

Data 102, Fall 2022 Midterm 2

- You have 110 minutes to complete this exam. There are 6 questions, totaling 41 points.
- You may use one 8.5×11 sheet of handwritten notes (front and back). No other notes or resources are allowed.
- You should write your solutions inside this exam sheet.
- You should write your name and 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 partial solutions so we can give you partial credit.
- We have provided two blank pages of scratch paper, one at the beginning of the exam and one near the end. No work on these pages will be graded.
- You may, without proof, use theorems and facts that were given in the discussions or lectures, **but please cite them**.
- There will be no questions allowed during the exam: if you believe something is unclear, clearly state your assumptions and complete the question.
- 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	
Calcentral email (@berkeley.edu)	
Name of person to your left	
Name of person to your right	

Honor Code

I will respect my classmates and the integrity of this exam by following this honor code.

I affirm:

- All of the work submitted here is my original work.
- I did not collaborate with anyone else on this exam.

Signature: _____

This page intentionally left blank for scratch work. No work on this page will be graded.

Name: _____

SID: _____

1. (6 points) For each of the following, circle either **TRUE** or **FALSE**.
 - (a) (1 point) Confounding variables only have a causal effect on the outcome
 - (b) (1 point) SUTVA is violated in this scenario: determining if there is a casual relationship between people working out at the gym and getting Covid.
 - (c) (1 point) Matching is usually the most data-efficient way to prove the existence of a casual relationship.
 - (d) (1 point) A neural network with linear activation function cannot be used to find non-linear relationships between input features and targets.
 - (e) (1 point) For multi-armed bandits using the explore-then-commit algorithm, a good prior finds the optimal arm faster than a uniform prior.
 - (f) (1 point) Having a larger learning rate in Q-learning will cause the learner to converge to the optimal policy more quickly.

2. (7 points) You're playing a game with your friend that goes like this. Your friend asks you a question. If you answer it correctly, you move 1 meter forward. If you answer it incorrectly, you stay in place. You have a 0.5 chance of answering each question correctly, and whether you answer one question (in)correctly has no effect on whether you answer the other questions (in)correctly.
- (a) (1 point) Let X_i equal 1 if you answer question i correctly, and 0 if you answer it incorrectly. Let Z be your distance (in meters) from the starting location after answering 20 questions. Write down a formula for Z in terms of the X_i .
- (b) (2 points) Compute the mean and variance of X_i .
- (c) (2 points) Using Markov's inequality, find an upper bound on the probability that, after you have answered 20 questions, you are 16 meters or more from your starting place. (In other words, upper bound $\mathbb{P}[Z \geq 16]$.)
- (d) (2 points) What upper bound on the same probability in part 1 does Chebyshev's inequality give you?

3. (9 points) Suppose that it is March 2020 and you are trying to predict the number of COVID-19 cases. You have data on the number of cases in Alameda County for each of the first 20 days of March, which we denote as Y_1, \dots, Y_{20} . You decide to model this using the following generalized linear model:

$$Y_t \sim \text{Poisson}(\exp(\beta_0 + \beta_1 t)) \quad (1)$$

- (a) (2 points) In the above equation, what is the distribution family? What is the covariate? What is the inverse link function? (You will get credit if you write down either the link or inverse link function for this problem.)

Distribution family: _____

Covariate: _____

Inverse link function: _____

- (b) (1 point) What would go wrong if, instead of the function from part (a), you used a linear link function? Your answer should be one to two sentences.

- (c) (2 points) You want to determine the growth rate r of COVID-19 cases (e.g. if cases increase 50% each day $r = 0.5$ and if they double each day then $r = 1$).

Which parameter (β_0 or β_1) helps you estimate the growth rate? Write down a formula that estimates r in terms of that parameter.

- (d) (4 points) It is usually recommended to use a Negative Binomial instead of Poisson distribution for count data. For each of the following, answer true/false on whether it is a possible problem that could arise from using the Poisson distribution:

1. (TRUE / FALSE) The confidence intervals returned by a package such as `statsmodels` might be overly narrow.
2. (TRUE / FALSE) The log-likelihood might no longer be a good test of model fit.
3. (TRUE / FALSE) The posterior predictive distribution will be overly narrow.
4. (TRUE / FALSE) We will no longer be able to use Bayesian methods such as MCMC.

