

SECOND MIDTERM EXAMINATION
ECON 103, STATISTICS FOR ECONOMISTS

MARCH 21ST, 2017

You will have 70 minutes to complete this exam. Graphing calculators, notes, and textbooks are not permitted.

I pledge that, in taking and preparing for this exam, I have abided by the University of Pennsylvania's Code of Academic Integrity. I am aware that any violations of the code will result in a failing grade for this course.

Name: _____

Signature: _____

Student ID #: _____ Recitation #: _____

Question:	1	2	3	4	5	6	Total
Points:	20	25	20	20	30	25	140
Score:							

Instructions: Answer all questions in the space provided, continuing on the back of the page if you run out of space. Show your work for full credit but be aware that writing down irrelevant information will not gain you points. Be sure to sign the academic integrity statement above and to write your name and student ID number on *each page* in the space provided. Make sure that you have all pages of the exam before starting.

Warning: If you continue writing after we call time, even if this is only to fill in your name, twenty-five points will be deducted from your final score. In addition, a point will be deducted for each page on which you do not write your name and student ID.

1. Let Y and Z be discrete RVs with the following joint pmf:

		Z		
		0	1	2
Y	0	1/4	3/8	1/8
	1	0	1/8	1/8

- 5 (a) Write down the support set and marginal pmf of Y .

Solution: $p_Y(0) = 3/4, p_Y(1) = 1/4$ with support set $\{0, 1\}$

- 5 (b) Write down the support set and marginal pmf of Z .

Solution: $p_Z(0) = 1/4, p_Z(1) = 1/2, p_Z(2) = 1/4$ with support set $\{0, 1, 2\}$

- 5 (c) Calculate $E[YZ]$.

Solution: $E[YZ] = 1 \times 1 \times 1/8 + 1 \times 2 \times 1/8 = 3/8$.

- 5 (d) Write down the conditional pmf of Y given that $Z = 1$.

Solution: $p_{Y|Z}(0|1) = (3/8)/(4/8) = 3/4$ and $p_{Y|Z}(1|1) = (1/8)/(4/8) = 1/4$.

2. Answer each part. No explanation is required.

- 4 (a) If X and Y are RVs with variance one and correlation $1/2$ what is $Var(X - Y)$?

Solution: $Var(X - Y) = \sigma_X^2 + \sigma_Y^2 - 2\sigma_X\sigma_Y\rho_{XY} = 1 + 1 - 2 \times 1/2 = 1$

- 4 (b) Suppose $X_1, \dots, X_n \sim \text{iid}$ with mean μ and variance σ^2 . Write down the formula for the standard error of \bar{X}_n .

Solution: $SE(\bar{X}_n) = \sigma/\sqrt{n}$

- 4 (c) If $X \sim N(\mu = 3, \sigma^2 = 9)$ approximately what is the value of $P(0 \leq X \leq 6)$?

Solution: 0.68

- 4 (d) Write R code to calculate the probability that a standard normal RV takes on a value between -1 and 3 .

Solution: `pnorm(3) - pnorm(-1)`

- 4 (e) Let $Z \sim N(0, 1)$. Write R code to calculate c such that $P(-c \leq Z \leq c) = 0.8$.

Solution: We can calculate $-c$ using `qnorm(0.1)` or c using `qnorm(0.9)`.

- 5 (f) Suppose we observe a sample of 16 iid observations from a $N(\mu, \sigma^2)$ population and we know that $\sigma^2 = 1$. Our sample mean \bar{x} turns out to equal 3. Write down an approximate 95% confidence interval for μ .

Solution: $\bar{x} \pm 2 \times \sigma / \sqrt{n} = 3 \pm 2 \times 1/4 = 3 \pm 0.5$ or equivalently (2.5, 3.5).

3. Let X be a continuous RV with support set $[0, 1]$ and pdf $f(x) = a + bx^2$ where $a, b > 0$.

- 6 (a) Calculate the CDF of X , $F(x_0)$, in terms of a and b .

