Stat 643 Exam 2

May 3, 2007

I have neither given nor received unauthorized assistance on this examination.

____ ________________

signature date

There are 10 "small" problems on this exam. Some are easier than others. All will be scored out of 10 points to make a total possible score of 100 pints.

1. Suppose that the figure below shows risk vectors for all non-randomized rules in a 2-state decision problem. Identify (give the coordinates of) the (possibly randomized) risk vectors that are

a) minimax
b) Bayes versus a prior that is uniform on $\Theta = \{\theta_1, \theta_2\}$

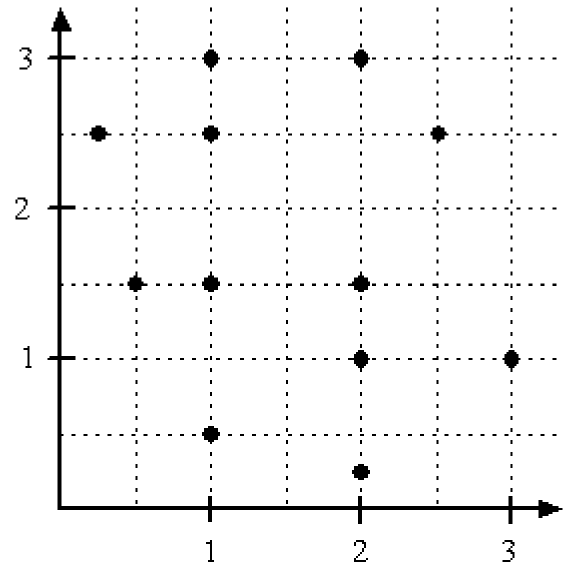

2. Suppose that X and Y are independent, $X \sim \text{Bi}(n, p)$ and $Y \sim \text{Ber}(p)$. Consider squared error loss estimation of p and an estimator $\hat{p} = \delta(X)$. Let $W = X + Y$. Find an estimator $\hat{p}^* = \gamma(W)$ that improves upon $\hat{p}$. (Carefully say why your new estimator is better in terms of risk.)

3. Consider a model for an observable X with real parameter $\alpha \in (0,1)$ and pdf on $(0,1)$

$$f(x \mid \alpha) = 1 + \alpha(2x - 1)$$

Write out explicitly as possible the Fisher information in X about α at α_0, $I_X(\alpha_0)$. (You'll probably have to leave this in terms of an integral that would need to be evaluated numerically.)

4. In the context of problem 3, $I_X(\alpha)$ can be computed numerically and looks as in the plot below. Based on the information in this plot, if $X_1, X_2, \ldots, X_{100}$ are iid with marginal density $f(x|\alpha)$,

a) what do you propose as an approximation to the distribution of an "MLE" of α if it is the case that $\alpha = .5$, and

b) what do you propose as (realizable/implementable) "large sample conservative" approximately 95% two-sided confidence limits for α, based on an "MLE" of α, say $\hat{\alpha}_{100}$?

