STAT 713 hw 6

Asymptotic distributions of estimators, size and power of some classical tests

1. Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Poisson}(\lambda)$.

   (a) Find the ARE of $\bar{X}_n$ with respect to $\hat{\sigma}_n^2 = n^{-1}\sum_{i=1}^n (X_i - \bar{X}_n)^2$ for estimating λ.

   The central limit theorem immediately gives $\sqrt{n}(\bar{X}_n - \lambda) \xrightarrow{\text{D}} \text{Normal}(0, \lambda)$. Noting that $\hat{\sigma}_n^2 = \hat{m}_2 - \hat{m}_1^2$, we can use

$$\sqrt{n}((\hat{m}_1, \hat{m}_2)^T - (m_1, m_2)^T)^T \xrightarrow{\text{D}} \text{Normal}((0,0)^T, (m_{i+j} - m_i m_j)_{1 \le i,j \le 2})$$

   as $n \to \infty$ with the delta method for $g(m_1, m_2) = m_2 - m_1^2$ to obtain

$$\sqrt{n}(\hat{\sigma}_n^2 - \sigma^2) \xrightarrow{\text{D}} \text{Normal}(0, \vartheta),$$

   where

$$\begin{aligned} \vartheta &= (-2m_1)^2(m_2 - m_1^2) + 2(-2m_1)(m_3 - m_1 m_2) + m_4 - m_2^2 \\ &= m_4 - m_2^2 + 8m_1^2 m_2 - 4m_1^4 - 4m_1 m_3. \end{aligned}$$

   For the Poisson(λ) distribution we have

$$\begin{aligned} m_1 &= \lambda \\ m_2 &= \lambda(1 + \lambda) \\ m_3 &= \lambda(1 + 3\lambda + \lambda^2) \\ m_4 &= \lambda(1 + 7\lambda + 6\lambda + \lambda^2), \end{aligned}$$

   which can be found using the mgf. Plugging these into our expression for ϑ, we obtain

$$\vartheta = \lambda + 2\lambda^2.$$

   So we have

$$\text{ARE}(\bar{X}_n; \hat{\sigma}_n^2) = (\lambda + 2\lambda^2)/\lambda = 1 + 2\lambda > 1.$$

   (b) Which estimator is more "efficient"?

   Since $\text{ARE}(\bar{X}_n; \hat{\sigma}_n^2) = 1 + 2\lambda > 1$, $\bar{X}_n$ is a more efficient estimator than $\hat{\sigma}_n^2$; that is it has a smaller asymptotic variance.

2. Let Z and W be independent rvs such that $Z(0,1)$ and $W \sim \chi^2_\nu$ and let ϕ be a constant. Then

$$\frac{Z + \phi}{\sqrt{W/\nu}} \sim t_{\nu,\phi},$$

where $t_{\nu,\phi}$ denotes the non-central t-distribution with degrees of freedom ν and non-centrality parameter ϕ.

Now, suppose $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$, and let μ_0 be a fixed constant.

(a) Give the distribution of $\sqrt{n}(\bar{X}_n - \mu_0)/S_n$.

We have

$$\sqrt{n}(\bar{X}_n - \mu_0)/S_n \sim t_{n-1,\phi},$$

where $\phi = \sqrt{n}(\mu - \mu_0)/\sigma$.

(b) Letting $F_{\nu,\phi}$ denote the cdf of the $t_{\nu,\phi}$ distribution, give the power function for the test of H_0: $\mu = \mu_0$ vs H_1: $\mu \neq \mu_0$ with decision rule $\phi(\mathbf{X}) = 1 \iff \sqrt{n}|\bar{X}_n - \mu_0|/S_n > t_{n-1,\alpha/2}$.

The power is given by

$$\gamma(\mu) = 1 - [F_{t_{\phi,n-1}}(t_{n-1,\alpha/2}) - F_{t_{\phi,n-1}}(-t_{n-1,\alpha/2})],$$

where $\phi = \phi = \sqrt{n}(\mu - \mu_0)/\sigma$.