Solution:

$$F(x_0) = \int_{-\infty}^{x_0} f(x) dx = \int_0^{x_0} (a + bx^2) dx = \left(ax + \frac{bx^3}{3} \right) \Big|_0^{x_0} = ax_0 + \frac{bx_0^3}{3}$$

for $x_0 \in [0, 1]$. For $x_0 < 0$ we have $F(x_0) = 0$ and for $x_0 > 1$ we have $F(x_0) = 1$.

- 6 (b) Calculate $E[X]$ in terms of a and b .

Solution:

$$E[X] = \int_{-\infty}^{\infty} xf(x) dx = \int_0^1 (ax + bx^3) dx = \left(\frac{ax^2}{2} + \frac{bx^4}{4} \right) \Big|_0^1 = \frac{a}{2} + \frac{b}{4}$$

- 8 (c) Since a and b are positive, $f(x) \geq 0$. This is one of the two conditions required to ensure that $f(x)$ is a valid pdf. What is the other condition? What restriction must we place on a and b to ensure that the condition is satisfied?

Solution: To be a valid pdf $f(x)$ must be non-negative and integrate to one over $(-\infty, +\infty)$. We already know that $f(x) \geq 0$ so the condition we need is

$$\int_0^1 f(x) dx = F(1) = a + \frac{b}{3} = 1 \text{ which holds so long as } b = 3(1 - a).$$

- 20 4. Let $X_1, \dots, X_n \sim \text{iid } N(\mu, \sigma^2 = 100)$ and define $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$ as usual. Find the

Name: _____

Student ID #: _____

value of n for which $P(\mu - 5 \leq \bar{X}_n \leq \mu + 5) \approx 0.95$. Your answer should give a specific numeric value for n and should *not* involve any R commands.

Solution: Since the X_i are iid $N(\mu, \sigma^2)$ it follows that $\bar{X}_n \sim N(\mu, \sigma^2/n)$ where $\sigma^2/n = 100/n$. Thus the standard deviation of \bar{X}_n , aka the standard error of the mean, equals $10/\sqrt{n}$. Hence

$$\begin{aligned} P(\mu - 5 \leq \bar{X}_n \leq \mu + 5) &= P(-5 \leq \bar{X}_n - \mu \leq 5) \\ &= P(-\sqrt{n}/2 \leq (\bar{X}_n - \mu)/(10/\sqrt{x}) \leq \sqrt{n}/2) \\ &= P(-\sqrt{n}/2 \leq Z \leq \sqrt{n}/2) \end{aligned}$$

where $Z \sim N(0, 1)$. So the problem reduces to finding the value of n such that a standard normal has approximately a 95% chance of taking on a value between $-\sqrt{n}/2$ and $\sqrt{n}/2$. There is approximately a 95% chance that a standard normal takes on a value between -2 and 2 so we solve $\sqrt{n}/2 = 2$ for n , yielding $n = 16$.

This problem was part of your homework for Lectures 13–14. I have reworded the problem slightly to make it clearer, but the details and solution are unchanged.

5. Let S denote the total number of successes in n iid Bernoulli trials, each with probability of success π . Consider two estimators of π : the sample proportion $P = S/n$ and an alternative estimator $P^* = \frac{S+1}{n+1} = \left(\frac{n}{n+2}\right)P + \left(\frac{1}{n+2}\right)$.

- 5 (a) Calculate the bias of P .

Solution: Since S is a Binomial(n, π) RV, $E[S] = n\pi$ so by the linearity of expectation $E[P] = E[S]/n = \pi$. Thus $\text{Bias}(P) = E[P] - \pi = 0$.

- 5 (b) Calculate $\text{Var}(P)$.

