Non-Central Chi-squared[1]

STA442/2101 Fall 2016

[1] See last slide for copyright information.

Definition

If $X \sim N(\mu, 1)$ then $Y = X^2$ is said to have a *non-central chi-squared* distribution with degrees of freedom one and *non-centrality parameter* $\lambda = \mu^2$.

Write $Y \sim \chi^2(1, \lambda)$.

Facts about the non-central chi-squared distribution

$Y = X^2$ where $X \sim N(\mu, 1)$

$$Y \sim \chi^2(1, \lambda), \text{ where } \lambda \geq 0$$

- $Pr\{Y > 0\} = 1$, of course.
- If $\lambda = 0$, the non-central chi-squared reduces to the ordinary central chi-squared.
- The distribution is "stochastically increasing" in λ, meaning that if $Y_1 \sim \chi^2(1, \lambda_1)$ and $Y_2 \sim \chi^2(1, \lambda_2)$ with $\lambda_1 > \lambda_2$, then $Pr\{Y_1 > y\} > Pr\{Y_2 > y\}$ for any $y > 0$.
- $\lim_{\lambda \to \infty} Pr\{Y > y\} = 1$
- There are efficient algorithms for calculating non-central chi-squared probabilities. R's `pchisq` function does it.

Why is it called "non-central?"

Recall that if $X \sim N(\mu, 1)$ then $Y = X^2 \sim \chi^2(1, \lambda = \mu^2)$

If $X \sim N(\mu, \sigma^2)$, then if we center it and scale it,

- $Z^2 = \left(\frac{X-\mu}{\sigma}\right)^2 \sim \chi^2(1)$, the *central* chi-squared.
- What if we scale it correctly, but center it to the wrong place?
- $Z = \frac{X-\mu_0}{\sigma} \sim N(\frac{\mu-\mu_0}{\sigma}, 1)$
- And Z^2 is chi-squared with *df* $= 1$ and non-centrality parameter

$$\lambda = \left(\frac{\mu - \mu_0}{\sigma}\right)^2$$

This applies whether the normality is exact or asymptotic.

An example

Back to the coffee taste test

$Y_1, \ldots, Y_n \stackrel{i.i.d.}{\sim} B(1, \theta)$

$H_0 : \theta = \theta_0 = \frac{1}{2}$

Reject H_0 if $|Z_2| = \left| \frac{\sqrt{n}(\overline{Y} - \theta_0)}{\sqrt{\overline{Y}(1-\overline{Y})}} \right| > z_{\alpha/2}$

Suppose that in the population, 60% of consumers would prefer the new blend. If we test 100 consumers, what is the probability of obtaining results that are statistically significant?

That is, if $\theta = 0.60$, what is the power for $n = 100$? Earlier, got 0.53 with a direct standard normal calculation.

Non-central chi-squared

Recall that if $X \sim N(\mu, \sigma^2)$, then $\left(\frac{X-\mu_0}{\sigma}\right)^2 \sim \chi^2\left(1, \left(\frac{\mu-\mu_0}{\sigma}\right)^2\right)$.

For large n, the sample proportion $\overline{Y}$ is approximately normal with mean $\mu = \theta$ and variance $\sigma^2 = \frac{\theta(1-\theta)}{n}$. So,

$$
\begin{aligned}
Z_2^2 &= \left(\frac{\sqrt{n}(\overline{Y}-\theta_0)}{\overline{Y}(1-\overline{Y})}\right)^2 \\
&\approx \frac{(\overline{Y}-\theta_0)^2}{\theta(1-\theta)/n} \\
&= \left(\frac{\overline{Y}-\theta_0}{\sigma}\right)^2 \\
&\overset{approx}{\sim} \chi^2\left(1, n\,\frac{(\theta-\theta_0)^2}{\theta(1-\theta)}\right)
\end{aligned}
$$

We have found that

The Wald chi-squared test statistic of $H_0 : \theta = \theta_0$

$$Z_2^2 = \frac{n(\overline{Y} - \theta_0)^2}{\overline{Y}(1 - \overline{Y})}$$

has an asymptotic non-central chi-squared distribution with $df = 1$ and non-centrality parameter

$$\lambda = \frac{n(\theta - \theta_0)^2}{\theta(1 - \theta)}$$

Notice the similarity of test statistic and non-centrality parameter, and also that

- If $\theta = \theta_0$, then $\lambda = 0$ and Z_2^2 has a central chi-squared distribution.
- The probability of exceeding any critical value (power) can be made as large as desired by making λ bigger.
- There are 2 ways to make λ bigger. What are they?

Power calculation with R

For $n = 100$, $\theta_0 = 0.50$ and $\theta = 0.60$

```
> # Power for Wald chisquare test of H0: theta=theta0
> n=100; theta0=0.50; theta=0.60
> lambda = n * (theta-theta0)^2 / (theta*(1-theta))
> critval = qchisq(0.95,1)
> power = 1-pchisq(critval,1,lambda); power
[1] 0.5324209
```

Earlier, had

```
> Z2power(0.60,100)
[1] 0.5324209
```

Check power calculations by simulation

First develop and illustrate the code

```
# Try a simulation to test it.
set.seed(9999) # Set seed for "random" number generation
theta = 0.50; theta0 = 0.50; n = 100; m = 10
critval = qchisq(0.95,1); critval
p = rbinom(m,n,theta)/n; p
Z2 = sqrt(n)*(p-theta0)/sqrt(p*(1-p))
rbind(p,Z2)
sig = (Z2^2>critval); sig
sum(sig)/n
```

Output from the last slide

```
> # Try a simulation to test it.
> set.seed(9999) # Set seed for  "random" number generation
> theta = 0.50; theta0 = 0.50; n = 100; m = 10
> critval = qchisq(0.95,1); critval
[1] 3.841459
> p = rbinom(m,n,theta)/n; p
 [1] 0.40 0.56 0.47 0.57 0.47 0.50 0.58 0.48 0.40 0.53
> Z2 = sqrt(n)*(p-theta0)/sqrt(p*(1-p))
> rbind(p,Z2)
        [,1]     [,2]       [,3]     [,4]       [,5] [,6]     [,7]       [,8]      [,9]
p   0.400000 0.560000  0.4700000 0.570000  0.4700000  0.5 0.580000  0.4800000  0.400000
Z2 -2.041241 1.208734 -0.6010829 1.413925 -0.6010829  0.0 1.620882 -0.4003204 -2.041241
       [,10]
p  0.5300000
Z2 0.6010829
> sig = (Z2^2>critval); sig
 [1]  TRUE FALSE FALSE FALSE FALSE FALSE FALSE FALSE  TRUE FALSE
> sum(sig)/n
[1] 0.02
```

Now the real simulation

First estimated probability should equal about 0.05 because $\theta = \theta_0$

```
> # Check Type I error probability
> set.seed(9999)
> theta = 0.50; theta0 = 0.50; n = 100; m = 10000
> critval = qchisq(0.95,1)
> p = rbinom(m,n,theta)/n; Z2 = sqrt(n)*(p-theta0)/sqrt(p*(1-p))
> sig = (Z2^2>critval)
> sum(sig)/m
[1] 0.0574

> # Exact power calculation for theta=0.60 gives power = 0.5324209
> set.seed(9998)
> theta = 0.60; theta0 = 0.50; n = 100; m = 10000
> critval = qchisq(0.95,1)
> p = rbinom(m,n,theta)/n; Z2 = sqrt(n)*(p-theta0)/sqrt(p*(1-p))
> sig = (Z2^2>critval)
> sum(sig)/m
[1] 0.5353
```

Conclusions from the power analysis

- Power for $n = 100$ is pathetic.
- As Fisher said, "To call in the statistician after the experiment is done may be no more than asking him to perform a postmortem examination: he may be able to say what the experiment died of."
- $n = 200$ is better.

```
> n=200; theta0=0.50; theta=0.60
> lambda = n * (theta-theta0)^2 / (theta*(1-theta))
> power = 1-pchisq(critval,1,lambda); power
[1] 0.8229822
```

- What sample size is required for power of 90%?

