Statistics 582, Problem Set 4

Wellner; 1/28/2015

Reading: Lehmann, TPE, Chapter 5, sections 5.1 and 5.2.
Course Notes, Chapter 5, sections 4-7.

Due: Wednesday, February 4, 2015.
Reminder: Make up lecture 2, 11:30 - 12:20, Wednesday, 4 February, SIG 226.
Reminder: Midterm exam: Friday, 13 February.

1. Consider Example 1.3 as given on page 4 of the Chapter 5 notes. (a) Consider calculation of the Bayes rule for the prior $\lambda_1 = \lambda$, $\lambda_2 = 1 - \lambda$ with $\lambda \in [0,1]$ via computation of the posterior risk(s) as in Theorem 5.1: in particular, compute the posterior risks
$$E\{L(\boldsymbol{\theta}, d(\cdot|X))|X = x\}$$
for $x \in \{R, B, G\}$.
(b) Specialize your calculation in (a) to the particular prior $\lambda = 1/2$ and thereby show that the non-randomized rule $d_2 = (0, 0, 1)$ is the Bayes rule.
(c) Determine a prior which yields all the risk points on the the line between $(1, 3.6) = (R(1, d_7), R(2, d_7))$ and $(0, 6) = (R(1, d_8), R(2, d_8))$ as risks of Bayes rules.

2. Suppose that $X_1, \ldots, X_n$ are i.i.d. Exponential(θ), so the X's have density $p_\theta(x) = \theta e^{-\theta x} 1_{(0,\infty)}(x)$. with respect to Lebesgue measure on R, and that $\theta \sim \Gamma(\alpha, \beta)$:
$$\lambda(\theta) = \beta \frac{(\beta\theta)^{\alpha-1}}{\Gamma(\alpha)} \exp(-\beta\theta) 1_{[0,\infty)}(\theta).$$
(a) Find the Bayes rule $d_B(\underline{X})$ for estimation of θ with squared error loss $L(\theta, a) = |\theta - a|^2$. Find the Bayes rule $d_{Bw}(\underline{X})$ for estimation of θ with weighted squared error loss $L(\theta, a) = (\theta - a)^2/\theta$. Is the maximum likelihood estimator among either of these families of Bayes estimators?
(b) Are the Bayes estimators d_B and d_{Bw} consistent? What are the limit distributions of d_B and d_{Bw}? Compare them with the maximum likelihood estimator.
(c) Suppose that instead of the Gamma prior distribution, θ has the Pareto(θ_0, α) distribution with density λ given by
$$\lambda(\theta) = (\frac{\alpha}{\theta_0})(\frac{\theta_0}{\theta})^{\alpha+1} 1_{(\theta_0,\infty)}(\theta);$$
here $E(\theta) = \frac{\alpha}{\alpha-1}\theta_0$ where $\alpha > 1$ and $\theta_0 > 0$ are known. What can you say about the Bayes estimator for squared error loss with this prior? For what values of θ_0 is the Bayes rule consistent?

3. Suppose that $X_n \equiv X \sim \text{Multinomial}_k(n, \underline{\theta})$.
(a) Suppose that the prior distribution on θ is given by a Dirichlet distribution, Dirichlet($\underline{\alpha}$):
$$\lambda(\underline{\theta}) = \frac{\Gamma(\alpha_1 + \cdots + \alpha_k)}{\prod_{j=1}^k \Gamma(\alpha_j)} \theta_1^{\alpha_1 - 1} \cdots \theta_k^{\alpha_k - 1} 1_{[\underline{\theta}:\ \sum \theta_i = 1]}\,.$$

Verify the computation of the Bayes estimator for squared error loss given in example 4.3.4

(b) What is the posterior distribution for θ? Find the mode of the posterior distribution (along the lines of our computations of the MLE of the multinomial) and compare it with the MLE.

(c) Find a minimax estimator d_M of $\underline{\theta}$.

4. Find the limit distribution of the minimax estimator d_M in problem 3 (i.e. $\sqrt{n}(d_M(X_n) - p) \to_d$ "something" and find "something"). Is d_M a regular estimator of p?

5. **Optional bonus problem 1:** Let $\Theta = (0, \infty)$, $\mathbf{A} = [0, \infty)$, let X have the discrete distribution

$$p(x, \theta) = \binom{r + x - 1}{x} \theta^x (\theta + 1)^{-(r+x)}, \qquad x = 0, 1, 2, \ldots$$

where r is some known positive integer; this is the negative binomial distribution reparametrized so that $E_\theta X = r\theta$. Suppose that

$$L(\theta, a) = \frac{(\theta - a)^2}{\theta(\theta + 1)}.$$

(a) Show that the usual estimator, $d_0(X) = X/r$ is an equalizer rule; i.e. show that it has a risk function $R(\theta, d_0)$ which is constant in θ.

(b) Show that the usual estimator d_0 is generalized Bayes with respect to Lebesgue measure on $(0, \infty)$ provided $r > 1$. (A generalized Bayes rule is a rule that minimizes the posterior Bayes risk even when starting with an improper prior; see e.g. Ferguson, MS, page 50.) (What happens if $r = 1$?)

(c) Find Bayes decision rules with respect to the prior distributions $\Lambda_{\alpha,\beta}$ with densities

$$\lambda_{\alpha,\beta}(\theta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha - 1} (\theta + 1)^{-(\alpha+\beta)} 1_{(0,\infty)}(\theta),$$

the distribution of $\theta = Z/(1 - Z)$ where $Z \sim$ Beta(α, β).

(d) Show that $d(X) = X/(r + 1)$ is minimax. [Note that d_0 is not minimax, hence not admissible.]

6. **Optional bonus problem 2:** (Compare with Lehmann and Casella, TPE, Examples 5.1 and 5.2, pages 254-255.)

(a) Let $(X|\sigma^2) \sim N(0, \sigma^2)$. Show that the conjugate prior for σ^2 is the distribution of $1/Y$ where Y has a gamma distribution.

(b) Suppose that $(X|\theta, \kappa) \sim N(\theta, 1/\kappa)$, $(\theta|\kappa) \sim N(\mu, \tau/\kappa)$, and $\kappa \sim$ Gamma(α, β). Show that the posterior distribution of (θ, κ) has the same form as the prior.

(c) Find the marginal posterior distribution for θ in (b).

(d) If $X_1, \ldots, X_n$ are i.i.d. as X in (b), find the limiting distribution of the Bayes estimator of θ for squared-error loss.