Solution: Since S is a Binomial(n, π) RV, $\text{Var}(S) = n\pi(1 - \pi)$ and thus $\text{Var}(P) = \text{Var}(S/n) = \text{Var}(S)/n^2 = n\pi(1 - \pi)/n^2 = \pi(1 - \pi)/n$.

- 5 (c) Is P consistent for π ? Explain briefly.

Solution: $\text{MSE}(P) = \text{Bias}(P)^2 + \text{Var}(P)^2 = 0 + \pi(1 - \pi)/n$. Since $\text{MSE}(P) \rightarrow 0$ as $n \rightarrow \infty$, P is consistent for π .

- 5 (d) Calculate the bias of P^* .

Solution: First, $E[P^*] = E\left[\left(\frac{n}{n+2}\right)P + \left(\frac{1}{n+2}\right)\right] = \left(\frac{n\pi+1}{n+2}\right)$ by the linearity of expectation. Hence $\text{Bias}(P^*) = \frac{n\pi+1}{n+2} - \pi = \frac{1-2\pi}{n+2}$.

- 5 (e) Calculate $\text{Var}(P^*)$.

Solution: $\text{Var}(P^*) = \text{Var}\left[\left(\frac{n}{n+2}\right)P + \left(\frac{1}{n+2}\right)\right] = \left(\frac{n}{n+2}\right)^2 \text{Var}(P) = \frac{n\pi(1-\pi)}{(n+2)^2}$

- 5 (f) Is P^* consistent for π ? Explain briefly.

Solution: $\text{MSE}(P^*) = \text{Bias}(P^*)^2 + \text{Var}(P^*) = \left(\frac{1-2\pi}{n+2}\right)^2 + \frac{n\pi(1-\pi)}{(n+2)^2}$. Now if we take $n \rightarrow \infty$, $\text{MSE}(P^*) \rightarrow 0$ so P^* is consistent for π .

6. At the time of this writing approximately 50% of Americans disapprove of President Trump according to a weighted average of polls constructed by FiveThirtyEight. Assume for the purposes of this question that this estimate is *exactly correct* so that precisely half of the US population disapproves of Trump. This means that we can model a poll

based on a random sample of n Americans as an iid sequence of n Bernoulli($1/2$) RVs where a 1 indicates that a given individual in the poll *disapproves* of Trump.

- 10 (a) Write an R function called `disapprove_prop` that simulates a poll asking a random sample of n Americans if they disapprove of Trump. Your function should take a single input argument `n` the sample size of the poll and then make n iid simulation draws from the population described in the problem statement. The output of `disapprove_prop` should be the sample proportion of individuals who disapprove of Trump calculated from your `n` simulation draws.

Solution: Various possible solutions. Here is one possibility:

```
disapprove_prop <- function(n){
  draws <- sample(0:1, size = n, replace = TRUE)
  return(mean(draws))
}
```

- 5 (b) Suppose that you wanted to learn about the accuracy of the sample proportion in the poll described in the problem statement when $n = 10$. Write R code to approximate its sampling distribution based 10000 Monte Carlo simulations using your function `disapprove_prop` from part (a). Store your simulations in a vector called `poll_sims`.

Solution: `poll_sims <- replicate(10000, disapprove_prop(10))`

- 10 (c) Continuing from (b) suppose I entered `mean(poll_sims)` and `var(poll_sims)` at the R console. Approximately what result would I get for each? Explain briefly.

Solution: The sample proportion is exactly the same thing as the sample mean in a sample that only contains zeros and ones, as we have in this question. We know from class that, under random sampling, the sample mean is an unbiased estimator of the population mean and the variance of its sampling distribution equals σ^2/n where σ^2 is the population variance. In this example the population is Bernoulli($1/2$), so the population mean is $1/2$, so `mean(poll_sims)` will be approximately 0.5. Since the variance of a Bernoulli($1/2$) RV is $1/4$ and $n = 10$, `var(poll_sims)` will be approximately $0.25/10 = 0.025$.