Testing the General Linear Hypothesis

The General Linear Hypothesis

For the linear model

$$\mathbf{Y} = \mathbf{X}\beta + \epsilon$$

we wish test the null hypothesis

$$\mathrm{H}_0 \ : \ \mathbf{A}\beta = \mathbf{c}$$

that places r linearly independent linear constraints on the $p \times 1$ vector β – the alternative hypothesis leaves β unconstrained.

$\mathbf{A}$ is an $r \times p$ matrix whose rows are linearly independent.

The General Linear Hypothesis (cont.)

Let $\widehat{\beta}$ denote the usual least squares estimator. Under Normality, we have as usual

$$\widehat{\beta}|\mathbf{X} \sim \text{Normal}_p(\beta, \sigma^2(\mathbf{X}^\top\mathbf{X})^{-1})$$

so therefore by properties of the Normal distribution.

$$\mathbf{A}\widehat{\beta} \mid \mathbf{X} \sim \text{Normal}_r(\mathbf{A}\beta, \sigma^2\mathbf{A}(\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{A}^\top).$$

Let $\mathbf{D} = \mathbf{A}(\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{A}^\top$. Then again by standard properties

$$\frac{1}{\sigma^2}(\mathbf{A}\widehat{\beta} - \mathbf{A}\beta)^\top\mathbf{D}^{-1}(\mathbf{A}\widehat{\beta} - \mathbf{A}\beta) \sim \text{Chisquared(r)}$$

The General Linear Hypothesis (cont.)

Finally, we have from earlier results that, under the null hypothesis

$$\mathbf{A}\beta = \mathbf{c}$$

and $\widehat{\beta}$ and $\widehat{\sigma}^2$ are independent, so therefore

$$F = \frac{(\mathbf{A}\widehat{\beta} - \mathbf{c})^\top \mathbf{D}^{-1} (\mathbf{A}\widehat{\beta} - \mathbf{c})/r}{\mathrm{SS}_{\mathrm{Res}}(\beta)/(n-p)} \sim \mathrm{Fisher}(r, n-p)$$

which may be used to test H_0.

The General Linear Hypothesis (cont.)

In the special case $\mathbf{c} = \mathbf{0}_r$, then the test statistic becomes

$$F = \frac{\widehat{\beta}^\top \mathbf{A}^\top \mathbf{D}^{-1} \mathbf{A} \widehat{\beta}}{r \mathrm{SS}_{\mathrm{Res}}(\beta)/(n-p)}$$

This is a "global" test; we may also consider individual or "univariate" tests of the r null hypotheses

$$\mathrm{H}_0 \ : \ \boldsymbol{a}_j \beta = 0 \qquad j = 1, \ldots, r$$

where $\boldsymbol{a}_j$ is the jth row of $\mathbf{A}$.

Simulation

Data generating model: $n = 1000$,

$$Y_i = 2 + 2x_{i1} - 0.5x_{i2} - 2x_{i3} + \epsilon_i$$

i.e. $\beta = (\beta_0, \beta_1, \beta_2, \beta_3)^\top$, with $\sigma = 1$.

All predictors are influential in the model.

The General Linear Hypothesis in R

Suppose that the null hypothesis of interest is specified

$$H_0 \ : \ \beta_1 + \beta_3 = 0, \beta_1 + \beta_3/2 = \beta_2$$

that is

$$\mathbf{A}\beta = 0_2$$

with

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 1 \\ 0 & 1 & -1 & 1/2 \end{bmatrix}$$

Here the $m = 2$ rows of $\mathbf{A}$ are linearly independent, so we have $r = 2$ constraints under H_0.

The General Linear Hypothesis in R

```
1  > library(multcomp)
2  > fit.tmp<-lm(Y~x1+x2+x3)
3  > A<-matrix(c(0,1,0,1,0,1,-1,1/2),byrow=T,nrow=2,ncol=4)
4  > A
5       [,1] [,2] [,3] [,4]
6  [1,]    0    1    0  1.0
7  [2,]    0    1   -1  0.5
8  > summary(glht(fit.tmp,linfct=A),Ftest())
9
10          General Linear Hypotheses
11
12 Linear Hypotheses:
13        Estimate
14 1 == 0 0.004762
15 2 == 0 1.513291
16
17 Global Test:
18      F DF1 DF2    Pr(>F)
19 1 1200   2 996 5.03e-266
```

The `glht` function in the `multcomp` library performs the general linear hypothesis F-test. We specify the matrix $\mathbf{A}$ (line 3); lines 17-19 tests this 'global' hypothesis, and here the hypothesis is rejected as the p-value is 5.03×10^{-266}

