Queen's University
Department of Mathematics and Statistics

MTHE/STAT 353

Final Examination, Part 2 Solutions, 2020
Instructor: G. Takahara

INSTRUCTIONS:

- The exam is in two parts. This is Part 2 of the exam.
- Part 2 has 3 questions, each worth 10 marks. For each question, begin each solution at the start of a fresh page, and put your student number at the start of each solution.
- The 3 solutions for Part 2 are to be submitted through crowdmark. You should have received an email inviting you to submit your solutions for Part 2 to crowdmark. Upload each solution separately.
- THE DEADLINE FOR SUBMISSION OF PART 2 IS 12:00 NOON ON APRIL 15. **THERE WILL BE ABSOLUTELY NO EXTENSIONS**. YOU MUST ALLOW TIME TO PREPARE YOUR SOLUTIONS FOR UPLOAD TO CROWDMARK, SO PLAN ON COMPLETING THE EXAM BY 11AM ON APRIL 15. SO YOU HAVE 24 HOURS TO WRITE YOUR SOLUTIONS STARTING FROM 11AM ON APRIL 14, THOUGH YOU CAN WORK ON PART 2 FROM THE TIME IT IS POSTED TO THE COURSE WEBPAGE AT 7PM ON APRIL 13.
- Part 2 of the exam is open book. This means that you can use your notes, the textbook, and your computer.
- **ABSOLUTELY ZERO COLLABORATION IS ALLOWED ON THE EXAM.** There is to be no collaboration in any form on any question on any part of the exam, either in person or remotely. All work on the exam must be completed *on your own*.

Instructions continued on page 2.

- "The candidate is urged to submit with the answer paper a clear statement of any assumptions made if doubt exists as to the interpretation of any question that requires a written answer."

- You may write your solutions in the space provided, continuing on your own paper if needed, or you may write your solutions using your own paper.

- SHOW YOUR WORK CLEARLY. Correct answers without clear work showing how you got there will not receive full marks. Marks per part question are shown in brackets at the right margin.

Problem 1 [10] Problem 2 [10] Problem 3 [10]

Total: [30] ____________

____
Student Number

1(a). Let $X_1, \ldots, X_n$ be random variables, each with mean 0 and variance σ^2, but not necessarily independent or identically distributed. Assume that the correlation coefficient between any pair of the X_i's is the same, and given by ρ. Show that

$$\rho \geq -\frac{1}{n-1} + P(|\overline{X}| > \sigma\sqrt{1 - 1/n}),$$

where $\overline{X} = \frac{1}{n}\sum_{i=1}^{n} X_i$. [5]

Solution: The variance of $\overline{X}$ is

$$\mathrm{Var}(\overline{X}) = \frac{1}{n^2}\left(\sum_{i=1}^{n}\mathrm{Var}(X_i) + \sum_{i\neq j}\mathrm{Cov}(X_i, X_j)\right) = \frac{1}{n^2}\Big(n\sigma^2 + n(n-1)\rho\sigma^2\Big) = \frac{\sigma^2}{n}(1+(n-1)\rho).$$

Since the mean of $\overline{X}$ is zero, by Chebyshev's inequality we have

$$P\left(|\overline{X}| > \sigma\sqrt{1 - \frac{1}{n}}\right) \leq \frac{\mathrm{Var}(\overline{X})}{\sigma^2(1 - 1/n)} = \frac{1 + (n-1)\rho}{n-1} = \frac{1}{n-1} + \rho$$

Therefore, the desired inequality follows.

(b) Let $X_1, \ldots, X_n$ be independent Poisson(1) random variables and let $\overline{X} = \frac{1}{n}\sum_{i=1}^n X_i$. Let $k > 1$ be given. Show that

$$P(\overline{X} \geq k) \leq \left(\frac{e^{k-1}}{k^k}\right)^n. \qquad [5]$$

Solution: We find the mgf of $\overline{X}$ and then apply Chernoff's bound. The mgf of each X_i is

$$M_{X_i}(t) = E[e^{tX_i}] = \sum_{j=0}^{\infty} e^{tj}\frac{1}{j!}e^{-1} = e^{-1}e^{e^t} = e^{e^t-1}.$$

Then the mgf of $\overline{X}$ is given by

$$M_{\overline{X}}(t) = E[e^{t\overline{X}}] = M_{X_i}\left(\frac{t}{n}\right)^n = e^{n(e^{t/n}-1)}.$$

Chernoff's bound gives

$$P(\overline{X} > k) \leq \min_{t>0} \frac{e^{n(e^{t/n}-1)}}{e^{tk}} = e^{n(e^{t/n}-1)-tk}. \qquad (1)$$

Differentiating the exponent with respect to t gives $e^{t/n} - k$. Setting this to 0 and solving for t gives $t = n\ln k$. Plugging this into the RHS of Eq.(1) then gives

$$P(\overline{X} > k) \leq e^{n(k-1)-\ln k^{kn}} = \frac{e^{n(k-1)}}{k^{kn}} = \left(\frac{e^{k-1}}{k^k}\right)^n,$$

as desired.

