STA 3000F (Fall, 2013)

Homework 4 due December 6

1. CH 4.7 Assume $y_1, \dots, y_n$ is an independent sample from the exponential distribution with density $\theta \exp(-\theta y), y \geq 0, \theta > 0$.

   (a) Obtain the uniformly most powerful test of $H_0 : \theta = \theta_0$ against alternatives $\theta < \theta_0$, and derive the power function of the test.

   (b) To reduce dependence on possible outliers, it is proposed to reject the largest observation. Derive the uniformly most powerful test from the remaining observations, and show that the loss of power corresponds exactly to the replacement of n by $n-1$.
   The *Renyi representation* of the order statistics $(Y_{(1)}, \dots, Y_{(n)})$ is

$$Y_{(r)} \stackrel{d}{=} \frac{1}{\theta} \sum_{j=1}^{r} \frac{E_j}{n+1-j},$$

   where $E_1, \dots, E_n$ are independent exponential random variables with rate 1. (SM, §2.3).

   (a) By the Neyman-Pearson lemma, the most powerful test of $H_0 : \theta = \theta_0$ against $H_1 : \theta - \theta_1$ has rejection region

$$\mathcal{R} = \{y : f(y; \theta_1)/f(y; \theta_0) > k\},$$

   which in this case simplifies to $(\theta_1/\theta_0)^n \exp(\theta_0 - \theta_1)\Sigma y_i > k$, and this in turn simplifies to $t > k'$ where $t(y) = \Sigma y_i$, because with $\theta_1 < \theta_0$, the likelihood ratio is monotone increasing in t. Since this region is the same for all $\theta_1 < \theta_0$, the test with this rejection region is uniformly most powerful against the composite alternate $H_1' : \theta < \theta_0$. The cutoff k' is determined from the null distribution of T, which is $\Gamma(n, \theta_0)$, and the power function is

$$\beta(\theta) = \Pr(T > k'; \theta) = \Pr(\Gamma(n, \theta) > k'),$$

   with these expressions having equivalent versions based on the χ^2_{2n} distribution. Nearly everyone got this correct, with occasional $t < k'$ critical regions. But since $E_\theta(T) = 1/\theta$, small values of θ go with large values of T. Note that the likelihood ratio of the NP Lemma is not the same as the likelihood ratio statistic $w(\theta) = 2\log\{L(\hat{\theta})/L(\theta)\}$, and the NP Lemma does not tell us anything about the power of tests with critical regions determined by $w(\theta)$. (Confusingly.)

(b) As noted in the revisions to the HW, the test statistic $t'(Y) = \Sigma_{i=1}^{n-1} Y_{(i)} + Y_{(n-1)}$, which replaces the largest observation by the second largest, is the one that has the property stated in the question, i.e. T' has the same distribution as a sum of $n-1$ independent exponential random variables with rate parameter θ, $\Gamma(n-1, \theta)$. Everywhere here the θ in the gamma density is the rate parameter θ (= 1/scale, in `R`-speak). Evgeny gets a gold star for deriving the joint distribution of the first $n-1$ order statistics by starting from the joint distribution of all the order statistics ($n!\Pi f(y_{(i)})$), and marginalizing over the nth; this calculation is do-able, and does lead to $T' > k$ as the critical regions for the uniformly most powerful test of H_0 vs H_1'.

2. The p-value function is not the same as the power function: the latter depends on both the null value and the alternative value(s) of the parameter. From the Neyman-Pearson lemma, we know that for $H_0 : \theta = \theta_0$ vs. $H_A : \theta = \theta_A$, the most powerful test based on a sample $y = (y_1, \ldots, y_n)$ from $f(y; \theta)$ is based on the likelihood ratio

$$\frac{L(\theta_A; y)}{L(\theta_0; y)}.$$

(a) For θ a scalar parameter, let $\theta_A = \theta_0 + \delta$, for fixed $\delta > 0$. Use Taylor series expansions to show that the locally most powerful test, where 'local' means δ near 0, is a function of $U(\theta_0; y)$, the score statistic. slightly amended

(b) Compare the local power of the test based on U to the power of the test based on the likelihood root $r = \pm\sqrt{[2\{\ell(\hat{\theta}) - \ell(\theta)\}]}$, for independent, identically distributed observations from the Cauchy distribution with location parameter θ: $f(y_i; \theta) = [\pi\{1 + (y_i - \theta)^2\}]^{-1}$.

This can only be done numerically, as far as I can tell. For example, simulate a sample of size 10 from the standard Cauchy with `rcauchy(10)`.

UMP Invariant tests

A composite null hypothesis $H_0 : \theta \in \Theta_0$ does not specify the distribution of Y under H_0, so some means of reduction to a test statistic with a known distribution under H_0 is needed. In some classes of models a group of transformations can be found which leaves both the null and alternative hypotheses unchanged, and a statistic exists which is a maximal invariant for the group. The resulting test, constructed using the Neyman-Pearson lemma on the marginal density of this invariant, is called most powerful invariant. Some notes on the classical theory of testing are posted, see also CH Ch. 5.2,3 and SM Ch. 7.3.3. The usual textbook example of this is the t-statistic for comparing two normal means with common unknown variance, but many tests in multivariate analysis are also motivated by considerations of invariance.

The next question is based on CH Example 5.17.

3. Assume $(x_1, y_1), \ldots, (x_n, y_n)$ are independent bivariate normal distributions with mean vector (μ_1, μ_2) and covariance matrix the identity. Suppose we are interested in testing $H_0 : \mu_1^2 + \mu_2^2 = 1$, against the alternative $H_1 : \mu_1^2 + \mu_2^2 > 1$.

   (a) Show H_0 and H_1 are invariant under the rotation group on the parameter space, which has elements

   $$g \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix} = \begin{bmatrix} \sin a & \cos a \\ \cos a & -\sin a \end{bmatrix} \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}, \quad 0 < a \leq 2\pi.$$

   (b) Show that $t = \bar{x}^2 + \bar{y}^2$ is a maximal invariant statistic under the rotation group on the sample space of $(\bar{x}, \bar{y})$, the minimal sufficient statistic, that corresponds to the group on the parameter space, and that the distribution of t depends only on $\phi = \mu_1^2 + \mu_2^2$.

   (c) Find the most powerful test of H_0 based on the marginal distribution of t, and describe the rejection region $\mathcal{R}$ for a test of size α.

   (d) (not required, but potentially of interest): In the more usual case where the variance is σ^2 and unknown, the maximal invariant is

   $$t' = \frac{t}{\Sigma(x_i - \bar{x})^2 + \Sigma(y_i - \bar{y})^2},$$

   and $u = n(2n-2)t$ follows a non-central F-distribution with 2 and $n-2$ degrees of freedom, and non centrality parameter $n\phi/\sigma^2$. In time series this is the Schuster test for the amplitude of a given component of the periodogram.