the other variables using logistic regression. She fits two models, one with all of the other variables and one with all of them except for `tea`. She does not use an intercept or constant term for either model. Here are the coefficients from the models, along with the log-likelihood for each one:

Name	Features used	<code>gpa</code>	<code>ed</code>	<code>hw</code>	<code>tea</code>	Log-likelihood
Model 1	<code>gpa, ed, hw</code>	-0.02	1.32	0.74	N/A	-1011.4
Model 2	<code>gpa, ed, hw, tea</code>	0.03	2.51	0.83	-5.5	-1003.6

(b) [4 Pts] Which of the following are valid conclusions from these results? Select all answers that apply **by filling in the square next to each correct answer**.

- There is a negative association between drinking coffee and drinking tea among these students.
- If a student's GPA increases by 0.1, then Model 2 predicts that their probability of drinking coffee should increase by 0.003.
- If the professor wants to choose the best model according to the AIC, she should choose Model 1.

(c) [2 Pts] Consider a student with a GPA of 3.1, who made 14 Ed posts and got an average score of 6.4 on the HW assignments. Assuming that this student does not drink tea, write a mathematical expression for predicted probability that the student drinks coffee **according to Model 1**.

Your answer should contain only arithmetic expressions and the `sigmoid()` and/or `logit()` functions, and should not contain any variables.

2 Causality and Streaming (11 Points)

Data scientists at a music streaming app company are trying to evaluate whether a new recommendation system causes users to listen to more new music than the old system. They define the following variables.

- Treatment S : whether a user uses the new system ($S = 1$) or the old system ($S = 0$)
- Outcome M : the number of times the user listened to new songs over the following week
- Demographic and other data D : age, gender, favorite genres of music, level of interest in new music, etc.
- Promotion P (only used in Plan 2): whether a user receives email promotions and in-app banners for the new system ($P = 1$) or receives no promotions ($P = 0$)

They come up with three proposed plans for studying the effect of the new system:

Plan 1 Randomly assign 5000 users to see recommendations from the new system ($S = 1$) and 5000 users to see recommendations from the old system ($S = 0$).

Plan 2 Allow all users to choose between the new and old systems. Randomly assign 5000 users to receive promotions ($P = 1$) and 5000 users to receive no promotions ($P = 0$).

Plan 3 Allow all users to choose between the new and old system.

(a) [2 Pts] Draw a causal DAG for Plan 3. Your graph should have no variables other than S , M , D , and P . Note that you may not need all of these variables for a correct answer.

(b) [2 Pts] Which of the following statements are true **about Plan 1**? Select all answers that apply **by filling in the square next to each correct answer**.

- The treatment is independent of the observed outcome because users are randomly assigned to the new or old system.
- The treatment is independent of the pair of potential outcomes because users are randomly assigned to the new or old system.
- Under this plan, the simple difference in mean outcomes (SDO) estimator will be a biased estimator of the average treatment effect (ATE), because it does not adjust for potential confounding variables such as D .

- (c) [4 Pts] For this part only, assume the data scientists choose Plan 2. They find that $\text{Cov}(P, S) = 0.6$.

Based on the information given, is the promotion P a valid instrumental variable (IV)? Justify your answer by listing each assumption necessary for IVs, and explaining whether or not P satisfies the assumption.

- (d) [3 Pts] For this part only, assume the data scientists use Plan 3 with a much smaller group of users. They correctly identify a list of confounding variables from the demographic information D and estimate the propensity scores based on those variables. After correctly removing users with propensity scores outside of the range $[0.1, 0.9]$ (as you did in lab), they find that only five users remain:

Z	Y	propensity
1	20	0.8
0	8	0.2
1	3	0.3
0	16	0.8
0	12	0.4

Assuming their list of confounders is complete, use the information above to compute the inverse propensity score-weighted (IPW) estimate for the ATE.

3 Waiting Times (9 points)

A company is analyzing the time customers wait in a bank queue, T . Before collecting data, the company knows the following information about the wait times.