Student Number

2(a). Let $X_1, X_2, \ldots$ and $Y_1, Y_2, \ldots$ be two sequences of random variables. Nothing is assumed about these random variables other than that they all have finite mean and variance. Let X and Y also be random variables, both with finite mean and variance.

(a) Give an example where $X_n \to X$ in distribution and $Y_n \to Y$ in distribution as $n \to \infty$, but $X_n + Y_n$ does not converge to $X + Y$ in distribution as $n \to \infty$. [2]

(b) If X is independent of all the X_i's and $\mathrm{Var}(X) > 0$, show that X_n cannot converge to X in mean square. [3]

Solution:

(a) Let X and Y be independent $N(0,1)$ random variables. Let $X_n = X$ and $Y_n = -X$ for all n. Then, clearly $X_n \to X$ in distribution. But $-X$ and X have the same distribution since the $N(0,1)$ distribution is symmetric about 0. So $Y_n \to Y$ in distribution as well. But $X_n + Y_n$ is equal to 0 for all n and all ω, so clearly $X_n + Y_n \to 0$ in distribution. But the distribution of $X + Y$ is $N(0,2)$ (the distribution of a sum of 2 independent $N(0,1)$ random variables). So $X_n + Y_n$ does not converge to $X + Y$ in distribution.

(b) Let μ_n and μ denote the mean of X_n and X, respectively, and let σ_n^2 and σ^2 denote the variance of X_n and X, respectively. If $X_n \to X$ in mean square as $n \to \infty$ then, by definition, we have $E[(X_n - X)^2] \to 0$ as $n \to \infty$. But if X_n and X are independent, then

$$
\begin{aligned}
E[(X_n - X)^2] &= E[X_n^2] - 2E[X_n X] + E[X^2] \\
&= E[X_n^2] - 2E[X_n]E[X] + E[X^2] \\
&= \sigma_n^2 + \mu_n^2 - 2\mu_n\mu + \sigma^2 + \mu^2 \\
&= \sigma_n^2 + \sigma^2 + (\mu_n - \mu)^2 \\
&\geq \sigma^2
\end{aligned}
$$

for all n. By assumption $\sigma^2 > 0$ so it is not possible that $E[(X_n - X)^2] \to 0$ as $n \to \infty$.

(b) Let $Y_1, Y_2, \ldots$ be a sequence of discrete random variables such that the joint probability mass function of $(Y_1, \ldots, Y_n)$ is

$$P(Y_1 = y_1, \ldots, Y_n = y_n) = \begin{cases} \left[(n+1)\binom{n}{\sum_{i=1}^n y_i}\right]^{-1} & \text{for } y_i \in \{0,1\},\ i = 1, \ldots, n \\ 0 & \text{otherwise} \end{cases}$$

Let X_n be the sample mean of $Y_1, \ldots, Y_n$, i.e., $X_n = \frac{1}{n}\sum_{i=1}^n Y_i$. Show that X_n converges in distribution to a limit X and find the distribution of X. [5]

Solution: Let $S_{k,n} = \{(y_1, \ldots, y_n) : y_i \in \{0,1\} \text{ and } \sum_{i=1}^n y_i = k\}$. Then $S_{k,n}$ contains $\binom{n}{k}$ elements, and

$$\begin{aligned} P\Big(\sum_{i=1}^n Y_i = k\Big) &= P((Y_1, \ldots, Y_n) \in S_{k,n}) \\ &= \sum_{(y_1,\ldots,y_n)\in S_{k,n}} \left[(n+1)\binom{n}{\sum_{i=1}^n y_i}\right]^{-1} \\ &= \sum_{(y_1,\ldots,y_n)\in S_{k,n}} \left[(n+1)\binom{n}{k}\right]^{-1} \\ &= \binom{n}{k}\left[(n+1)\binom{n}{k}\right]^{-1} = \frac{1}{n+1}. \end{aligned}$$

In other words, $\sum_{i=1}^n Y_i$ has a discrete uniform distribution on $\{0, 1, \ldots, n\}$. Now let $x \in [0,1]$. Then

$$P(X_n \le x) = P\Big(\sum_{i=1}^n Y_i \le nx\Big) = \frac{\lfloor nx \rfloor + 1}{n+1} = \frac{nx - c(n,x) + 1}{n+1} \to x$$

as $n \to \infty$, where $c(n,x)$ is some value satisfying $0 \le c(n,x) < 1$ for all n and all $x \in [0,1]$. Clearly, $P(X_n \le x) = 1$ for all $x > 1$ and $P(X_n \le x) = 0$ for all $x < 0$. Therefore, the cdf of X_n converges to the cdf of a uniform distribution on $(0,1)$, at all x. In other words, $X_n \to X$ in distribution, where $X \sim U(0,1)$.

