

Tools for Teaching Econometrics

Pierre Chausse*

Abstract

This vignette explains how to use the different tools I used in my econometrics courses at the University of Waterloo. It includes functions to generate solutions from inference questions, to print regression results and more.

1 Introduction

I use many R functions (R Core Team (2025)) in my courses to print regression results in a nice format, generate solutions for inference question on the mean, the variance, and least squares models, to simulate data to illustrate concepts, and so on. This document explains how to use them using examples from the courses. To run the different functions available in the package, you first need to load it:

```
library(metricsUW)
```

2 Printing regression results in equation format

To print a regression result from `lm` or `glm`, the method `printReg` applied to `lm` and `glm` objects generates a latex equation with the coefficient estimates and their standard errors. To have the equation printed in Latex format in a R-Markdown document, the chunk option `results='asis'` must be added. For example, the following is an estimated wage equation using the PSID1976 dataset from the AER package (Kleiber and Zeileis (2008)):

```
data(PSID1976, package="AER")
fit1 <- lm(wage~education+age+I(age^2), PSID1976)
printReg(fit1)
```

$$\widehat{wage} = \underset{(3.2900)}{-8.5590} + \underset{(0.0495)}{0.4538} education + \underset{(0.1524)}{0.2558} age - \underset{(0.0018)}{0.0029} I(age^2)$$
$$n = 753, R^2 = 0.1047, SSR = 7075.373$$

The function offers different options. You can add stars for significant coefficients (`stars=TRUE`), adjusted R^2 (`adjrsq=TRUE`), limit the number of variables per line when the equation does not fit the page (e.g. `maxpl=3`), replace default standard errors by robust ones (`se=newse`), replace the default t-distribution used to compute p-values by the N(0,1) distribution (`dist="n"`) or omit variables. We present here a few examples.

- Adding adjusted R^2 and stars, and reducing the number of variables per line. I also increase the number of digits.

```
fit2 <- lm(wage~education*city+heducation+age+I(age^2), PSID1976)
printReg(fit2, maxpl=3, adjrs=TRUE, stars=TRUE, digits=5)
```

*University of Waterloo, pchausse@uwaterloo.ca

$$\widehat{wage} = -6.83341 + \frac{0.45828}{(0.09077)^{***}} education - \frac{1.42391}{(1.27155)} cityyes - \frac{0.12139}{(0.04801)^{**}} heducation$$

$$+ \frac{0.24075}{(0.15272)} age - \frac{0.00284}{(0.00176)} I(age^2) + \frac{0.13427}{(0.10318)} education : cityyes$$

$$n = 753, R^2 = 0.11471, SSR = 6996.553, \bar{R}^2 = 0.10759$$

$$*pv < 0.1; **pv < 0.05; ***pv < 0.01$$

- Replacing default standard errors by robust ones, and using the standard normal distribution for p-values (Zeileis (2004)).

```
library(sandwich)
newse <- sqrt(diag(vcovHC(fit2)))
printReg(fit2, maxpl=3, se=newse, stars=TRUE)
```

$$\widehat{wage} = -6.8334 + \frac{0.4583}{(0.0719)^{***}} education - \frac{1.4239}{(1.1819)} cityyes - \frac{0.1214}{(0.0419)^{***}} heducation$$

$$+ \frac{0.2408}{(0.1559)} age - \frac{0.0028}{(0.0018)} I(age^2) + \frac{0.1343}{(0.1002)} education : cityyes$$

$$n = 753, R^2 = 0.1147, SSR = 6996.553 \text{ (Robust S-E)}$$

$$*pv < 0.1; **pv < 0.05; ***pv < 0.01$$

For GLM estimation, the R^2 is replaced by residual deviance and AIC is printed. Also, the left-hand side specifies that it is the link of \hat{Y} that has the linear representation. For example, the following is the result from a Poisson regression with the log link, using the dataset `fertil2` from the `wooldridge` package (Shea (2024)):

```
data(fertil2, package="wooldridge")
fit3 <- glm(children~educ+age+I(age^2)+catholic+electric+radio+tv+heduc,
            family=poisson(link=log), data=fertil2)
printReg(fit3, maxpl=3, stars=TRUE)
```

$$\text{link}[\widehat{children}] = -3.8200 - \frac{0.0187}{(0.0042)^{***}} educ + \frac{0.2650}{(0.0148)^{***}} age - \frac{0.0031}{(0.0002)^{***}} I(age^2)$$

$$+ \frac{0.0104}{(0.0414)} catholic - \frac{0.0803}{(0.0456)^*} electric + \frac{0.0444}{(0.0279)} radio$$

$$- \frac{0.1229}{(0.0548)^{**}} tv - \frac{0.0126}{(0.0036)^{***}} heduc$$

$$n = 1953, AIC = 7344.376, \text{ Residual Deviance} = 1834.816,$$

$$\text{Null Deviance} = 3246.903, \text{ Family} = \text{poisson}, \text{ Link} = \text{log}$$

$$*pv < 0.1; **pv < 0.05; ***pv < 0.01$$

If we just want to create a regression equation from a formula, the method admits objects of class `formula`. Here is an example:

```
printReg(log(wage)~education*female+age+I(age^2), maxpl=3)
```

$$\log(wage) = \beta_0 + \beta_1 education + \beta_2 female + \beta_3 age$$

$$+ \beta_4 I(age^2) + \beta_5 education : female + \epsilon$$

3 Solution to inference questions

The way the functions are organized is as follows. First the inference function generates the solution, which is an object of class `metricsSol`, and the `print` method generates the answer in Latex format. It is meant to be printed inside an R-Markdown chunk, with the option `results="asis"`.

