EXERICISE ON BAYESIAN AND FREQUENTIST INFERENCE

CHRISTOPHER A. SIMS
PRINCETON UNIVERSITY
SIMS@PRINCETON.EDU

(1) In our rain gear example, suppose the list of feasible ϕ's (i.e. the whole Φ set) consisted of the following list of points:

	$\phi(\omega_r)$	$\phi(\omega_c)$
1	0.95	0.62
2	0.23	0.79
3	0.61	0.92
4	0.49	0.74
5	0.89	0.18
6	0.76	0.41
7	0.46	0.94
8	0.02	0.92
9	0.82	0.41
10	0.44	0.89

Which are admissible? Which are Bayes? Which, if any, are admissible but not Bayes? Which, if any, are Bayes but not admissible? Which, if any, are admissible but would become inadmissible if we allowed randomized decision rules? [This question can be answered quickly with a graph (and maybe a ruler) if you get the computer to plot points. In R you would use `plot()` followed by `text()` (which labels the points, sequentially by default.]

(2) Suppose we have i.i.d. observations on $\{x_t,\ t = 1, \ldots, T\}$, with each x_t distributed as $N(\mu, \mu^2)$. That is, μ is both the mean and the standard deviation of the observations. It is known that $\mu > 0$. This exercise compares Bayesian with frequentist methods of constructing point estimates and interval estimates for μ.
   (a) Show that $\sum x_t/T$ and $\sum x_t^2$ form a two-dimensional sufficient statistic here.
   (b) Show that $\bar{x} = \sum x_t/T$ is an unbiased and consistent estimator for μ.
   (c) Show that $s^2 = \sum(x_t - \bar{x})^2/T$ is an unbiased estimator for μ^2 and that its square root is a consistent estimate of μ.

Date: September 17, 2006.

(d) Show that $\bar{x}/\mu$ and s^2/μ^2 are what is known as **pivotal quantities**, or just plain **pivots**, meaning that they each have a distribution that does not depend on the unknown parameter μ. Show also that the two are independent for each μ.

(e) Derive the form of three confidence intervals, based on $\bar{x}/\mu$, on s^2/μ^2, and on $\sum x_t^2/\mu^2$ (which is also pivotal).

(f) For each of the following samples, find $\bar{x}$, $\sqrt{s^2}$, the maximum likelihood estimate of μ, the flat-prior posterior mean of μ, and the posterior mean of μ when the prior is proportional to $\mu^{-2}\exp(-1/(10\mu))$. The posterior means probably require numerical integration. There are functions in Matlab and R that do numerical integration, or it is fairly easy to code this yourself in a couple of lines.
Also find for each sample the 95% confidence intervals you derived above and 95% HPD (highest posterior density) regions under the flat prior and the proper prior. If $p(\mu \mid \vec{x})$ is the posterior density function, the 95% HPD region for μ is a set of the form $\{\mu \mid p(\mu \mid \vec{x}) > \bar{p}\}$ for a $\bar{p}$ such that the set's posterior probability is .95. You will need the computer to find it numerically.
The samples:
   (i) $\{-5, 0, 5\}$
   (ii) $\{0, 1, 2\}$
   (iii) $\{5, 5.1, 5.2\}$
   (iv) $\{-2, -2, -2.1\}$

(g) In deciding which of the three confidence intervals to use, would it make sense in a given sample to pick whichever is smallest, on the grounds that that is the one that is giving the most precise information? Why or why not?

(h) How should empty confidence intervals be interpreted? Should there be a big difference between the interpretation of an empty confidence interval and of a very short (nearly empty?) confidence interval?

(i) Show that the Bayesian posterior mean and HPD region can't be computed for the flat prior, if the sample size is one.