(c) Suppose $\sigma = 1$, $\mu_0 = 3$, and $\alpha = 0.05$.

i. Use R to make a plot of the power function in part (b) under $n = 5$.

```
mu0 <- 3
sigma <- 1
n <- 5
alpha <- 0.05
mu <- seq(mu0 - 3,mu0 + 3,length=201)

power2 <- 1-(pt(qt(1-alpha/2,n-1),n-1,sqrt(n)*(mu-mu0)/sigma)
            -pt(-qt(1-alpha/2,n-1),n-1,sqrt(n)*(mu-mu0)/sigma))

plot(power2 ~ mu,
     type = "l",
     ylim = c(0,1),
     xlab = "mu",
     ylab = "power",
     main = paste("Power of t-test under mu0 = ",mu0,", n = ",n,",
                sigma = ",sigma,", alpha = ",alpha,sep=""),
     font.main = "1")

abline(h = alpha, lty = 3) # horiz line at size
```

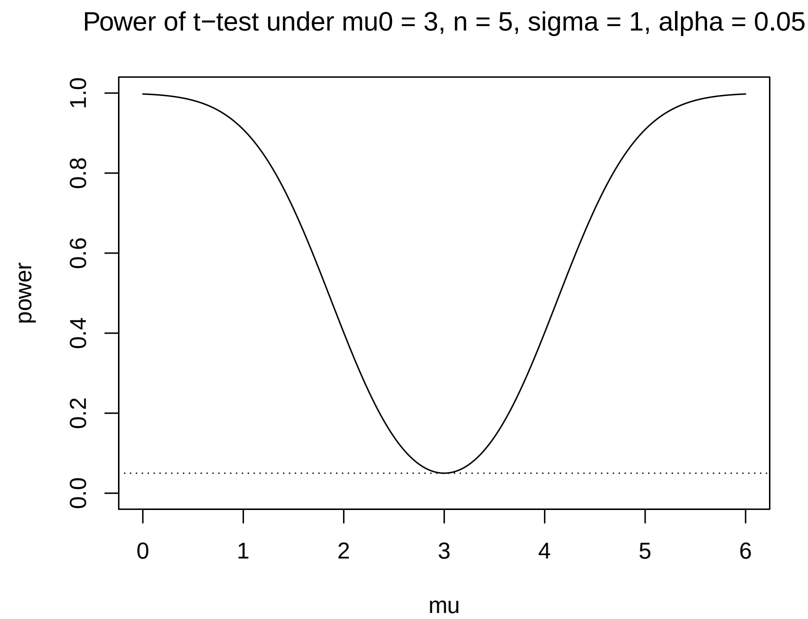

ii. Suppose $\mu = 4$. What is the smallest sample size n under which the test will reject H_0 with probability at least 0.90?

A sample of size $n = 14$ is needed. We can find this with the R code

```
mu <- 4
n <- 2:40
power2 <- 1-(pt(qt(1-alpha/2,n-1),n-1,sqrt(n)*(mu-mu0)/sigma)
             -pt(-qt(1-alpha/2,n-1),n-1,sqrt(n)*(mu-mu0)/sigma))
n[min(which(power2 >= 0.90))]
```

3. Let $X_1, \dots, X_n \overset{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2)$ and consider the test of H_0: $\sigma^2 \leq \sigma_0^2$ versus H_1: $\sigma^2 > \sigma_0^2$ with decision rule $\phi(\mathbf{X}) = 1 \iff (n-1)S_n^2/\sigma_0^2 > c$.

(a) Give an expression for the power function $\gamma(\sigma^2)$ of the test.

We have

$$\gamma(\sigma^2) = 1 - F_{\chi^2_{n-1}}((\sigma_0^2/\sigma^2)c),$$

where $F_{\chi^2_{n-1}}$ is the cdf of the χ^2_{n-1} distribution.

(b) Find c such that the test has size α.

Setting $c = \chi_{n-1,\alpha}$ gives the test size α.

(c) Suppose $\sigma_0^2 = 2$ and $\alpha = 0.05$.

i. Use R to make a plot of the power curve of the test under $n = 10$.

```
n <- 10
sigma0 <- sqrt(2)
alpha <- 0.05
sigma <- seq(1/8,3,length = 500)
powerR <- 1 - pchisq(qchisq(1-alpha,n-1)*sigma0^2/sigma^2,n-1)

sigma.sq <- sigma^2
plot(powerR ~ sigma.sq,
     type="l",ylim=c(0,1),xlab="sigma^2",ylab="power")

abline(v=sigma0^2,lty = 2) # vert line at null value
abline(h=alpha,lty = 3)         # horiz line at size
```

