Study question set #2 for STA 3000, Spring 2014

These are for study only, not to hand in for credit. They don't cover all the topics to be covered on the second test, just mostly those not covered by the assignments.

Question 1: Data $X_1, \ldots, X_n \in (0,1)$ are modelled using a parameter $\theta \in (0,1)$, with $X_1, \ldots, X_n$ i.i.d. given θ, each with density

$$f(x|\theta) \; = \; \begin{cases} (1-\theta)/2 & \text{if } x \le 1/2 \\ \theta/2 & \text{if } x > 1/2 \end{cases}$$

We wish to estimate θ with squared error loss. The estimator $\delta_0(x) = (1/n)\sum x_i$ has been proposed. Using the Rao-Blackwell theorem, find a better estimator.

Question 2: Suppose we model binary observations $Y_1, Y_2, \ldots, Y_n$ as being i.i.d. given $\theta \in (0,1)$, with $P(Y_i = 1 \,|\, \theta) \; = \; \theta$. Suppose that our actual observations are that $y_1 = y_2 = \cdots = y_n = 0$.

A) Suppose we do a frequentist hypothesis test of $H_0: \; \theta = 1/2$ versus $H_1: \; \theta \neq 1/2$ in the standard way. What will be the p-value computed from this dataset (as a function of n)?

B) Suppose we compare two Bayesian models, differing only in their prior for θ, with the prior for model H_0 being that θ is 1/2 with probability one, and the prior for θ in model H_1 being uniform over $(0,1)$. Suppose that our prior probabilities for H_0 and H_1 are equal (both 1/2). What will be the posterior probability of model H_0 (as a function of n)?

Question 3: Suppose we model Y_1 and Y_2 as being i.i.d. given $\theta \in (0,1)$, with $P(Y_i = 1|\theta) = \theta$. We first observe y_1. If y_1 is 0, we stop; otherwise we observe y_2 as well.

A) Suppose we estimate θ by the average of the y_i that we observed (ie, either y_1 or $(y_1 + y_2)/2$). Is this estimate unbiased?

B) Supoose we use a uniform prior on $(0,1)$ for θ, and estimate θ by its posterior mean given the observed y_i. What is the prior expectation of this estimate (ie, the expected value of the estimate averaging over the prior for θ and the distribution of the observed data)?

Question 4: Suppose we model positive real observations $Y_1, Y_2, \ldots, Y_n$ as being i.i.d. given a positive real parameter θ, with each observation having density $f(y|\theta) = \theta \exp(-y\theta)$.

A) What will Jeffreys' prior be for θ in this model?

B) Suppose we switch to parameterizing the model by $\mu = 1/\theta$. Derive Jeffreys' prior in this parameterization, and confirm that it is the same as would be obtained from the prior in (A) after transforming the prior density in the usual way.