The General Linear Hypothesis in R

Lines 14 and 15 inspect the $r = 2$ components of the linear constraints; these individual components can also be tested.

```
20  > summary(glht(fit.tmp,linfct=A),test=univariate())
21
22           Simultaneous Tests for General Linear Hypotheses
23
24  Fit: lm(formula = Y ~ x1 + x2 + x3)
25
26  Linear Hypotheses:
27         Estimate Std. Error t value Pr(>|t|)
28  1 == 0 0.004762   0.069963   0.068    0.946
29  2 == 0 1.513291   0.036303  41.685   <2e-16 ***
30  ---
31  Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
32  (Univariate p values reported)
```

In this case:

- line 28: the first null hypothesis $H_0 : \beta_1 + \beta_3 = 0$ is **not rejected** ($p = 0.946$)
- line 29: the second null hypothesis $H_0 : \beta_1 + \beta_3/2 = \beta_2$ is **rejected** ($p < 2 \times 10^{-16}$)

Multiple Testing

Recall that if K independent hypothesis tests are carried out, each using significance level α, then if all the related null hypothesis are in fact true, the probability that none of the nulls is rejected is

$$(1-\alpha)^K.$$

Therefore the probability that at least one is (incorrectly) rejected is

$$\alpha_{\text{FWE}} = 1-(1-\alpha)^K.$$

This is the 'familywise' Type I error rate across the 'family' of K tests.

Recall that α specifies the 'per test' Type I error rate (the proportion of true null hypotheses that are incorrectly rejected).

Multiple Testing (cont.)

If $\alpha = 0.05$, then we have

K	α_{FWE}
2	0.0975
5	0.2262
10	0.4013
50	0.9231
100	0.9941

Multiple Testing (cont.)

If $\alpha = 0.01$, then we have

K	α_{FWE}
2	0.0199
5	0.0490
10	0.0956
50	0.3950
100	0.6340

Multiple Testing (cont.)

Therefore, as the per test α is reduced, the familywise α_{FWE} is also reduced.

This suggests that if we want to control the familywise error rate we need to use a smaller per test α.

- If we require $\alpha_{\text{FWE}} = 0.05$, we need to set

$$\alpha = 1 - (1 - \alpha_{\text{FWE}})^{1/K}$$

- The Bonferroni correction approximates this by

$$\alpha^{(\text{Bonf})} = \alpha_{\text{FWE}}/K.$$

Multiple Testing (cont.)

With $\alpha_{\text{FWE}} = 0.05$

K	α	$\alpha^{(\text{Bonf})}$
2	0.0253	0.0250
5	0.0102	0.0100
10	0.0051	0.0050
100	0.0005	0.0005

The General Linear Hypothesis: Multiple Testing

In the earlier analysis, no correction for multiplicity is made: we may also make the Bonferroni adjustment to the *p*-values

```
> summary(glht(fit.tmp,linfct=A),test=adjusted("bonf"))

         Simultaneous Tests for General Linear Hypotheses

Fit: lm(formula = Y ~ x1 + x2 + x3)

Linear Hypotheses:
       Estimate Std. Error t value Pr(>|t|)
1 == 0 0.004762   0.069963   0.068        1
2 == 0 1.513291   0.036303  41.685   <2e-16 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
(Adjusted p values reported -- bonferroni method)
```

In this case, there is no difference to the test result.

Confidence Ellipses

Confidence Ellipses

In the multiple regression problem, we could construct hypothesis tests and confidence intervals for the individual β_j parameters. However, we may also construct *simultaneous* confidence statements.

As in the previous calculation, under the Normality assumption,

$$\widehat{\beta}|\mathbf{X} \sim \text{Normal}_p(\beta, \sigma^2(\mathbf{X}^\top\mathbf{X})^{-1})$$

and so

$$F(\beta) = \frac{(\widehat{\beta}-\beta)^\top(\mathbf{X}^\top\mathbf{X})(\widehat{\beta}-\beta)/p}{\text{SS}_{\text{Res}}(\beta)/(n-p)} \sim \text{Fisher}(p, n-p)$$

Confidence Ellipses (cont.)

We may therefore construct a simultaneous $(1-\alpha)100\%$ confidence region as the set $\mathcal{C}_{\alpha,p,n-p} \in \mathbb{R}^p$ defined by

$$\mathcal{C}_{\alpha,p,n-p} = \left\{\beta : F(\beta) \leq F_{\alpha,p,n-p}\right\}.$$

where $F_{\alpha,p,n-p}$ is the $1-\alpha$ quantile of the Fisher$(p, n-p)$ distribution. This set is an ellipsoid in p-dimensional space.

We can compute similar ellipsoids for sub-vectors of β.

Confidence Ellipses (cont.)

Two dimensional confidence ellipses can be constructed using the `ellipse` package.

