EXERCISE ON MODEL CHECKING, DUE TUESDAY NOVEMBER 9

(1) Consider the example we discussed in class, where the two competing models are
A: $X_t \sim N(0,\sigma^2)$ with σ^2 unknown and
B: $X_t \sim N(\mu,1)$ with μ unknown.
In both models we assume the observations are i.i.d. over $t = 1,\ldots,10$. Use an exponential prior (pdf $\alpha e^{-\alpha/\sigma^2} d(1/\sigma^2)$) on $1/\sigma^2$ and a $N(0,\nu^2)$ prior on μ.
  (a) Setting $\nu = 100$, $\alpha = .01$, (i.e. very flat priors), display contour maps of the log odds ratio for A over B and of the posterior probability of of model A. (Of course these two quantities are monotonically related, but their contour maps, with the default equi-spaced contour lines of Matlab, R, or S, will look quite different.) Use as axes for the contour maps $\bar{X}$, the sample mean of the observed X's, and $s^2 = \sum(X_t - \bar{X})^2/9$, the usual unbiased estimate of the variance. Show the same two plots also for the case $\alpha = \nu = 1$. [Hint: Matlab, R and S all have `contour` functions that produce contour plots automatically. In R or S it is particularly convenient to combine this with the `outer` function. If you set `xb <- seq(-4,4,.1)`, `s2 <- seq(0,.05,4)`, and define a function (say `LR <- function(x,s){...}`) that takes two arguments and produces the value you want contours for, you then get the contour plot from `contour(xb,s2,outer(xb,s2,FUN=LR))`. Note that if you have not forgotten any factors in your formula, the posterior probability plot should show levels ranging from near one to near zero.]
  (b) A natural decision rule, if we had to choose between the two models, would be to choose model A when its posterior probability exceeds .5 (or, equivalently, when the log odds ratio exceeds 0). We could then interpret this as a test of A as the null hypothesis, with the rejection region being samples in which the posterior probability of A is less than .5. Find the size of this test by drawing artificial samples from Model A with a few well-chosen values of σ^2 and using them to estimate rejection probabilities. Your calculation of size this way does not have to be extremely precise. Note that we might expect rejection probabilities under model A to be greatest when $\sigma^2 \doteq 1$, i.e. in a version of model A that is actually close to being consistent with model B. Use just the flatter prior.
  (c) From your calculations in (1b), determine whether it appears that this test is unbiased.

(d) We could consider the same test reversed, in which B is treated as the null and rejected when the posterior of A exceeds .5. What would be its size? would it be unbiased?

(e) Suppose

(i) H_0 and H_1 are two subsets of the parameter space Θ that contain at least one point in common, with $H1 \cup H2 = \Theta$.

(ii) $\delta(X)$ is a rule for rejecting H_0 as a function of the data that is unbiased as a test of H_0 as null hypothesis.

Show that $\bar{\delta}(X)$, a decision rule that rejects H_1 whenever $\delta(X)$ accepts H_0, is an unbiased test of H_1 and that at least one of $\delta, \bar{\delta}$ has size .5 or larger.

(2) Suppose we are estimating a SNLM with two coefficient parameters, β_1, β_2, and an unknown residual variance σ^2. We proceed by first estimating the regression with both coefficient parameters free, then using a a t test to check whether β_2 is significantly different from zero, retaining this two-parameter estimate if the test rejects the null of $\beta_2 = 0$, but using instead an estimate of β_1 from the univariate regression with β_2 constrained to zero if the $\beta_2 = 0$ null hypothesis is accepted.

(a) Show that if in fact $\beta_2 = 0$, if the test for $\beta_2 = 0$ is one-tailed (rejecting only for β_2 "significantly greater than 0", not when it is "significantly less than zero"), and if the two variables in the X matrix are positively correlated in the sample, the estimate of β_1 resulting from this two-step procedure is biased.

(b) Show that if $\beta_2 \neq 0$, then even if the first-stage test is two-sided, the resulting estimator for β_2 is biased. (This estimator is the first-stage estimator of β_2 when the $\beta_2 = 0$ null is rejected, and 0 if the null is accepted.)

(c) Suppose that in a particular sample the p-value of the test statistic for $\beta_2 = 0$ is .005, and that the estimate we are using is the two-coefficient regression estimate, because we set the size of our test at some $\alpha > .005$. Does the size of the bias depend on whether we specified in advance $\alpha = .05$ or $\alpha = .01$? Would it make sense to try to correct the estimate for this bias? Would your answers to these questions be different if the p-value were .15, so that we were using the 1-coefficient OLS estimate for either value of α?

[Hints: Note that to show bias, you only have to show that the estimator's expected value is not the true value for some single value of the true parameter vector. You can answer these questions either analytically, or by programming the computer to give you the answer by Monte Carlo methods. If you take the latter course, to keep the results easy to grade, assume

$$X'X = \begin{bmatrix} 5 & 4 \\ 4 & 5 \end{bmatrix}, \quad T = 1000, \quad \sigma^2 = 1\,.$$

With this large number of observations T, the t and F distributions ordinarily used in constructing distributions of SNLM estimators and test statistics become essentially exactly Normal and chi-squared distributions. In other words the model becomes essentially equivalent to one in which σ^2 is known. Of course you can assume this σ^2-known simplification also if you attack the problem analytically.