The purpose of these functions is to generate questions for assignments, exercises or exams and have a fast way to produce a detailed solution document. The questions can be generated by datasets, in which case the solutions are derived from the same dataset, or by made up numbers.

3.1 Introduction to Statistics

This section covers inference problems typically covered in an introductory statistics course.

3.1.1 Test on the mean

The solution is based on the following properties:

- If $X_i \sim N(\mu, \sigma^2)$, then

$$\frac{\sqrt{n}(\bar{X} - \mu)}{\hat{\sigma}} \sim t_{n-1},$$

where n is the sample size, \bar{X} is the sample mean and $\hat{\sigma}$ is the sample standard errors defined as $\sum_{i=1}^n (X_i - \bar{X})^2 / (n-1)$.

- If the distribution of X_i is unknown, the true distribution of the test is also unknown, but the distribution it converges to is known. In my course, the t-distribution is only used when it is the exact distribution. If not, we use the $N(0,1)$ as an approximation. The solution generator applies this rule as well.

Suppose we have a series and want to test the hypothesis $H_0 : \mu = c$ against $H_1 : \mu \neq c$, $H_1 : \mu > c$ or $H_1 : \mu < c$. The function `testm` generates the solution. Consider the `wage` series from the PSID1976 dataset, restricted to workers with positive hours:

```
wage <- subset(PSID1976, hours>0)$wage
```

We want to test if the population average is equal to 5 dollars per hour against the alternative that it is not equal to 5. We just insert the following code in the chunk:

```
testMean(wage, h0=5)
```

Testing $H_0 : \mu = 5$ against $H_1 : \mu \neq 5$ at 5%

$$test = \frac{(\bar{x} - 5)}{s/\sqrt{n}} = \frac{(4.1777 - 5)}{(3.3103)/\sqrt{428}} = -5.1392 \sim t_{427}$$

Since $|-5.1392| > 1.9655$ (the 97.5% quantile of the t_{427}), we reject H_0 .

It is not necessary to use the `print` method directly, unless we want to add something to the solution. For example, I like to add mark distribution in my exam solutions:

```
sol1 <- testMean(wage, h0=5)
print(sol1, addMess="1 point for the statistic and 1 point for the conclusion")
```

Testing $H_0 : \mu = 5$ against $H_1 : \mu \neq 5$ at 5%

$$test = \frac{(\bar{x} - 5)}{s/\sqrt{n}} = \frac{(4.1777 - 5)}{(3.3103)/\sqrt{428}} = -5.1392 \sim t_{427}$$

Since $|-5.1392| > 1.9655$ (the 97.5% quantile of the t_{427}), we reject H_0 . **(1 point for the statistic and 1 point for the conclusion)**

By default, we assume normality of the data, a size of 5% and a two-sided alternative. When we assume normality, the distribution used for the critical value is the t-distribution with $(n - 1)$ degrees of freedom. If we don't, the critical value is based on the asymptotic $N(0,1)$ property of the test. It is also possible to change the alternative hypothesis by either "greater" or "less" and the size of the test:

```
testMean(wage, h0=5, size=0.10, alter="less", assume="nonNormal")
```

Testing $H_0 : \mu = 5$ against $H_1 : \mu < 5$ at 10%

$$test = \frac{(\bar{x} - 5)}{s/\sqrt{n}} = \frac{(4.1777 - 5)}{(3.3103)/\sqrt{428}} = -5.1392 \approx N(0, 1)$$

Since $-5.1392 < -1.2816$ (the 10% quantile of the $N(0, 1)$), we reject H_0 . (The $N(0, 1)$ is an approximation based on the C.L.T. because the distribution of the data is unknown)

By default, all decimals are kept to compute the statistic. This could lead to slightly different solutions when the question is asked in an exam, because the printed numbers are rounded. It is possible to round the sample mean and standard errors before computing the test, through the argument `dround`. For example, suppose an exam question was generated directly from the data and the following table was printed using `stargazer`:

```
library(stargazer)
```

```
##
```

```
## Please cite as:
```

```
## Hlavac, Marek (2022). stargazer: Well-Formatted Regression and Summary Statistics Tables.
```

```
## R package version 5.2.3. https://CRAN.R-project.org/package=stargazer
```

```
stargazer(data.frame(wage), digits=3, header=FALSE, float=FALSE)
```

Statistic	N	Mean	St. Dev.	Min	Max
wage	428	4.178	3.310	0.128	25.000

We would generate the solution as follows:

```
testMean(wage, h0=5, size=0.10, alter="less", dround=3)
```

Testing $H_0 : \mu = 5$ against $H_1 : \mu < 5$ at 10%

$$test = \frac{(\bar{x} - 5)}{s/\sqrt{n}} = \frac{(4.178 - 5)}{(3.31)/\sqrt{428}} = -5.1377 \sim t_{427}$$

Since $-5.1377 < -1.2835$ (the 10% quantile of the t_{427}), we reject H_0 .

