Ph.D. Comprehensive Exam, August 2020

Instructions: *The exam has 6 problems, most with multiple parts. All of the problems will be graded. Write your code words on each of your answer sheets. Do not put your name or UNM ID on any of the sheets. Be clear, concise, and complete. All solutions should be rigorously* ***explained****.*

Problem 1. Consider a linear model

$$Y = X\beta + e, \quad \mathrm{E}[e] = 0, \quad \mathrm{Cov}\,[e] = \sigma^2 I.$$

Let M be the perpendicular projection operator onto $C(X)$. Consider the parameter $\lambda'\beta$ where λ is known. For known vectors a and ρ, suppose $a'Y$ and $\rho'Y$ are both unbiased estimates of $\lambda'\beta$.

(a) Show that $\lambda'\beta$ is estimable.

(b) Show that $a'X = \rho'X$.

(c) Show that $\mathrm{Var}[a'Y] = \mathrm{Var}[a'Y - \rho'MY] + \mathrm{Var}[\rho'MY]$.

(d) State and prove the Gauss-Markov Theorem.

Problem 2. For a linear model $Y \sim \mathrm{N}(X\beta, \sigma^2 I)$, consider least squares estimates $\hat{\beta}$ and the MSE and suppose that independently a random r vector has $\tilde{Y} \sim \mathrm{N}\left(\tilde{X}\beta, \sigma^2 I_r\right)$. Also assume that $\tilde{X} = BX$ for some $r \times n$ matrix B and specify an r vector ρ.

(a) Find the distribution of $\rho'\tilde{Y}$.

(b) Find the distribution of $\rho'\tilde{Y} - \rho'\tilde{X}\hat{\beta}$.

(c) Why did we require that $\tilde{X} = BX$?

(d) Let A be the constant such that $\mathrm{Var}\left[\rho'\tilde{Y} - \rho'\tilde{X}\hat{\beta}\right] = \sigma^2 A$. Find the distribution of $(\rho'\tilde{Y} - \rho'\tilde{X}\hat{\beta})/\sqrt{MSE\,A}$.

(e) Find the form of a $1-\alpha\%$ prediction interval for $\rho'\tilde{Y}$.

(f) Assuming that the model for Y is a regression model, find an explicit form for the matrix B in $\tilde{X} = BX$.

Problem 3. Let $X_1, X_2, \ldots$ be independent with X_n taking the values $-\sqrt{n-1}, -1, 1, \sqrt{n-1}$ each with probability $1/4$. Use the Lindeberg condition to show that the sample mean converges in distribution to a $\mathrm{N}(0, .25)$. Hint: $\sum_{i=1}^{r} i = r(r+1)/2$; $\sum_{i=1}^{r} i^2 = r(r+1)(2r+1)/6$.

Problem 4. Consider a Bayesian model for 10 conditionally independent Bernoulli trials, where the probability of success on a trial is $\theta \sim \text{Beta}(1, 1)$ and the total number of successes on 10 trials is $X|\theta \sim \text{Bin}(10, \theta)$. (For this problem, remember that $\Gamma(a) = (a-1)!$ for all $a \in \mathbb{N}$.)

(a) Find the marginal probability that X, the number of successes on 10 trials, is odd: $\Pr[X \in \{1, 3, 5, 7, 9\}]$.

(b) Assume that $X = 7$ successes are observed—and then a further 4 trials are performed, exchangeable with the first 10 trials. Let Y be the total number of successes on these 4 new trials. Find the predictive probability that $Y = 3$.

Problem 5. Let $X_1, X_2, \ldots, X_n$ be a random sample from a scale-uniform distribution $X_i \sim \text{Uniform}((1-k)\theta, (1+k)\theta)$, with unknown mean $\text{E}[X_i] = \theta$ and known design parameter $k \in (0, 1)$. Let $X_{(1)} \equiv \min_i\{X_i\}$ and $X_{(n)} \equiv \max_i\{X_i\}$. $\text{E}\left[X_i|X_{(1)}, X_{(n)}\right] = (X_{(1)} + X_{(n)})/2$.

(a) What is the mean and variance of X_i?

(b) Show that $(X_{(1)}, X_{(n)})$ is a sufficient statistic.

(c) Give an intuitive argument for why $(X_{(1)}, X_{(n)})$ is minimal sufficient. (Extra credit for a rigorous argument.)

(d) Why is $\text{E}\left[X_i|X_{(1)}, X_{(n)}\right] = (X_{(1)} + X_{(n)})/2$ a better estimate of θ than X_i.

(e) The variance of $(X_{(1)} + X_{(n)})/2$ is $2k^2\theta^2/[n+1](n+2)$. Establish whether this is larger or smaller than the variance of the sample mean $\overline{X}$.

(f) For some C there is an unbiased estimate of the form $C[(1-k)X_{(1)} + (1+k)X_{(n)}]$ with variance $2k^2\theta^2/[(1+k^2)(n-1)+2](n+2)$. Establish whether this is larger or smaller than the variance of $(X_{(1)} + X_{(n)})/2$.

(g) Prove that $(X_{(1)}, X_{(n)})$ is not complete. (Hint: What would happen if it were complete?)

Problem 6. Let $X \sim \text{N}(0, \sigma^2)$. We want to use this one observation to test the null $H_0 : \sigma^2 = 1$ versus the composite alternative $H_A : \sigma^2 < 1$.

(a) Ignoring the alternative, what data are least consistent with the null hypothesis.

(b) Find the most powerful test for the simple alternative $H_A : \sigma^2 = .5$.

(c) Show that the problem has monotone likelihood ratio.

(d) Find the uniformly most powerful test for the composite alternative.

(e) Does rejecting this test actually suggest that the null hypothesis is wrong or does it merely suggest that there are better alternatives.