- **Standard Deviation:** The standard deviation in each wait time, $\sigma = \text{SD}[T] = 5$ minutes.
- **Bounds:** According to bank policy, the bank will immediately help any customer who has waited 40 minutes.

The company collected a sample of n i.i.d. wait times, T_1, T_2, \dots, T_n . Based on that data, they observe the following:

- **Sample Average:** The sample average of the observed wait times is, $\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^n T_i = 20$ minutes.

For all the parts below, you must show your calculation and/or reasoning to receive credit.

- (a) [3 Pts] For this part only, the company assumes that the sample average $\hat{\mu}_n$ equals the true mean wait time $\mu = \text{E}[T]$. Under this assumption, use Chebyshev's inequality to calculate a lower bound on the probability that a customer waits more than 10 minutes and less than 30 minutes.
- (b) [3 Pts] In general, the observed average waiting time, $\hat{\mu}_n$, will differ from the true expected waiting time, $\mu = \text{E}[T]$. The bank would like to find a two-sided tail bound on the difference $\hat{\mu}_n - \mu$ that fails with a small probability $\epsilon \in (0, 1)$. Find α as a function of ϵ and n such that

$$\Pr(|\hat{\mu}_n - \mu| \geq \alpha(\epsilon, n)) \leq \epsilon$$

via Hoeffding's inequality.

(c) [3 Pts] Using your answer to part (c), or by referencing the general form for the Chebyshev and Hoeffding inequalities, complete the following sentences. Select *the single best* option per question by completely filling in the circle next to the correct answer.

(i) For fixed ϵ , the width of the confidence interval $\alpha(\epsilon, n)$ is, as a function of the sample-size n , proportional to

n $n^{-1/2}$ n^{-1} 1 (does not depend on n)

when using

Chebyshev. Hoeffding. both Chebyshev and Hoeffding.

(ii) For fixed sample size n , the width of the confidence interval $\alpha(\epsilon, n)$ is, as a function of the failure probability ϵ , proportional to $\sqrt{\epsilon^{-1}}$ using Chebyshev's inequality and $\sqrt{\ln(2/\epsilon)}$ using Hoeffding's inequality.

Therefore, the width of the confidence interval, $\alpha(\epsilon, n)$, (*increases*, *decreases*) as ϵ approaches zero, and does so more quickly using the (*Chebyshev*, *Hoeffding*) bound.

4 Advertisement Design (10 Points)

A data scientist at a marketing firm is testing variants of an online advertisement. Each time the firm releases a variant of an advertisement, they measure its “payout” in terms of the fraction of users exposed to the ad who clicked on the ad. They assume that the payout of a given trial is independent of the payout of every other trial, the payout distribution depends only on the variant chosen, and the payout distributions are constant in time.

The firm wants to maximize the number of times their ad is clicked on, while minimizing the number of extra ad slots purchased.

Terminology: Use the variables defined here. Answers using alternate variable names for variables defined below will not receive full credit.

- m indicates the m^{th} ad slot purchased after initialization.
- $v(m)$ denotes the variant selected for testing at the m^{th} slot.
- $P_j(m)$ is the payout received if the j^{th} variant is chosen at slot m . Because the payouts are defined as the fraction of users who clicked through, they are bounded between 0 and 1.
- $M_j(m)$ denotes the number of times the j^{th} ad has been chosen before the m^{th} slot.
- $\bar{P}_j(m)$ is the observed average of the payouts received by the j^{th} variant in the $M_j(m)$ slots in which it was tested before the m^{th} slot.

- (a) [2 Pts] The data scientist suggests the Upper-Confidence Bound (UCB) algorithm. The data-scientist suggests using confidence bounds that fail with probability $\leq \epsilon$.

Selection rule: According to the UCB algorithm, at ad slot m , the firm should select the variant:

$$v(m) = \operatorname{argmax}_j \left\{ \begin{array}{l} \bar{P}_j(m) + \sqrt{\frac{\ln m}{M_j(m)}} \\ P_j(m) \end{array} \right\}$$

- (b) [1 Pt] In the formula above, draw a **circle around the exploitation term** and draw a **box around the exploration term**.

