Statistics 582, Problem Set 6

Wellner; 2/10/2010

Due: Wednesday, February 17, 2010.
Reminder: Midterm exam, Friday February 12.
Reading: Chapter 5, sections 7-8.

1. (a) Suppose that $X \sim F$ on $\mathbb{R}$. Let $\rho_p(x) \equiv x(p - 1_{(-\infty,0]}(x))$ for $p \in (0,1)$ and $x \in \mathbb{R}$. Show that $h_p(t) \equiv E_F \rho_p(X - t)$ is minimized by $F^{-1}(p)$.(Note that $\rho_{1/2}(x) = |x|/2$, so $h_{1/2}(t)$ is minimized by any median of F as we already know.)
(b) Suppose that $(X|\boldsymbol{\theta} = \theta) \sim P_\theta$ on $\mathbb{R}$ and that $\boldsymbol{\theta} \sim \Lambda$. Fix $p \in (0,1)$, and let $L(\theta, a) \equiv \rho_p(\theta - a)$. Show that the Bayes rule with respect to the prior Λ for this loss function is the $p-$th quantile of the posterior distribution $\Lambda(\theta|X)$.

2. Let $\mathcal{X} = \{0, 1\}$, $\mathcal{A} = \Theta = \{1, 2\}$, and assume that the losses are given by $L(1,1) = L(2,2) = 0$, $L(1,2) = a$, $L(2,1) = b$. Suppose that the statistician can observe either X or Y where

$$
\begin{aligned}
p_1(1) &= P_1(X = 1) = 2/3, & p_2(1) &= P_2(X = 1) = 1/2, \\
p_1^*(1) &= P_1(Y = 1) = 3/4, & p_2^*(1) &= P_2(Y = 1) = 1/2.
\end{aligned}
$$

Let $\underline{\lambda} = (\lambda, 1 - \lambda)$, $\lambda \in [0,1]$ be the prior distribution over Θ.
(a) Find the Bayes risk when X is observed, and similarly for Y.
(b) In the case $a = b$, $\lambda = 1/2$, would the statistician prefer to observe X or Y?
(c) For general $a \neq b$, $\lambda \in (0,1)$ would the statistician prefer to observe X or Y?

3. Consider Example 5.5.4 on pages 16 and 17 of the Chapter 5 notes.
(a) Show that the variance of $\hat{\psi}$ is given by

$$
Var(\hat{\psi}_n) = \frac{1}{n}\left\{\frac{1}{B}\sum_{j=1}^{B}\frac{\theta_j}{\xi_j} - \psi(\theta)^2\right\}.
$$

[Hint: use the formula $Var(Y) = EVar(Y|X) + Var[E(Y|X)]$ twice.]
(b) Use the result of (a) to show that

$$
Var(\hat{\psi}_n) \leq \frac{1}{n\delta}
$$

under the assumption that $\xi_j \geq \delta > 0$ for all $1 \leq j \leq B$.

4. **Optional bonus problem 1:** (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 $\kappa \sim \text{Gamma}(\alpha, \beta)$. 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.

5. **Optional bonus problem 2:** Suppose that $X \sim P_\theta$ for $\theta \in \Theta \subset R^k$ has well-defined Fisher information matrix $I(\theta)$ for θ. The *Jeffreys prior* distribution Λ_J has density $\lambda_J(\theta) = \det(I(\theta))^{1/2}$ with respect to Lebesgue measure on Θ. Note that Λ_J may not be a finite measure, and even if Λ_J is a finite measure, it may not have total mass 1. If a prior distribution is a finite measure, then call it a *proper prior distribution*, and correspondingly if it is not a finite measure, call it an *improper prior distribution*. If the resulting posterior distribution is a finite measure, call it a *proper posterior distribution*, and (by convention) normalize it to have total mass 1. See Lehmann and Casella, TPE, pages 230, 234, 287, 305.
(a) Suppose that $X \sim \text{Bernoulli}(\theta)$. Find the Jeffrey's prior density λ_J for θ. Is Λ_J a finite measure? If it is finite, what is $\Lambda_J((0,1))$? Find the corresponding posterior distribution of Θ starting with the Jeffrey's prior.
(b) Suppose that $X \sim \text{Poisson}(\theta)$ with $\theta \in (0, \infty)$. Find the Jeffrey's prior density λ_J for θ. Is Λ_J a finite measure? If it is finite, what is $\Lambda_J((0,\infty))$? Find the corresponding posterior distribution of Θ starting with the Jeffrey's prior. Is it ever a proper posterior distribution?
(c) Suppose that $X \sim \text{Geometric}(\theta)$, i.e. the number of trials until the first success in i.i.d. Bernoulli trials with probability θ of success for each trial – recall Chapter 1, section 1. Find the Jeffrey's prior density λ_J for θ. Is Λ_J a finite measure? If it is finite, what is $\Lambda_J((0,1))$? Find the corresponding posterior distribution of Θ starting with the Jeffrey's prior. If we observe $X_1, \ldots, X_n$ i.i.d. Geometric(θ), so that $\sum X_i \sim$ Negative Binomial(n, θ) is the posterior distribution "proper" for some n?
(d) Suppose that $X \sim \text{Weibull}(\theta)$ with $\theta = (\alpha, \beta) \in (0,\infty) \times (0,\infty)$ as in chapters 3 and 4. Find the Jeffrey's prior density λ_J for θ. Is Λ_J a finite measure? If it is finite, what is $\Lambda_J((0,\infty)^2)$? Find the corresponding posterior distribution of Θ starting with the Jeffrey's prior.