S&DS 242/542: Homework 5

Due Wednesday, February 19, at 1PM

1. Sign test and its local power. Let $X_1, \ldots, X_n \stackrel{IID}{\sim} f$ be distributed according to an unknown PDF f on the real line. Consider testing

 $H_0: f$ has median 0 vs. $H_1: f$ has median greater than 0.

Let S be the number of positive values among X_1, \ldots, X_n , and consider the test statistic

$$T = \sqrt{\frac{4}{n}} \Big(S - \frac{n}{2} \Big).$$

The test of H_0 vs. H_1 based on T (or equivalently, based on S) is called the sign test.

(a) Apply the Central Limit Theorem to provide a normal approximation for the distribution of T under H_0 , and explain how you would use this approximation to test H_0 vs. H_1 at a given significance level α .

(b) Consider the specific alternative $H'_1: X_1, \ldots, X_n \stackrel{IID}{\sim} \mathcal{N}(\frac{h}{\sqrt{n}}, 1)$ for some h > 0. Show that

$$\mathbb{P}_{H_1'}[X_i > 0] = \Phi\left(\frac{h}{\sqrt{n}}\right)$$

where Φ is the CDF of $\mathcal{N}(0, 1)$. Assuming that h is a small fixed value and that the sample size n is large, explain why

$$\mathbb{P}_{H_1'}[X_i > 0] \approx \frac{1}{2} + \frac{1}{\sqrt{2\pi}} \cdot \frac{h}{\sqrt{n}}$$

(c) Using your result in part (b), derive a normal approximation for the distribution of T under H'_1 that depends only on h and not on n. A heuristic argument is fine, and you do not need to formalize convergence in distribution.

(d) Apply your normal approximation in part (c) to show that the power of your test in part (a) against this alternative H'_1 is approximately given by $\Phi(\sqrt{\frac{2}{\pi}} \cdot h - z^{(\alpha)})$ where $z^{(\alpha)}$ is the upper- α point of the standard normal distribution.

2. Power comparisons. For testing

$$H_0: X_1, \dots, X_n \stackrel{IID}{\sim} \mathcal{N}(0, 1)$$
 vs. $H_1: X_1, \dots, X_n \stackrel{IID}{\sim} \mathcal{N}(\mu, 1)$ where $\mu > 0$

consider two different tests at significance level $\alpha = 0.05$: the sign test from Problem 1, and the standard one-sample t-test from lecture. Let us fix the sample size n = 100.

(a) Perform a simulation of the Type I error probability and power of these two tests, in the following way: For each value $\mu \in \{0, 0.1, 0.2, 0.3, 0.4\}$ (where $\mu = 0$ corresponds to the null hypothesis H_0) simulate 10,000 random samples of $X_1, \ldots, X_{100} \stackrel{IID}{\sim} \mathcal{N}(\mu, 1)$, carry out both tests for each sample, and record whether each test accepts or rejects H_0 .

Report the simulated probability of Type I error and simulated power against each alternative $\mu \in \{0.1, 0.2, 0.3, 0.4\}$ for both tests.

[The following commands may be helpful in R:

```
qnorm(0.95)
qt(0.95,df=99)
```

give the upper-0.05 points of the standard normal distribution and t-distribution with 99 degrees of freedom. To decide if a test accepts or rejects H_0 , you may use an if-else statement such as

```
if (T > qnorm(0.95)) {
   reject = 1
} else {
   reject = 0
}
```

The commands

mean(X)
sd(X)
length(which(X>0))

compute the sample mean, standard deviation, and number of positive observations in X.]

(b) The power of the one-sample z-test for this testing problem was derived in lecture to be $\Phi(\sqrt{n\mu} - z^{(\alpha)})$. How does your simulated power of the t-test compare with this power? Substituting $\mu = \frac{h}{\sqrt{n}}$, the power of the sign test was derived in Problem 1(d) to be approximately $\Phi(\sqrt{\frac{2}{\pi}} \cdot \sqrt{n\mu} - z^{(\alpha)})$. How does your simulated power of the sign test compare with this approximation? How does it compare with the above powers of the z-test and t-test?

3. **FWER vs. FDR.** (a) Consider any multiple testing procedure for n null hypotheses $H_0^{(1)}, \ldots, H_0^{(n)}$ that controls the familywise error rate (FWER) at a level $\alpha \in (0, 1)$. Does

this procedure necessarily control the false discovery rate (FDR) at level α ? Explain why or why not.

(b) Suppose the Bonferroni method applied to control FWER $\leq \alpha$ rejects a subset of these null hypotheses $H_0^{(1)}, \ldots, H_0^{(n)}$. Would these hypotheses necessarily be rejected by the Benjamini-Hochberg procedure applied to control FDR $\leq \alpha$? Explain why or why not.

4. Improving on Bonferroni for independent tests.

(a) Let P_1, \ldots, P_n be the p-values from n different hypothesis tests. Suppose these p-values are independent, and $P_i \sim \text{Uniform}(0,1)$ if the i^{th} null hypothesis is true. Consider the multiple testing procedure which rejects those null hypotheses where $P_i \leq t$. Show that if there are n_0 true null hypotheses, then for any $t \in (0, 1)$,

 $\mathbb{P}[\text{reject any true null hypothesis}] = 1 - (1 - t)^{n_0}.$

(b) Show that if we choose $t = 1 - (1 - \alpha)^{1/n}$, then this controls the FWER at level α . Would this procedure reject fewer or more hypotheses than the Bonferroni method which uses $t = \alpha/n$? How does it differ from the Bonferroni method in its assumptions?