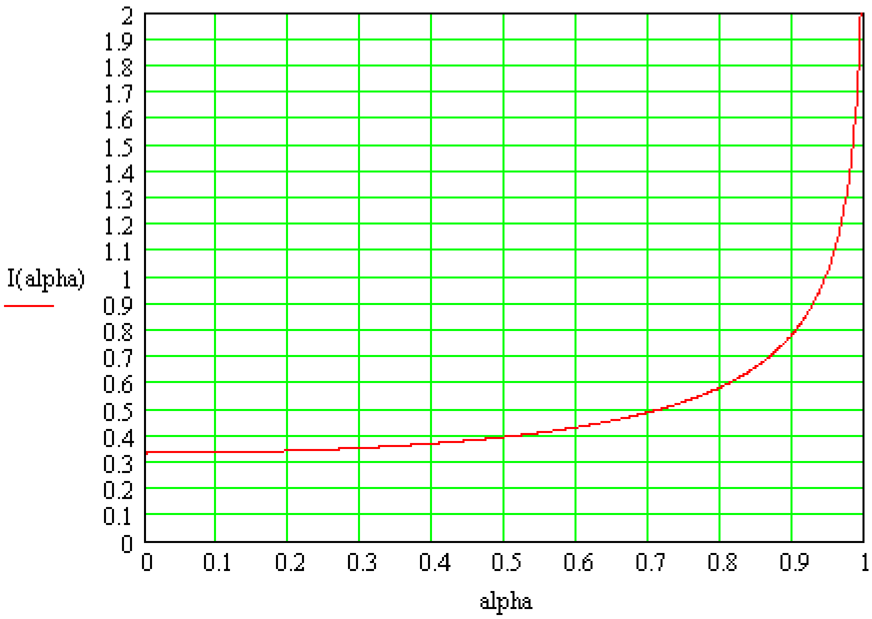

5. Again in the context of problem 3, if $X_1, X_2, \ldots, X_n$ are iid with marginal density $f(x \mid \alpha)$,

a) what is the likelihood equation (that would need to be solved in order to find an MLE of α)?

b) give an *explicit* formula for an estimator of α that will be "asymptotically efficient" here.

6. Consider the problem of likelihood ratio testing in the *non-regular* family of $\text{Uniform}(0,\theta)$ distributions. That is, suppose that $X_1, X_2, \ldots, X_n$ are iid $\text{Uniform}(0,\theta)$ and consider testing $\text{H}_0\text{:}\theta = \theta_0$ versus $\text{H}_\text{a}\text{:}\theta \neq \theta_0$. Let

$$\Lambda_n = 2\left(\sup_\theta L_n(\theta) - L_n(\theta_0)\right)$$

where (as usual) $L_n(\theta)$ is the n observation log-likelihood. What is the large sample distribution of Λ_n under the null hypothesis? (Hint: You may use without proof the facts that the MLE of θ is $\max_{i=1,2,\ldots,n} X_i$, and that under the null hypothesis, $n\left(\theta_0 - \max_{i=1,2,\ldots,n} X_i\right)$ is asymptotically Exponential with mean θ_0.)

7. Suppose that as on the Exponential Families handout, for $\eta \in \Gamma \subset \Re$, distributions P_η have densities

$$f_\eta(x) = K(\eta)\exp(\eta T(x))h(x)$$

If $X_1, X_2, \ldots, X_n$ are iid P_{η_0} and $L_n(\eta)$ is the n observation log-likelihood, to what does $\frac{1}{n}L_n(\eta)$ converge in probability? (Give a formula in terms of the factors of f_η.)

8. Bayesians sometimes argue that their "posteriors are consistent." Consider the simplest possible version of this. Suppose that $\Theta = \{1,2\}$ and for some σ-finite measure μ, f_1 and f_2 are densities for two different probability distributions P_1 and P_2. For a prior distribution G on Θ, with $g_i = G(\{i\}) \in (0,1)$ for $i = 1,2$ and $X_1, X_2, \ldots, X_n$ iid with marginal one of P_1 and P_2,

a) what is the posterior probability (based on the n observations) that $\theta = 1$?
b) argue carefully that under the $\theta = 1$ model, the posterior probability that $\theta = 1$ from a) converges to 1 in probability.

9. Suppose that $X \sim \mathrm{Bi}(2,p)$ and that squared error loss of $p \in [0,1]$ is under consideration. Argue directly that for any $c \in (0,1)$ the estimator

$$\delta_c(X) = \begin{cases} 0 & \text{if } x = 0 \\ c & \text{if } x = 1 \\ 1 & \text{if } x = 2 \end{cases}$$

is admissible. (Hint: What must be the $p=0$ risk and the $p=1$ risk of any ϕ improving upon a δ_c. What does that say about the form of ϕ?)

10. An alternative to 0-1 loss for decision-theoretic treatments of some testing problems is the so-called "linear loss." That is, for an interval $\Theta \subset \Re$ and action space $\mathcal{A} = \{0,1\}$, if $\theta_0 \in \Theta$, testing $\mathrm{H}_0{:}\theta \le \theta_0$ vs $\mathrm{H}_a{:}\theta > \theta_0$ might be phrased in decision theoretic terms using a loss function

$$L(\theta,a) = \begin{cases} \max\{0, \theta - \theta_0\} & \text{for } a = 0 \\ \max\{0, \theta_0 - \theta\} & \text{for } a = 1 \end{cases}$$

For a prior G on Θ, and distributions P_θ of an observable X, what is the form of a Bayes rule for G? (Hint: Consider $L(\theta,0) - L(\theta,1)$.)