

18.443 Problem Set 5, due Friday, March 21, 2008

Here are some comments on “ p -values,” a common notion in statistics that Rice mentions, and which we’ve already seen in the output of the Shapiro-Wilk test. If a hypothesis H_0 will be rejected for large values of a statistic V such as X^2 or W , and a value V_{obs} is observed, then the p -value is the probability that V is greater or equal V_{obs} , assuming H_0 .

If V has approximately a chi-squared distribution, then you can determine whether the p -value is less than α , or larger than $1 - \alpha$, for a few usual values of α , from the chi-squared tables.

By the same computational methods that have been used to compute the numbers in the tables (which was done some decades before computers were invented, on very slow desk calculator machines) it is also possible to compute p -values for any V_{obs} , and Rice mentions some such p -values. But you aren’t asked to do such computations, just to answer questions about α and $1 - \alpha$ comparisons, and the tables are enough for that.

1. (This relates to the handout “Mean-square errors of estimators...”.) For the three possible choices of c_n in estimating sample variance, namely $1/(n - 1)$, $1/n$, and the newer Yatracos choice $c'_n = (n + 2)/[n(n + 1)]$,

(a) clearly $1/n < 1/(n - 1)$. Where does c'_n fit into the ordering?

Recall that $a_n \sim b_n$ means that $a_n/b_n \rightarrow 1$ as $n \rightarrow \infty$. For each of the following D_n , find a constant C and power k such that $D_n \sim C/n^k$ as $n \rightarrow \infty$.

(b) $D_n = \frac{1}{n-1} - \frac{1}{n}$; (c) $D_n = \frac{1}{n-1} - c'_n$;

(d) Let the true variance of the distribution of X_j be σ^2 and D_n the bias of Yatracos’ estimator $c'_n \sum_{j=1}^n (X_j - \bar{X})^2$ as an estimator of σ^2 .

2. Some random numbers were generated in R as follows. They are independent and normally distributed, all with mean 5.7, but with a variance chosen at random: with probability 0.9 the variance is 1, and with probability 0.1 it is 36, so that the standard deviation is 6. Let’s consider some numbers actually output by R and see if we can determine which variance they were generated from. For each of the following numbers, find the likelihood ratio that they came from $N(5.7, 1)$ relative to coming from $N(5.7, 36)$. Also find the posterior probability in each case. If answers are very large or small, give them in scientific notation $r \cdot 10^m$ where m is an integer that may be positive, negative or 0 and $1 \leq r < 10$, giving r to three significant digits.

(a) $X = 7.43$ (b) $X = 10.46$ (c) $X = 17.28$.

3. In Buffon’s needle problem, a needle of length L is thrown at random onto ruled paper with lines at distance D apart. If $L < D$, the probability that the needle hits a line is $2L/(\pi D)$. Lazzarini in 1901 reported on an experiment with $L/D = 5/6$, where the needle was thrown 3408 times and it hit the line 1808 times.

Usually χ^2 tests are one-sided, and the hypothesis is rejected for large values of the statistic X^2 . But, in this case, do a two-sided test for random sampling with the given hitting probability, where we’d reject the hypothesis if X^2 is too large (the hitting probability is wrong) or if X^2 is too small (the sampling may have been non-random), at the $\alpha = 0.05$ level.

4. Consider the class of all normal distributions $N(\mu, 1)$, so that H_1 is the set of all $(\mu, 1)$ for any real μ and $\sigma = \sigma^2 = 1$. Let H_0 be the subset of H_1 where μ has a fixed value $\mu = \mu_0$. Show that in this case the Wilks statistic $W = -2 \log \Lambda$ (defined in the handout in general, as opposed to the multinomial case as in Rice) has exactly a $\chi^2(1)$ distribution for all n , not only asymptotically as $n \rightarrow \infty$.

5. Rice, §9.11 Problem 42(a), (b).