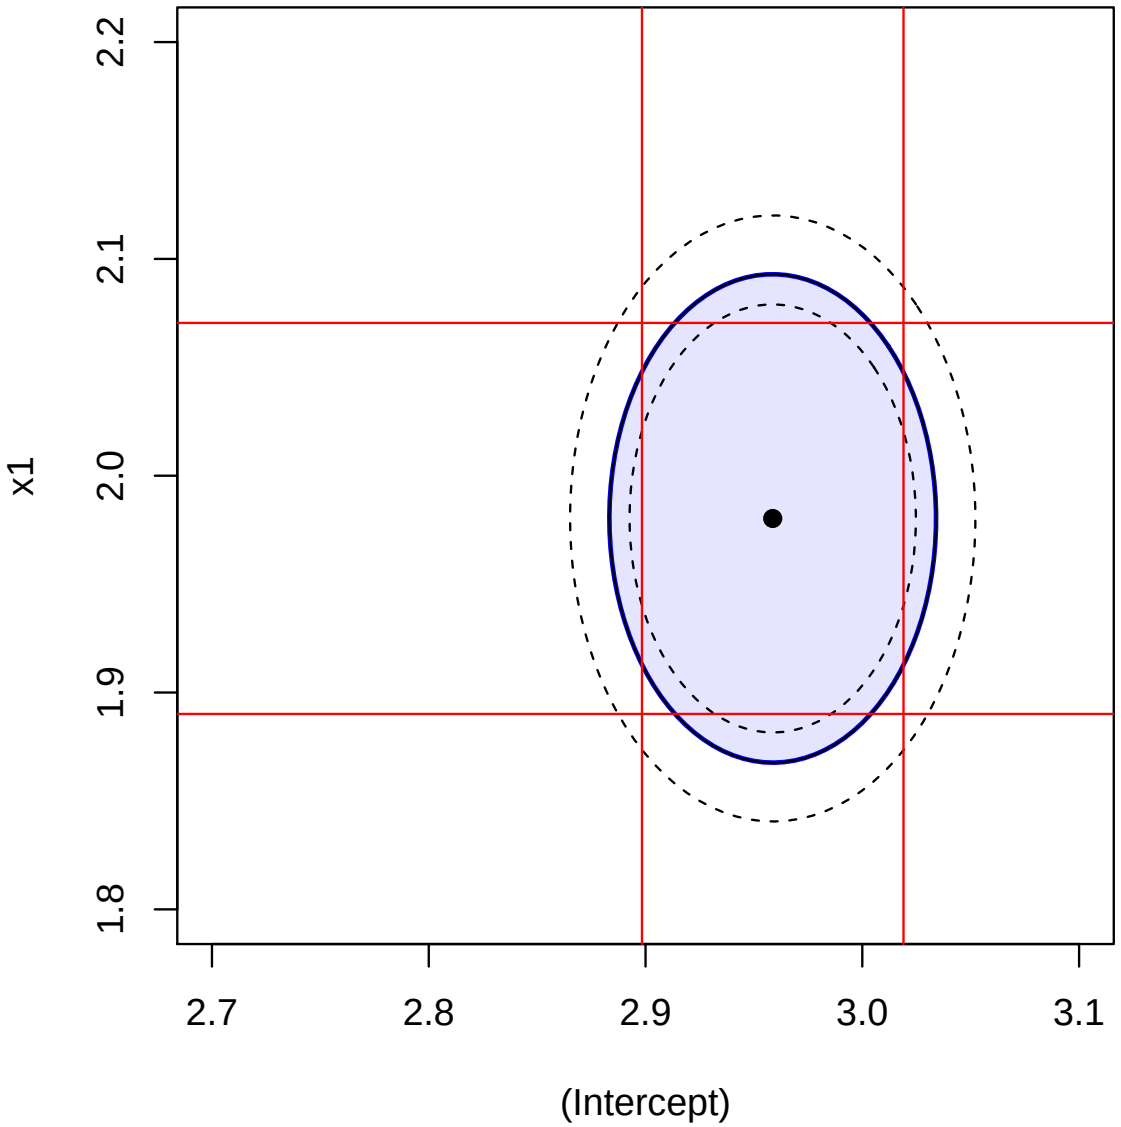
```
> library(ellipse)
> plot(ellipse(fit.tmp,level=0.99),
+ type='l',lty=2,xlim=range(2.70,3.10),ylim=range(1.80,2.20))
> par(pty='s')
> plot(ellipse(fit.tmp,level=0.99),
+ type='l',lty=2,xlim=range(2.70,3.10),ylim=range(1.80,2.20))
> lines(ellipse(fit.tmp,level=0.95),col='blue',lwd=2)
> polygon(ellipse(fit.tmp,level=0.95,which=c(1,2)),
+ col=rgb(0,0,1,alpha=.1))
> lines(ellipse(fit.tmp,level=0.90),lty=2)
> points(fit.tmp$coefficients[1], fit.tmp$coefficients[2],pch=19)
> abline(v=confint(fit.tmp)[1,],col='red')
> abline(h=confint(fit.tmp)[2,],col='red')
```

Confidence Ellipses: Key

- dot: parameter estimates;
- shaded area: 95% confidence ellipse;
- dotted lines: 99% and 90% confidence ellipses;
- red lines: boundaries of 95% marginal intervals.