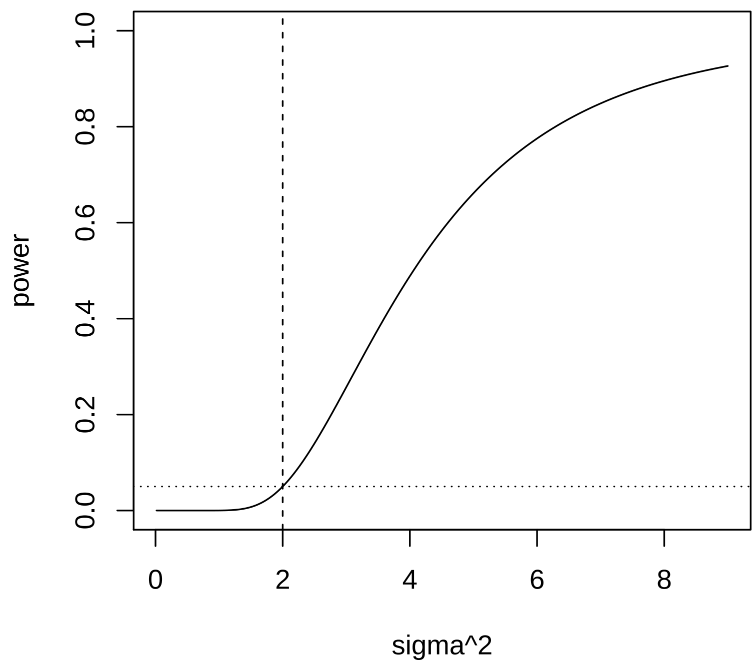

ii. Suppose $\sigma^2 = 2.1$. Find the sample size necessary to reject the null hypothesis with probability at least 0.80.

One needs a sample of size $n = 5169$. We can see this with the R code

```
sigma0 <- sqrt(2)
sigma <- sqrt(2.1)
```

```
n <- 2:6000
powerR <- 1 - pchisq(qchisq(1-alpha,n-1)*sigma0^2/sigma^2,n-1)
n[min(which(powerR >= .80))]
```

4. Let $X_{i1},\dots,X_{in_i} \overset{\text{ind}}{\sim} \text{Normal}(\mu_i,\sigma^2)$, $i = 1,2$, be two independent random samples. Let $\bar{X}_1$, $\bar{X}_2$, S_1^2, and S_2^2 be the means and variances of the two samples, respectively. Moreover, let $S^2_{\text{pooled}} = ((n_1-1)S_1^2 + (n_2-1)S_2^2)/(n_1+n_2-2)$.

   (a) Give the distribution of
   $$\frac{\bar{X}_1 - \bar{X}_2 - \delta}{S_{\text{pooled}}\sqrt{1/n_1 + 1/n_2}},$$
   where δ is a constant.

   This has the $t_{n_1+n_2-2,\phi}$ distribution with $\phi = \delta/(\sigma\sqrt{1/n_1+1/n_2})$.

   (b) For testing H_0: $\mu_1 - \mu_2 = 0$ versus H_1: $\mu_1 - \mu_2 \neq 0$, give the value of c such that the test with decision rule
   $$\phi(\mathbf{X}) = 1 \iff \frac{|\bar{X}_1 - \bar{X}_2|}{S_{\text{pooled}}\sqrt{1/n_1 + 1/n_2}} > c$$
   has size α.

   The value is $c = t_{n_1+n_2-2,\alpha/2}$.

   (c) Use R to plot the power curve as of the test in (b) with size $\alpha = 0.05$ under $n_1 = 10$, $n_2 = 15$, and $\sigma^2 = 4$. Make the power a function of $\delta = \mu_1 - \mu_2$.

   We can plot the power curve with the following R code:

```
delta <- seq(-4,4,length=100)
delta0 <- 0
n1 <- 10
n2 <- 15
alpha <- 0.05
sigma <- 2

phi <- (delta-delta0)/(sigma*sqrt(1/n1+1/n2))
power2 <- 1-(pt(qt(1-alpha/2,n1+n2-2),n1+n2-2,phi)
            - pt(-qt(1-alpha/2,n1+n2-2),n1+n2-2,phi))

plot(power2 ~ delta,type="l",ylim=c(0,1),xlab="delta",ylab="power")
abline(v=delta0,lty=2) # vert line at null value
abline(h=alpha,lty=3)         # horiz line at size
```