- (c) [2 Pts] The parameter ϵ is a tuning parameter that must be chosen by the data scientist. Complete the sentence below. Choose the single best answer **by filling in the circle next to it**.

To tune the algorithm to exploit more than it explores, the data scientist should

- decrease ϵ increase ϵ

(d) [3 Pts] Suppose that, instead of using UCB, the data scientist proposes Thompson sampling. Complete each statement below regarding Thompson sampling and UCB by completely filling in the circle next to each correct answer:

- (i) In Thompson sampling we learn from the observed history by iteratively updating the (Likelihood, Prior, Posterior) given the history of all observed outcomes.
- (ii) Consider the following criteria for selecting an ad: “Select the ad...
- (A) whose observed average payout is largest.”
 - (B) for which, the maximum over all true expected payouts such that the observed payouts could have plausibly occurred, is largest.”
 - (C) for which, the true expected payout could plausibly be largest given the observed payouts.”
 - (D) whose average payout is largest when sampling payout distributions from their posterior given all observations.”

Which criteria does UCB use? Choose the single best answer **by filling in the circle next to it.**

- (A) (B) (C) (D)

Which criterion does Thompson Sampling use? Choose the single best answer **by filling in the circle next to it.**

- (A) (B) (C) (D)

(e) [2 Pts] Regret $R(m)$ is defined as the difference between the total reward (payout) received up to ad slot m , and the largest possible expected reward up to time m . The expected regret is the expected value of $R(m)$ averaged over both the randomness in the outcome of each choice, and the sequence of choices, for a given algorithm.

Complete each statement below by completely filling in the circle next to each correct answer:

- (i) Using Explore-Then-Commit (ETC) with a fixed-length exploration round, the probability of selecting the wrong choice is (*shrinking*, *constant*, *growing*) in m and the expected regret is (*constant*, *logarithmic*, *linear*) in m .
- (ii) Using UCB, the probability of selecting the wrong choice is (*shrinking*, *constant*, *growing*) in m and the expected regret is (*constant*, *logarithmic*, *linear*) in m .

5 Markov Decision Processes (9 Points)

Consider a Markov decision process (MDP) with 4 states $\mathcal{S} = \{s_1, s_2, s_3, s_4\}$. The MDP obeys the following rules:

1. An agent receives the following rewards for moving *into* each state listed below.

$s' =$	s_1	s_2	s_3	s_4
$R(s') =$	0	1	1	2

2. The discount factor γ is 0.5.
3. If an agent is in state s_2 , then they have two possible actions $\mathcal{A}(s_2) = \{a_1, a_2\}$. Given that the agent selected action a , they move according to the following transition probabilities.

$s' =$	s_1	s_2	s_3	s_4
$T(s' s_2, a_1) =$	$\frac{1}{5}$	0	$\frac{4}{5}$	0
$T(s' s_2, a_2) =$	$\frac{2}{5}$	0	$\frac{1}{5}$	$\frac{2}{5}$

- (a) [4 Pts] **Definitions:** In the space below provide a complete definition of each specified function.

Your definitions should clearly state what each function returns, what it accepts as an argument, and what it, if anything, it conditions on. For example, if you were asked to define the Q-function, a correct answer would be “ $Q^\pi(s, a)$ returns the expected sum of future rewards given that the agent is in state s , takes the action a , and then acts according to policy π .”

- (i) Define the policy function π .

- (ii) Define the associated value function V^π .

- (b) [4 Pts] Suppose that a user runs value iteration to find the optimal value function of each state. After running the algorithm n times, they've found:

$s =$	s_1	s_2	s_3	s_4
$v_n^*(s) =$	3	3	2	1

In the space below, compute $v_{n+1}^*(s_2)$ using the value iteration algorithm.

- (c) [1 Pt] Suppose that $v_n^* \approx V^*$, the true value function. According to your calculation above, which action should an optimal agent choose when in state $s = s_2$?

This page has been intentionally left blank. No work on this page will be graded.