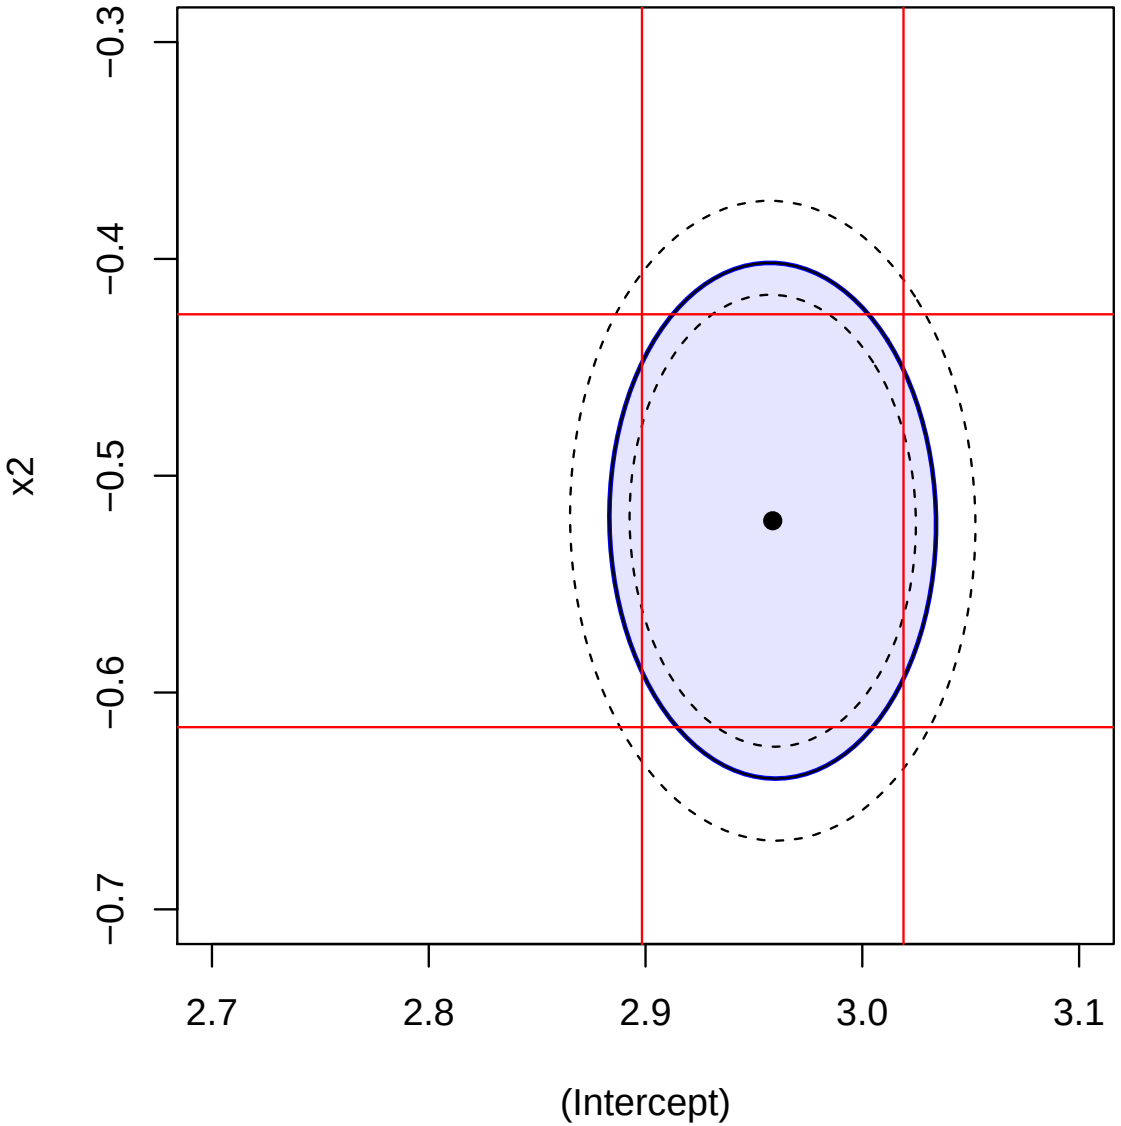
Confidence Ellipses

β_0 vs β_1:



Confidence Ellipses

β_0 vs β_2:



Confidence Ellipses

β_1 vs β_2:

