STAT 652: HW #3

due March 28, 2006

Instructions: Please present your solutions in order. As always, use paper which is approximately 8.5 inches by 11 inches, present tidy and easy-to-follow solutions, draw boxes around your final answers (but don't just give answers without supporting work), and staple all sheets together in the upper left hand corner. I don't like cover sheets, executive summaries, folders, binders, or paper clips.

I'll be willing to offer a *limited* amount of advice on 1(a) and 3(a), but not the other parts. There are 21 points in all, but I'll truncate scores at 20 and make this a 20 point assignment.

1) Suppose that we observe iid random variables $X_1, X_2, \ldots, X_n$ having pmf

$$p_X(x|\theta) = \theta(1-\theta)^{x-1} I_{\{1,2,3,\ldots\}}(x),$$

where $\theta \in \Theta = (0,1)$. Consider the prior density $\pi(\theta) = I_{(0,1)}(\theta)$.

(a) (6 points) Find the posterior density, and give the Bayes estimate of θ based on the squared-error loss function. (Note: Be sure to give both items that are requested — draw a box around (or highlight) each on your submitted solutions.)

(b) (3 points) Consider the squared-error loss function and give the Bayes estimate of $\tau(\theta) = (1-\theta)/\theta^2$. (The estimate is $E[\tau(\theta)|\mathbf{X} = \mathbf{x}]$ and for large n will be close to the mle of $\tau(\theta)$.)

The following parts may be useful to you in checking your work for this problem, but they will not be graded. There is no need to submit solutions for these parts.

(c) (0 points) The mode of the posterior density (the value of θ which maximizes $p(\theta|\mathbf{x})$) is sometimes called the *generalized maximum likelihood estimate*. (Note that the generalized mle doesn't depend on the choice of the loss function. We can employ this estimate even if we don't identify an appropriate loss function. Of course, if you can identify an appropriate loss function, then it makes sense to use it by finding the estimator which minimizes the Bayes risk. In part (a), we consider a squared-error loss function, and the corresponding Bayes estimator is the mean of the posterior distribution instead of its mode.) Use the posterior density requested in part (a) and find the generalized mle. (The estimate is $\frac{1}{\bar{x}}$, which is different from, but simliar to, the estimate from part (a) — for large sample sizes the two estimates will be very close to one another.)

(d) (0 points) Show that the Bayes estimators corresponding to parts (a) and (b) are both consistent estimators.

2) (6 points) Consider the N$(\theta, 1)$ random variable X, the prior density $\pi(\theta) = (2\pi)^{-1/2}\exp(-\theta^2/2)$ (for all $\theta \in (-\infty, \infty)$), and the loss function

$$l(\theta, \hat{\theta}) = \exp\left(\frac{\hat{\theta} - \theta}{2}\right) - \frac{(\hat{\theta} - \theta)}{2} - 1.$$

(Note: This is a special case of Zellner's LINEX loss function. It is asymmetric in that it penalizes overestimation differently than underestimation.) Find the posterior density, and give the Bayes estimate of θ based on the LINEX loss function given above. (Note: Be sure to give both items that are requested — draw a box around (or highlight) each on your submitted solutions.) *Note that the estimate is to be based on a single observation, x, and not a sample of size n.* (Note: With this loss function we do not have that the Bayes estimate of θ is $E(\theta \,|\, X = x)$. Instead, the Bayes estimate is the value of a which minimizes

$$\int_\Theta l(\theta, a)\, p(\theta|x)\, d\theta \;=\; \int_\Theta e^{(a-\theta)/2} p(\theta|x)\, d\theta - \frac{1}{2}a \int_\Theta p(\theta|x)\, d\theta + \frac{1}{2}\int_\Theta \theta p(\theta|x)\, d\theta - \int_\Theta p(\theta|x)\, d\theta.$$

While evaluating the first of these integrals may be a bit messy, the other ones should be relatively easy if you keep in mind that the $p(\theta|x)$ is a valid pdf.)

3) Suppose that we have a single observation u from a uniform $(0, \theta)$ distribution, with $\theta > 1$, and that the prior distribution for θ has the pdf

$$\pi(\theta) = \theta^{-2} I_{(1,\infty)}(\theta).$$

(a) (3 points) Give the Bayes estimate of θ based on the squared-error loss function. Make sure your estimate is correct no matter what value u is. (*Hint*: When obtaining the posterior pdf, you might want to reexpress an indicator function of the form $I_{A(\theta)}(u)$, where $A(\theta)$ represnts an interval having θ as an endpoint, as an indicator function of the form $I_{B(u)}(\theta)$, where $B(u)$ represents an interval having u as an endpoint, since one should view the posterior pdf as a function of θ. There will be another indicator fuction of the form $I_C(\theta)$ involved, where C represnts an interval having endpoints that do not involve θ or u. I suggest that you consider a product of two indicator functions, both having argument θ, and then express this product as a single indicator function.)

(b) (3 points) Give the Bayes estimate of θ based on the absolute error loss function. (Note: Make use of the fact that for this loss function, the Bayes estimate is just the median of the posterior distribution.) Make sure your estimate is correct no matter what value u is.