This option exists for all solution generator, so we won't discuss it further. It is also possible to generate a solution without data. We just need to provide the sample mean, the standard error and the sample size:

```
testMean(h0=4, xbar=3.7, se=0.8, n=40)
```

Testing $H_0 : \mu = 4$ against $H_1 : \mu \neq 4$ at 5%

$$test = \frac{(\bar{x} - 4)}{s/\sqrt{n}} = \frac{(3.7 - 4)}{(0.8)/\sqrt{40}} = -2.3717 \sim t_{39}$$

Since $|-2.3717| > 2.0227$ (the 97.5% quantile of the t_{39}), we reject H_0 .

3.1.2 Confidence interval for the mean

The $(1 - \alpha) \times 100\%$ confidence interval is defined as:

$$[\bar{X} - q_{1-\alpha/2}\hat{\sigma}, \bar{X} + q_{1-\alpha/2}\hat{\sigma}],$$

where $q_{1-\alpha/2}$ is the $(1 - \alpha/2) \times 100\%$ quantile of the t_{n-1} when $X_i \sim N(\mu, \sigma^2)$, in which case the coverage is exact, and the $N(0,1)$ when the distribution of X_i is unknown. For the latter case, the coverage is just an approximation based on the CLT. This rule is consistent with the rule for tests on the mean described in the previous section.

If we use the wage data from the previous section, the 95% confidence interval for the average wage, assuming normality, is:

```
ciMean(wage)
```

95% confidence interval for the mean

$$\begin{aligned} CI &= \left[\bar{x} - t^* \frac{s}{\sqrt{n}}, \bar{x} + t^* \frac{s}{\sqrt{n}} \right] \\ &= \left[4.1777 - 1.9655 \frac{3.3103}{\sqrt{428}}, 4.1777 + 1.9655 \frac{3.3103}{\sqrt{428}} \right] \\ &= [3.8632, 4.4922] \end{aligned}$$

The t^* is the 97.5% quantile of the t-distribution with 427 degrees of freedom.

As for `testMean`, we can choose not to assume normality and change the `size`. It is also possible to round the mean and standard deviation to match the solution of written questions.

```
ciMean(wage, size=0.15, assume="nonNormal", dround=3)
```

85% confidence interval for the mean

$$\begin{aligned} CI &= \left[\bar{x} - t^* \frac{s}{\sqrt{n}}, \bar{x} + t^* \frac{s}{\sqrt{n}} \right] \\ &= \left[4.178 - 1.4395 \frac{3.31}{\sqrt{428}}, 4.178 + 1.4395 \frac{3.31}{\sqrt{428}} \right] \\ &= [3.9477, 4.4083] \end{aligned}$$

The t^* is the 92.5% quantile of the $N(0,1)$ (The $N(0,1)$ is an approximation based on the C.L.T. because the distribution of the data is unknown).

Finally, we can specify the sample mean, standard deviation and sample size:

```
ciMean(xbar=3.2, se=0.9, n=32)
```

95% confidence interval for the mean

$$\begin{aligned} CI &= \left[\bar{x} - t^* \frac{s}{\sqrt{n}}, \bar{x} + t^* \frac{s}{\sqrt{n}} \right] \\ &= \left[3.2 - 2.0395 \frac{0.9}{\sqrt{32}}, 3.2 + 2.0395 \frac{0.9}{\sqrt{32}} \right] \\ &= [2.8755, 3.5245] \end{aligned}$$

The t^* is the 97.5% quantile of the t-distribution with 31 degrees of freedom.

3.1.3 Tests on the difference between two means

The following properties are assumed in the solution generator `testDiffMeans`. If we have two samples of sizes n_1 and n_2 for X_1 and X_2 , and want test the hypothesis $H_0 : \mu_1 - \mu_2 = c$, then:

- If $X_1 \sim N(\mu_1, \sigma^2)$ and $X_2 \sim N(\mu_2, \sigma^2)$, which implies that we assume equal variance, we have the following result under the null:

$$\frac{\bar{X}_1 - \bar{X}_2 - c}{\hat{\sigma} \sqrt{\frac{1}{n_1} + \frac{1}{n_2}}} \sim t_{n_1+n_2-2},$$

where

$$\hat{\sigma}^2 = \frac{1}{n_1 + n_2 - 2} [\hat{\sigma}_1^2(n_1 - 1) + \hat{\sigma}_2^2(n_2 - 1)]$$

and $\hat{\sigma}_1^2$ and $\hat{\sigma}_2^2$ are the usual bias corrected estimator of the variance of X_1 and X_2 .

- For any other cases, which include non-normality and/or non-equal variances, the exact distribution of the test is unknown, so we use the approximated $N(0,1)$ instead of the t-distribution. If the variances are not equal, the above test is not valid and must be replaced by:

$$\frac{\bar{X}_1 - \bar{X}_2 - c}{\sqrt{\frac{\hat{\sigma}_1^2}{n_1} + \frac{\hat{\sigma}_2^2}{n_2}}} \approx N(0, 1).$$

Note: The test presented here assumes that $Cov(X_1, X_2) = 0$. We do not cover the non-zero covariance case in my courses, so it is not yet implemented.

Suppose we want to test if the average wage is the same for workers living in a city and the ones not living in a city, we can proceed as follows. We first consider the normal case with equal variances.

```
wageCity <- subset(PSID1976, hours>0 & city=="yes")$wage
wageNoCity <- subset(PSID1976, hours>0 & city=="no")$wage
testDiffMeans(x1=wageCity, x2=wageNoCity, h0=0)
```

Testing $H_0 : \mu_1 - \mu_2 = 0$ against $H_1 : \mu_1 - \mu_2 \neq 0$ at 5%

$$\begin{aligned} s &= \sqrt{\frac{1}{n_1 + n_2 - 2} [\hat{\sigma}_1^2(n_1 - 1) + \hat{\sigma}_2^2(n_2 - 1)] \left(\frac{1}{n_1} + \frac{1}{n_2} \right)} \\ &= \sqrt{\frac{1}{426} [13.4736 \times (274 - 1) + 6.1053 \times (154 - 1)] \left(\frac{1}{274} + \frac{1}{154} \right)} \\ &= 0.3314 \end{aligned}$$

$$test = \frac{(\bar{x}_1 - \bar{x}_2)}{s} = \frac{(4.4735 - 3.6513)}{0.3314} = 2.4813 \sim t_{426}$$

Since $|2.4813| > 1.9655$ (the 97.5% quantile of the t_{426}), we reject H_0 .

For any other cases, the approximated $N(0,1)$ is used. Here is an example with other specifications:

```
testDiffMeans(x1=wageCity, x2=wageNoCity, h0=1, assumev="diff", size=0.10,
  alter="less")
```

Testing $H_0 : \mu_1 - \mu_2 = 1$ against $H_1 : \mu_1 - \mu_2 < 1$ at 10%

$$s = \sqrt{\frac{\hat{\sigma}_1^2}{n_1} + \frac{\hat{\sigma}_2^2}{n_2}} = \sqrt{\frac{13.4736}{274} + \frac{6.1053}{154}} = 0.298$$

$$test = \frac{(\bar{x}_1 - \bar{x}_2) - 1}{s} = \frac{(4.4735 - 3.6513) - 1}{0.298} = -0.5963 \approx N(0, 1)$$

Since $-0.5963 > -1.2816$ (the 90% quantile of the $N(0, 1)$), we do not reject H_0 . (The $N(0, 1)$ is an approximation based on the C.L.T. because the variances are not the same)

As for the other tests, we can input estimated means and standard errors instead of vectors. The arguments `xbar`, `se`, and `n` must be vectors of two:

```
testDiffMeans(h0=1, xbar=c(2.2, 3.3), se=c(3.4, 4.6), n=c(34, 76))
```

Testing $H_0 : \mu_1 - \mu_2 = 1$ against $H_1 : \mu_1 - \mu_2 \neq 1$ at 5%

$$\begin{aligned} s &= \sqrt{\frac{1}{n_1 + n_2 - 2} [\hat{\sigma}_1^2(n_1 - 1) + \hat{\sigma}_2^2(n_2 - 1)] \left(\frac{1}{n_1} + \frac{1}{n_2} \right)} \\ &= \sqrt{\frac{1}{108} [11.56 \times (34 - 1) + 21.16 \times (76 - 1)] \left(\frac{1}{34} + \frac{1}{76} \right)} \\ &= 0.8809 \\ test &= \frac{(\bar{x}_1 - \bar{x}_2) - 1}{s} = \frac{(2.2 - 3.3) - 1}{0.8809} = -2.3841 \sim t_{108} \end{aligned}$$

Since $|-2.3841| > 1.9822$ (the 97.5% quantile of the t_{108}), we reject H_0 .

3.1.4 Confidence intervals for the difference between two means

The theory from the previous section also applies here: we use the t-distribution only if the data is normally distributed and the variances of X_1 and X_2 are the same. Also, the standard deviation of $\bar{X}_1 - \bar{X}_2$ is

$$s = \sqrt{\frac{1}{n_1 + n_2 - 2} [\hat{\sigma}_1^2(n_1 - 1) + \hat{\sigma}_2^2(n_2 - 1)] \left(\frac{1}{n_1} + \frac{1}{n_2} \right)},$$

if the variances are the same and the following if they are not:

$$s = \sqrt{\frac{\hat{\sigma}_1^2}{n_1} + \frac{\hat{\sigma}_2^2}{n_2}}.$$

The confidence interval for the difference in average wage between workers from a city the ones not from a city, assuming normality and equal variance, is

```
ciDiffMeans(wageCity, wageNoCity)
```

95% confidence interval for $(\mu_1 - \mu_2)$

$$\begin{aligned} s &= \sqrt{\frac{1}{n_1 + n_2 - 2} [\hat{\sigma}_1^2(n_1 - 1) + \hat{\sigma}_2^2(n_2 - 1)] \left(\frac{1}{n_1} + \frac{1}{n_2} \right)} \\ &= \sqrt{\frac{1}{426} [13.4736 \times (274 - 1) + 6.1053 \times (154 - 1)] \left(\frac{1}{274} + \frac{1}{154} \right)} \\ &= 0.3314 \end{aligned}$$

$$\begin{aligned}
CI &= [(\bar{X}_1 - \bar{X}_2) - t^*s, (\bar{X}_1 - \bar{X}_2) + t^*s] \\
&= [(4.4735 - 3.6513) - 1.9655 \times 0.3314, (4.4735 - 3.6513) + 1.9655 \times 0.3314] \\
&= [0.1709, 1.4737]
\end{aligned}$$

The t^* is the 97.5% quantile of the t-distribution with 426 degrees of freedom.

If we relax the normality and/or the equal variance we obtain:

```
ciDiffMeans(wageCity, wageNoCity, assumev="diff", size=0.1)
```

90% confidence interval for $(\mu_1 - \mu_2)$

$$\begin{aligned}
s &= \sqrt{\frac{\hat{\sigma}_1^2}{n_1} + \frac{\hat{\sigma}_2^2}{n_2}} = \sqrt{\frac{13.4736}{274} + \frac{6.1053}{154}} = 0.298 \\
CI &= [(\bar{X}_1 - \bar{X}_2) - t^*s, (\bar{X}_1 - \bar{X}_2) + t^*s] \\
&= [(4.4735 - 3.6513) - 1.6449 \times 0.298, (4.4735 - 3.6513) + 1.6449 \times 0.298] \\
&= [0.3321, 1.3125]
\end{aligned}$$

The t^* is the 95% quantile of the $N(0, 1)$ (The $N(0, 1)$ is an approximation based on the C.L.T. because the distribution of statistic is unknown).

As for testing the difference between means, we can also replace the vectors of observations by **xbar**, **se**, and **n**.

3.1.5 Test on the variance

The only implemented test at the moment is the one that assumes normality. The assumption implies that under the null $H_0 : \sigma^2 = c$, we have the following distribution:

$$\frac{(n-1)\hat{\sigma}^2}{c} \sim \chi_{n-1}^2$$

Let's consider the following summary statistics from the **hprice1** dataset:

```
data(hprice1, package="wooldridge")
hprice1$lotsize <- hprice1$lotsize/1000
hprice1$sqrft <- hprice1$sqrft/1000
stargazer(hprice1[,1:5], digits=4, header=FALSE, float=FALSE)
```

Statistic	N	Mean	St. Dev.	Min	Max
price	88	293.5460	102.7134	111.0000	725.0000
assess	88	315.7364	95.3144	198.7000	708.6000
bdrms	88	3.5682	0.8414	2	7
lotsize	88	9.0199	10.1742	1.0000	92.6810
sqrft	88	2.0137	0.5772	1.1710	3.8800

Suppose we want to test $H_0 : \sigma^2 = 130$ against $H_1 : \sigma^2 \neq 130$ for the lot size. In the following I show what happens if we set **assume** to "nonNormal". The test is performed, but a note is added saying that the chi-square is not a valid distribution.

```
testVar(hprice1$lotsize, h0=130, assume="nonNormal")
```

Testing $H_0 : \sigma^2 = 130$ against $H_1 : \sigma^2 \neq 130$ at 5%

$$test = \frac{(n-1)\hat{\sigma}^2}{130} = \frac{(87)103.5133}{130} = 69.2743 \sim \chi_{87}^2$$

Since $69.2743 > 63.0894$ (the 2.5% quantile of the χ^2_{87}) and $69.2743 < 114.6929$ (the 97.5% quantile of the χ^2_{87}), we do not reject H_0 . (The χ^2_{87} is not valid in this case because the distribution of the data is unknown)

We can use one-sided tests and change the size:

```
testVar(hprice1$lotsize, h0=130, alter="less", size=0.1)
```

Testing $H_0 : \sigma^2 = 130$ against $H_1 : \sigma^2 < 130$ at 10%

$$test = \frac{(n-1)\hat{\sigma}^2}{130} = \frac{(87)103.5133}{130} = 69.2743 \sim \chi^2_{87}$$

Since $69.2743 < 70.581$ (the 10% quantile of the χ^2_{87}), we reject H_0 .

```
testVar(hprice1$lotsize, h0=70, alter="greater", size=0.01)
```

Testing $H_0 : \sigma^2 = 70$ against $H_1 : \sigma^2 > 70$ at 1%

$$test = \frac{(n-1)\hat{\sigma}^2}{70} = \frac{(87)103.5133}{70} = 128.6523 \sim \chi^2_{87}$$

Since $128.6523 > 120.591$ (the 99% quantile of the χ^2_{87}), we reject H_0 .

We can also replace the numeric vector by values of **se** and **n**. For example, we want to use the number from the table and test $H_0 : \sigma^2 = 0.25$ for square footage, we proceed as follows:

```
testVar(se=0.5772, n=88, h0=0.25)
```

Testing $H_0 : \sigma^2 = 0.25$ against $H_1 : \sigma^2 \neq 0.25$ at 5%

$$test = \frac{(n-1)\hat{\sigma}^2}{0.25} = \frac{(87)0.3332}{0.25} = 115.9396 \sim \chi^2_{87}$$

Since $115.9396 > 114.6929$ (the 97.5% quantile of the χ^2_{87}), we reject H_0 .

3.1.6 Confidence intervals for the variance

As for testing the variance in the previous section, the confidence intervals presented here are only valid under normality. It is based on the same chi-square distribution. For example, if we want to generate a confidence interval for the variance of **sqrft** using the data from the previous table, we can do it using the dataset (setting **dround** to 4 as in the table):

```
ciVar(x=hprice1$sqrft, dround=4)
```

95% confidence interval for the variance

$$\begin{aligned} CI &= \left[\frac{(n-1)S^2}{Q_u^*}, \frac{(n-1)S^2}{Q_l^*} \right] \\ &= \left[87 \times \frac{0.3332}{114.6929}, 87 \times \frac{0.3332}{63.0894} \right] \\ &= [0.2527, 0.4595] \end{aligned}$$

The Q_u^* is the 97.5% quantile of the χ^2_{87} and Q_l^* is its 2.5% quantile.

This is equivalent to using the numbers from the table:

```
ciVar(se=0.5772, n=88)
```

95% confidence interval for the variance

$$\begin{aligned}
 CI &= \left[\frac{(n-1)S^2}{Q_u^*}, \frac{(n-1)S^2}{Q_l^*} \right] \\
 &= \left[87 \times \frac{0.3332}{114.6929}, 87 \times \frac{0.3332}{63.0894} \right] \\
 &= [0.2527, 0.4594]
 \end{aligned}$$

The Q_u^* is the 97.5% quantile of the χ_{87}^2 and Q_l^* is its 2.5% quantile.

As for the test, we can change the `size` for a different confidence level, and set `assume` to "nonNormal" to add a note about the distribution not being valid under non-normality.

3.1.7 Testing equality between two variances

We want to test $H_0 : \sigma_1^2 = \sigma_2^2$ against $H_1 : \sigma_1^2 \neq \sigma_2^2$ at $\alpha \times 100\%$. The only test I cover in the intro to statistics is the one that assumes normality of the two samples and their independence. Under these assumptions and the null hypothesis, we have the following distribution:

$$\frac{\hat{\sigma}_1^2}{\hat{\sigma}_2^2} \sim F(n_1 - 1, n_2 - 1)$$

For this test to be two-sided, we define $\hat{\sigma}_1^2$ as the highest estimate of the variance between the two samples. We reject of the statistic exceeds the $(1 - \alpha) \times 100\%$ quantile of the distribution.

Consider the following summary statistics. The first table is for female workers and the second is not male workers. The dataset used to construct those tables is `wage1` form the `wooldridge` package.

Statistic	N	Mean	St. Dev.	Min	Max
wage	252	4.588	2.529	0.530	21.630
educ	252	12.317	2.473	0	18
exper	252	16.429	13.653	1	50
tenure	252	3.615	5.358	0	34

Statistic	N	Mean	St. Dev.	Min	Max
wage	274	7.099	4.161	1.500	24.980
educ	274	12.788	3.003	2	18
exper	274	17.558	13.500	1	51
tenure	274	6.474	8.369	0	44

If we assume that wage is normally distributed and that female wage is independent of male wage, we can test if the variance of wage is equal between the two groups. First, using the dataset `wage1` with `dround=3` to match the precision of the table:

```
wF <- subset(wage1, female==1)$wage
wM <- subset(wage1, female==0)$wage
testDiffVar(wF, wM, dround=3)
```

Testing $H_0 : \sigma_1^2 = \sigma_2^2$, against $H_1 : \sigma_1^2 \neq \sigma_2^2$ at 5%

$$test = \frac{\hat{\sigma}_1^2}{\hat{\sigma}_2^2} = \frac{17.3139}{6.3958} = 2.7071 \sim F(273, 251)$$

Since $2.7071 > 1.2267$ (the 95% quantile of the $F(273, 251)$), we reject H_0 .

We can also enter the information manually and change the size of the test:

```
testDiffVar(se=c(2.529, 4.161), n=c(252, 274), size=0.1)
```

Testing $H_0 : \sigma_1^2 = \sigma_2^2$, against $H_1 : \sigma_1^2 \neq \sigma_2^2$ at 10%

$$test = \frac{\hat{\sigma}_1^2}{\hat{\sigma}_2^2} = \frac{17.3139}{6.3958} = 2.7071 \sim F(273, 251)$$

Since $2.7071 > 1.1725$ (the 90% quantile of the $F(273, 251)$), we reject H_0 .

3.2 Introduction to Econometrics

In this section, we present solution generators for inference based on linear regressions. The general model is:

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_k X_k + u$$

By default, the estimated covariance matrix used to perform tests and construct confidence intervals is the one obtained by `vcov`. When a robust estimate is needed, it has to be obtained outside the solution generator. As for inference on the means, the normality assumption implies exact tests using either the t or the F distribution. Without this assumption, the asymptotic standard normal or chi-square distribution is used.

3.2.1 Tests on a single regression coefficient

Consider the following estimated model:

```
fit <- lm(wage~educ+exper, wage1)
printReg(fit)
```

$$\widehat{wage} = \underset{(0.7666)}{-3.3905} + \underset{(0.0538)}{0.6443} educ + \underset{(0.0110)}{0.0701} exper$$

$$n = 526, \quad R^2 = 0.2252, SSR = 5548.16$$

If we want to test the hypothesis $H_0 : \beta_1 = 0.5$ against the alternative $H_1 : \beta_1 \neq 0.5$ with the assumption that the errors are conditionally normal and homoskedastic, we can run the following. Note that β_1 is the second coefficient, so we set `which` to 2. Also, to generate a solution used the rounded coefficient and standard error shown in the above regression, we set `dround` to 4:

```
testRegCoef(fit, which=2, h0=0.5, dround=4)
```

Testing $H_0 : \beta_1 = 0.5$ against $H_1 : \beta_1 \neq 0.5$ at 5%

$$test = \frac{(\beta_1 - 0.5)}{se} = \frac{(0.6443 - 0.5)}{0.0538} = 2.6822 \sim t_{523}$$

Since $|2.6822| > 1.9645$ (the 97.5% quantile of the t_{523}), we reject H_0 .

We can also obtain the solution without the estimated model, by setting the coefficient estimate, its standard error, and sample size and the number of coefficients to the values we see in the printed equation:

```
testRegCoef(h0=0.5, beta=0.6443, se=0.0538, n=526, ncoef=3)
```

Testing $H_0 : \beta_j = 0.5$ against $H_1 : \beta_j \neq 0.5$ at 5%

$$test = \frac{(\beta_j - 0.5)}{se} = \frac{(0.6443 - 0.5)}{0.0538} = 2.6822 \approx N(0, 1)$$

Since $|2.6822| > 1.96$ (the 97.5% quantile of the $N(0, 1)$), we reject H_0 . (The $N(0, 1)$ is an approximation based on the C.L.T. because either the distribution of the error term is unknown and/or the errors are not homoskedastic.)

If we relax the normality assumption, the critical value is obtained using the $N(0, 1)$. If a vector of standard errors is provided, the solution generator assumes that the ones obtained with `vcov` is considered non-valid. In this case, there are no exact distribution, so the $N(0, 1)$ is also used. In the following the test is performed using the HCCM standard errors. For a solution based on a printed regression, we have to print the regression with the robust standard errors.

```
rse <- sqrt(diag(vcovHC(fit)))
printReg(fit, se=rse)
```

$$\widehat{wage} = \underset{(0.8748)}{-3.3905} + \underset{(0.0659)}{0.6443} educ + \underset{(0.0111)}{0.0701} exper$$

$$n = 526, R^2 = 0.2252, SSR = 5548.16 \text{ (Robust S-E)}$$

In the following test, the $N(0, 1)$ is used instead of the t distribution, because it is assume that the homoskedasticity assumption does not hold.

```
testRegCoef(fit, which=2, h0=0.5, dround=4, se=rse)
```

Testing $H_0 : \beta_1 = 0.5$ against $H_1 : \beta_1 \neq 0.5$ at 5%

$$test = \frac{(\beta_1 - 0.5)}{se} = \frac{(0.6443 - 0.5)}{0.0659} = 2.1897 \approx N(0, 1)$$

Since $|2.1897| > 1.96$ (the 97.5% quantile of the $N(0, 1)$), we reject H_0 . (The $N(0, 1)$ is an approximation based on the C.L.T. because either the distribution of the error term is unknown and/or the errors are not homoskedastic.)

As for tests on the mean, we can change the size with the argument `size` and the alternative with the argument `alter`.

3.2.2 Confidence interval on a regression coefficient

The same rule applies for confidence intervals. It is exact (using the t distribution) only if we assume normality and do not provide alternative standard errors. For example, we can obtain the 90% confidence interval for the coefficient of `bdrms` in the following regression as follows. First the estimated model is printed:

```
fit <- lm(log(price)~bdrms+log(lotsize)+log(sqrft), hprice1)
printReg(fit)
```

$$\widehat{\log(price)} = \underset{(0.1049)}{4.7003} + \underset{(0.0275)}{0.0370} bdrms + \underset{(0.0383)}{0.1680} \log(lotsize) + \underset{(0.0929)}{0.7002} \log(sqrft)$$

$$n = 88, R^2 = 0.643, SSR = 2.8626$$

Then we construct the interval:

```
ciRegCoef(fit, which=2, size=0.10, dround=4)
```

90% confidence interval for β_1 .

$$\begin{aligned} CI &= [\beta_1 - t^*se, \beta_1 + t^*se] \\ &= [0.037 - 1.6632 \times 0.0275, 0.037 + 1.6632 \times 0.0275] \\ &= [-0.0087, 0.0827] \end{aligned}$$

The t^* is the 95% quantile of the t-distribution with 84 degrees of freedom.