What sample size is required for power of 90%?

```
> # Find sample size needed for power = 0.90
> theta0=0.50; theta=0.60; critval = qchisq(0.95,1)
> effectsize = (theta-theta0)^2 / (theta*(1-theta))
> n = 0
> power=0
> while(power < 0.90)
+ {
+ n = n+1
+ lambda = n * effectsize
+ power = 1-pchisq(critval,1,lambda)
+ }
> n; power
[1] 253
[1] 0.9009232
```

General non-central chi-squared

Let $X_1, \ldots, X_n$ be independent $N(\mu_i, \sigma_i^2)$. Then

$$Y = \sum_{i=1}^{n} \frac{X_i^2}{\sigma_i^2} \sim \chi^2(n, \lambda), \text{ where } \lambda = \sum_{i=1}^{n} \frac{\mu_i^2}{\sigma_i^2}$$

- Density is a bit messy — a Poisson mixture of central chi-squares.
- Reduces to central chi-squared when $\lambda = 0$.
- Generalizes to $Y \sim \chi^2(\nu, \lambda)$, where $\nu > 0$ as well as $\lambda > 0$.
- Stochastically increasing in λ, meaning $Pr\{Y > y\}$ can be increased by increasing λ.
- $\lim_{\lambda \to \infty} Pr\{Y > y\} = 1$
- Probabilities are easy to calculate numerically.

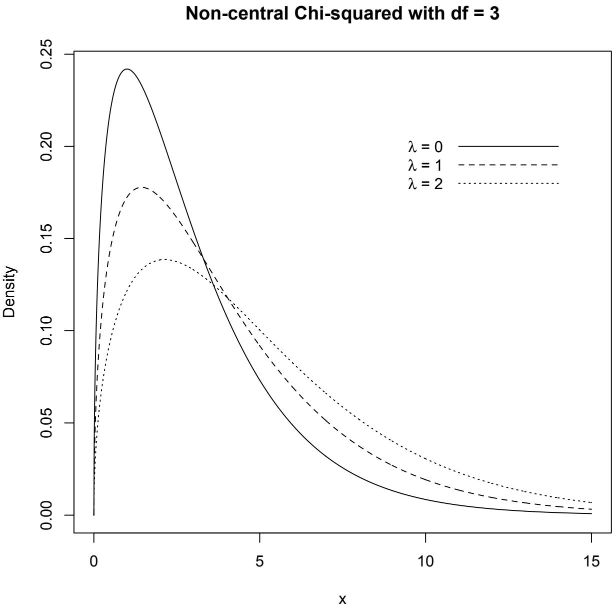
Non-central Chi-squared with df = 3
λ = 0
λ = 1
λ = 2
Density
0.00
0.05
0.10
0.15
0.20
0.25
0
5
10
15
x

R code for the record

```
# Plotting non-central chi-squared in R with Greek letters
lambda1 = 1; lambda2=2; top = 15; DF=3
titlestring = paste("Non-central Chi-squared with df =",DF)
x = seq(from=0,to=top,by=0.01)
Density = dchisq(x,df=DF)
d1 = dchisq(x,df=DF,ncp=lambda1); d2 = dchisq(x,df=DF,ncp=lambda2)
plot(x,Density,type="l",main=titlestring) # That's a lower case L
lines(x,d1,lty=2); lines(x,d2,lty=3)
# Make line labels
x1 <- c(11,14) ; y1 <- c(0.20,0.20) ; lines(x1,y1,lty=1)
x2 <- c(11,14) ; y2 <- c(0.19,0.19) ; lines(x2,y2,lty=2)
x3 <- c(11,14) ; y3 <- c(0.18,0.18) ; lines(x3,y3,lty=3)
caption0 = expression(paste(lambda," = 0")); text(10,0.20,caption0);
# Ugly but flexible: * means concatenation in expressions, and we
#                      need to substitute numerical values. First comes
#                      an expression, and then a list of substitutions.
caption1 = substitute( lambda * " = " * L1, list(L1=lambda1) )
text(10,0.19,caption1)
caption2 = substitute( lambda * " = " * L2, list(L2=lambda2) )
text(10,0.18,caption2)
```

Copyright Information