Student Number

3(a). Let $X_1, X_2 \ldots$ be an infinite sequence of independent and identically distributed random variables, each with finite mean μ and finite variance σ^2. Use the strong law of large numbers to show that, with probability 1, infinitely many of the X_i's must be greater than or equal to μ. [4]

Solution: If $\sigma^2 = 0$ then $X_i = \mu$ with probability 1 for all i. If $A_i = \{\omega \in \Omega : X_i(\omega) = \mu\}$ then $P(A_i) = 1$ for all i. Then $A = \cap_{i=1}^{\infty} A_i$ has probability 1 as well. But if $\omega \in A$ then $X_i(\omega) = \mu$ for all i. Therefore, infinitely many of the X_i (in fact all of them) are greater than or equal to μ (they are equal to μ) with probability 1. Now suppose $\sigma^2 > 0$. Then for any i, if $P(X_i \geq \mu) = 1$ then X_i must equal μ with probability 1 (since $E[X_i] = \mu$). But this contradicts $\sigma^2 > 0$. Therefore, $P(X_i \geq \mu)$ must be strictly less than 1. Similarly, $P(X_i \leq \mu)$ must be strictly less than 1. Together, these imply $0 < P(X_i \geq \mu) < 1$. Now define $Y_i = I_{[\mu,\infty)}(X_i)$. Then $E[Y_i] = P(X_i \geq \mu)$ and the Y_i's are i.i.d. (and they have finite variance). So by the strong law of large numbers $\overline{Y}_n = \frac{1}{n}\sum_{i=1}^{n} Y_i$ converges to $P(X_i \geq \mu)$ with probability 1. Let B denote the set on which this convergence occurs, so $P(B) = 1$. Let $\omega \in B$. Then $\overline{Y}_n(\omega)$ converges to $P(X_i \geq \mu)$, which is strictly between 0 and 1. Since each $Y_i(\omega)$ is either 0 or 1 the only way this convergence can happen is if infinitely many of the $Y_i(\omega)$'s are equal to 0 and infinitely many of them are equal to 1. But to say that infinitely many of the $Y_i(\omega)$'s are equal to 1 is the same as saying that infinitely many of $X_i(\omega)$'s are greater than or equal to μ. In other words, if $\omega \in B$, then $X_i(\omega) \geq \mu$ for infinitely many i's. Since $P(B) = 1$ we conclude that infinitely many of the X_i's are greater than or equal to μ with probability 1.

(b) Consider a sequence of independent experiments, where in each experiment we take k balls, labelled 1 to k and randomly place them into k slots, also labelled 1 to k, so that there is exactly one ball in each slot. For the ith experiment, let X_i be the number of balls whose label matches the slot label of the slot into which it is placed. So $X_1, X_2, \ldots$ is a sequence of independent and identically distributed random variables. Use the central limit theorem to approximate the probability that in the first 25 experiments the total number of balls whose label matches their slot label is greater than 30. [6]

Solution: We first find the mean and variance of X_i. Writing $X_i = X_{i1} + \ldots + X_{ik}$, where X_{ij} is the indicator that slot j receives ball j in the ith experiment, we have that

$$E[X_i] = \sum_{j=1}^{k} E[X_{ij}] = \sum_{j=1}^{k} P(X_{ij} = 1),$$

where $P(X_{ij} = 1)$ is the probability that slot j receives ball j. Since all assignments of the balls to the slots are equally likely, a simple counting argument gives that

$$P(X_{ij} = 1) = \frac{(k-1)!}{k!} = \frac{1}{k}.$$

Therefore, $E[X_i] = \sum_{j=1}^{k} \frac{1}{k} = 1$. Similarly,

$$E[X_i^2] = E[(X_{i1} + \ldots + X_{ik})^2] = \sum_{j=1}^{k} E[X_{ij}^2] + \sum_{j \neq r} E[X_{ij} X_{ir}] = 1 + \sum_{j \neq r} E[X_{ij} X_{ir}]. \quad (2)$$

Another counting argument yields

$$\begin{aligned} E[X_{ij} X_{ir}] &= P(X_{ij} = 1, X_{ir} = 1) = P(\text{ball } j \text{ goes in slot } j \text{ and ball } r \text{ goes in slot } r) \\ &= \frac{(k-2)!}{k!} = \frac{1}{k(k-1)}. \end{aligned}$$

Since there are $k(k-1)$ terms in the final sum in Eq.(2), we have $E[X_i^2] = 2$, and so

$$\text{Var}(X_i) = E[X_i^2] - E[X_i]^2 = 2 - 1^2 = 2 - 1 = 1.$$

The total number of balls in the first 25 experiments whose label matches their slot label is $\sum_{i=1}^{25} X_i$, where $X_1 \ldots, X_{25}$ are iid with $E[X_i] = 1$ and $\text{Var}(X_i) = 1$. By the CLT

$$P\left(\sum_{i=1}^{25} X_i > 30\right) = P\left(\frac{\sum_{i=1}^{25} X_i - 25}{5} > \frac{30 - 25}{5}\right) \approx P(Z > 1) = 1 - \Phi(1) = 0.1587,$$

where $Z \sim N(0, 1)$.