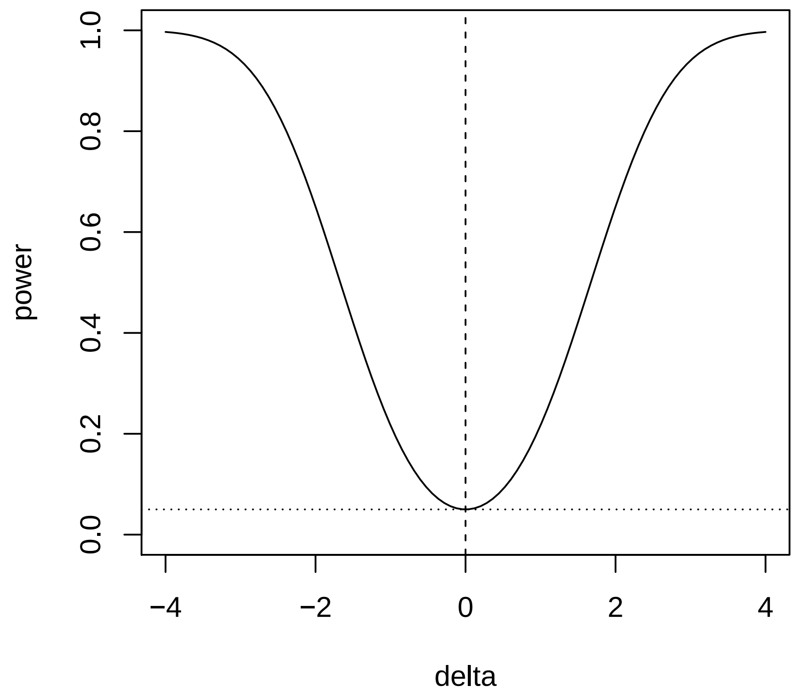

5. Let $X_{i1}, \ldots, X_{in_i} \overset{\text{ind}}{\sim} \text{Normal}(\mu_i, \sigma_i^2)$, $i = 1, 2$, be two independent random samples. Let $\bar{X}_1$, $\bar{X}_2$, S_1^2, and S_2^2 be the means and variances of the two samples, respectively.

   (a) Give the distribution of $(S_1^2/\sigma_1^2)/(S_2^2/\sigma_2^2)$.

   We have $(S_1^2/\sigma_1^2)/(S_2^2/\sigma_2^2) \sim F_{n_1-1,n_2-1}$.

   (b) Consider testing the hypotheses H_0: $\sigma_1^2/\sigma_2^2 \geq \vartheta_0$ versus H_1: $\sigma_1^2/\sigma_2^2 < \vartheta_0$ for some constant ϑ_0. Give an expression for the power function of the test with decision rule $\phi(\mathbf{X}) = 1 \iff (S_1^2/S_2^2)/\vartheta_0 < c$. Express the power as a function of $\vartheta = \sigma_1^2/\sigma_2^2$.

   The power is given by

$$\gamma(\vartheta) = F_{F_{n_1-1,n_2-1}}((\vartheta_0/\vartheta) \cdot c),$$

   where $F_{F_{n_1-1,n_2-1}}$ is the cdf of the F_{n_1-1,n_2-1} distribution.

   (c) Find c such that the test in (b) has size α.

   The size of the test is given by

$$\sup_{\vartheta \leq \vartheta} \gamma(\vartheta) = \gamma(\vartheta_0) = F_{F_{n_1-1,n_2-1}}(c).$$

Setting this equal to α gives the choice of c as the upper $1-\alpha$ quantile (which is the α quantile) of the F_{n_1-1,n_2-1} distribution, which we might denote by $F_{n_1-1,n_2-1,1-\alpha}$.

(d) Under $\alpha = 0.05$, $n_1 = 15$, and $n_2 = 5$, use R to plot the power function over a range of ϑ values when testing for equal variances.

We can plot the power curve with the following R code:

```
v <- seq(1/32,2,length=200)
v0 <- 1
n1 <- 15
n2 <- 5
alpha <- 0.05

powerR <- pf(qf(alpha,n1-1,n2-1)*v0/v,n1-1,n2-1)

plot(powerR ~ v, type = "l",ylim = c(0,1), xlab = "v", ylab = "power")
abline(v=v0,lty = 2)   # vert line at null value
abline(h=alpha,lty = 3) # horiz line at size
```