4. (5 points) Consider a Markov Decision Process whose state space is a 2×5 grid. The reward for entering each state is -2 , except for 3 terminal states. This is depicted below, with terminal states labeled with a $*$ and the start state labeled as \mathbf{S} .

-2	-2	-1*	-2	-2
10*	-2	-2 \mathbf{S}	-2	1*

There are four actions: up, down, left, right. Assume that transitions are deterministic: e.g., the left action always causes the agent to move one square to the left. Also suppose that the discount factor γ is equal to 0.9.

- (a) (2 points) Suppose the agent chooses the right action twice (after which they are in the bottom-right terminal state). What is their total discounted reward?
- (b) (2 points) Suppose instead that the agent plays the policy with optimal discounted reward. Which terminal state will the agent end up in (bottom-left, bottom-right, or top-middle)?
- (c) (1 point) What is the corresponding optimal value $V^*(\mathbf{S})$?

Name: _____

SID: _____

This page intentionally left blank for scratch work. No work on this page will be graded.

5. (7 points) An education company built an app to help students improve their scores on a standardized test. They want to know whether their app has a causal effect on student test performance. The students need to pay \$1000 to use the app. Z_j is a binary random variable indicating whether student j uses the application, and Y_j is student j 's test score. Assume that the education company knows the test score of every student. For questions a and b, you can assume a stable unit treatment value.

(a) (2 points) Can the company estimate the effect of their application by taking the average of the students who have used the app, and comparing it with the average of the rest? If not, list two confounding variables and state how they could affect the estimate.

(b) (2 points) The company randomly gives 50% of the students a coupon that lets them buy the app for free. Suppose that:

- Of the students offered the coupon, 70% use the app.
- Of those not offered the coupon, 10% use the app.
- The students who are offered the coupon on average score 1.5 points higher on the exam than the ones who are not.

Assuming the coupon is a valid instrumental variable, what is the average treatment effect of the app?

(c) (3 points) The stable unit treatment value assumption might not be true in practice. Explain the two criteria in this assumption, and for each, describe why the assumption might fail in this application.

Name: _____

SID: _____

6. (7 points) Suppose there are three advertisement formats (A, B, C) that a company is testing on its website, where the goal is to maximize the probability that a user clicks on the ad. Let Z denote the event that the user clicks on the ad (i.e. $Z = 1$ if the user clicks and 0 otherwise). We will model this as a multi-armed bandit problem with three arms (A, B , and C) and reward equal to Z (we get a reward of 1 if the user clicks and 0 otherwise).

The **ground truth** probabilities of clicking for each of the ad formats (arms) are as follows:

$$\mathbb{P}_A(Z = 1) = 0.3$$

$$\mathbb{P}_B(Z = 1) = 0.35$$

$$\mathbb{P}_C(Z = 1) = 0.4$$

For the first 12 rounds, we assign each user one of the three ad formats at random. After these 12 rounds, we observe the following results:

Table 1: Ad Click Rate Information

Ad Format	Number of views	Number of clicks
A	6	2
B	5	1
C	1	0

- (a) (1 point) What is the expected reward μ^* of the optimal arm?
- (b) (2 points) What is the pseudo-regret for these ad-display selections at the end of round 12?

- (c) (2 points) For the 13th user onward, suppose we use the upper confidence bound (UCB) with $C_t = 2 * \ln(t^3)$. Which ad will we display in round 13?
Note: $2 \cdot \ln(13^3) = 15.3897$. Recall that UCB picks the arm i with the highest value of $\hat{\mu}_i(t) + \sqrt{C_t/T_i(t)}$, where $\hat{\mu}_i(t)$ is the mean of the observed rewards for arm i and $T_i(t)$ is the number of times arm i has been pulled.

- (d) (2 points) For the 13th user onward, suppose we instead use Thompson Sampling, with an identical $\text{Beta}(\alpha = 1, \beta = 1)$ prior distribution for each ad format's click rate (eg. for $\mu_i = \mathbb{P}_i(Z = 1)$). Write down the sampling distributions for μ_A , μ_B , and μ_C in round 13 (you may express your answer in terms of common distribution families).