STA 3000, 2008/2009 — Assignment #4

Due January 15, at start of lecture. Worth 8% of the course grade.

This assignment is to be done by each student individually. You may discuss it in general terms with other students, but the work you hand in should be your own. In particular, you should not leave any discussion of this assignment with any written notes or other recordings, nor receive any written or other material from anyone (except your instructor) by other means such as email.

For all questions, show both the final answer and how you obtained it.

Question 1: A cosmic ray detector records the number of cosmic rays detected in successive one-minute periods. These observations over n minutes are $X_1, \ldots, X_n$. Suppose cosmic rays arrive independently, with uniform intensity over time. These counts should thus be IID from a Poisson distribution with some unknown mean, if the detector is functioning correctly.

Suppose, however, that you suspect that the software controlling the detector has a bug, which has the effect that the count recorded for some one-minute periods is not the correct count, but is rather a duplicate of a count recorded in some earlier one-minute period — so that X_i is exactly equal to X_j for some $j < i$.

Duplicate counts could also occur just by chance, of course, so it may not be obvious whether the software is really defective. For this question, you are to devise a pure significance test of the null hypothesis that the counts are correctly recorded, and are IID from a Poisson distribution (with mean unknown). This test should be sensitive to the sort of software flaw described above. You may assume that the duplicate counts occur fairly rarely (less than 10% of the time, maybe much less). There may be a tendency of the software, if it is actually flawed, to duplicate a fairly recent count, rather than a count from many minutes in the past. Although the mean of the Poisson distribution for counts is unknown, you should assume that it is neither very large nor very small (eg, in the the range 2 to 100). The number of counts, n, is also at least fairly large (greater than 20).

Hints: The method described by Cox and Hinkley in Section 3.3 of their book may be useful. For this method, you find a sufficient statistic for the model assumed by the null hypothesis, and then base a test statistic on the conditional distribution of the full data given the sufficient statistic, which would not depend on the unknown Poisson mean parameter. Do not worry about computational difficulties, as long as the computations are theoretically possible. Note that since the data is discrete, a test that produces exactly a $U(0, 1)$ distribution for the p-value if the null hypothesis is true may not be possible — very close to $U(0, 1)$ is good enough.

Question 2: Consider a model with $X_1, \ldots, X_n | \theta \sim U(\theta - 1, \theta + 1)$. One idea on how to do a "Bayesian hypothesis test" of $\theta = 0$ is to use a prior in which $\theta = 0$ has a positive probability, with a fairly broad prior otherwise. Consider such a prior, with $\theta \sim p_0\delta_0 + (1 - p_0)U(-w, +w)$, where δ_0 is a point mass at zero. The values of p_0 and w are known, but below you should work out the results for any $p_0 \in (0, 1)$ and any positive real w.

a) Find the posterior distribution of θ given $X_1, \ldots, X_n$, expressing it as a posterior density with respect to a suitable dominating measure, and find the posterior probability that $\theta = 0$.

b) Suppose that we must "accept" or "reject" $\theta = 0$, with loss of c (a positive real) if we reject when in fact $\theta = 0$, loss of 1 if we accept $\theta = 0$ when actually $\theta \neq 0$, and loss of 0 if we reject when $\theta \neq 0$ or accept when $\theta = 0$. Find the formal Bayes rule for this decision.

c) Find the probability that the decision rule above will reject if $\theta = 0$, as a function of p_0, w, and c.

c) Discuss the results you found above. In particular, by adjusting p_0, w, and c, can you make the accept/reject decision in (b) be what you would regard as a sensible frequentist test of $\theta = 0$ at some given significance level (eg, 0.05)? How does the behaviour of the formal Bayes rule change as w changes? Do you think this behaviour is reasonable?