To get confidence interval with robust standard errors or non-normality, we just change `se` and `assume` as for the solution of hypothesis tests.

3.2.3 Testing linear combinations of coefficients (direct-t)

Any test on linear combinations of coefficients can be written as $H_0 : R\beta = q$. In this section we consider single hypotheses, so R has one row and q is a scalar. For the function `testLinRegCoefs`, R can be a plain vector or a matrix. The only requirement is to have a length equal to the number of coefficients. Consider the following regression, followed by the estimated covariance matrix:

```
dat <- subset(wage1, female==1 & married==1 & smsa==1)

fit <- lm(wage~educ+exper+tenure, dat)
printReg(fit, digits=5)
```

$$\begin{aligned} \widehat{wage} &= 1.73810 + 0.24617 educ - 0.00470 exper - 0.00235 tenure \\ &\quad (1.37570) \quad (0.10038) \quad (0.01712) \quad (0.03666) \\ n &= 88, \quad R^2 = 0.07187, SSR = 273.8005 \end{aligned}$$

```
knitr::kable(round(vcov(fit),5))
```

	(Intercept)	educ	exper	tenure
(Intercept)	1.89256	-0.13345	-0.00777	0.00117
educ	-0.13345	0.01008	0.00026	-0.00015
exper	-0.00777	0.00026	0.00029	-0.00027
tenure	0.00117	-0.00015	-0.00027	0.00134

If we want to generate a reproducible solution, we need to set `dround` to 5, which is the number of digits shown in the regression and covariance matrix. In the following, we test the hypothesis $\beta_1 - 10\beta_2 = 0.6$, which implies $R = \{0, 1, -10, 0\}$ and $q = 0.6$.

```
R <- c(0,1,-10,0)
q <- 0.6
testLinRegCoefs(object=fit, R=R, q=q, dround=5, alter="less")
```

Testing $H_0 : \beta_1 - 10\beta_2 = 0.6$ against $H_1 : \beta_1 - 10\beta_2 < 0.6$ at 5%

$$test = \frac{\hat{\beta}_1 - 10\hat{\beta}_2 - 0.6}{se} = \frac{-0.3068}{0.1841} = -1.667 \sim t_{84}$$

where,

$$se = \sqrt{Var(\hat{\beta}_1) + 10^2 Var(\hat{\beta}_2) - 2 \times 10 Cov(\hat{\beta}_1, \hat{\beta}_2)}$$

Since $-1.667 < -1.6632$ (the 5% quantile of the t_{84}), we reject H_0 .

If we set `assume` to "nonNormal" , or input an alternative covariance matrix, the $N(0,1)$ will be used to compute the critical values.

```
testLinRegCoefs(object=fit, R=R, q=q, dround=5, alter="greater",
                vcoef=vcovHC(fit))
```

Testing $H_0 : \beta_1 - 10\beta_2 = 0.6$ against $H_1 : \beta_1 - 10\beta_2 > 0.6$ at 5%

$$test = \frac{\hat{\beta}_1 - 10\hat{\beta}_2 - 0.6}{se} = \frac{-0.3068}{0.1667} = -1.8406 \approx N(0, 1)$$

where,

$$se = \sqrt{Var(\hat{\beta}_1) + 10^2 Var(\hat{\beta}_2) - 2 \times 10 Cov(\hat{\beta}_1, \hat{\beta}_2)}$$

Since $-1.8406 < 1.6449$ (the 95% quantile of the $N(0, 1)$), we do not reject H_0 . (The $N(0,1)$ is an approximation based on the C.L.T. because either the distribution of the error term is unknown and/or the errors are not homoskedastic.)

It is also possible to generate a solution with specific covariance matrix and coefficient vector. In this case, we need to specify if we assume normality or not, in order to base the test on the t distribution or the $N(0,1)$. The following is the same solution without the `lm` object, but with the t-distribution:

```
v <- round(vcovHC(fit),5)
b <- round(coef(fit),5)
testLinRegCoefs(R=R, q=q, alter="greater", vcoef=v, beta=b, n=nobs(fit))
```

Testing $H_0 : \beta_1 - 10\beta_2 = 0.6$ against $H_1 : \beta_1 - 10\beta_2 > 0.6$ at 5%

$$test = \frac{\hat{\beta}_1 - 10\hat{\beta}_2 - 0.6}{se} = \frac{-0.3068}{0.1667} = -1.8406 \approx N(0, 1)$$

where,

$$se = \sqrt{Var(\hat{\beta}_1) + 10^2 Var(\hat{\beta}_2) - 2 \times 10 Cov(\hat{\beta}_1, \hat{\beta}_2)}$$

Since $-1.8406 < 1.6449$ (the 95% quantile of the $N(0, 1)$), we do not reject H_0 . (The $N(0,1)$ is an approximation based on the C.L.T. because either the distribution of the error term is unknown and/or the errors are not homoskedastic.)

3.2.4 Testing multiple restrictions with the F test

The solution generator `testRegF` is for testing joint hypotheses under the homoskedasticity assumption. I don't ask questions about the non-homoskedastic case in exams, because it involves matrix calculations that are not suite for questions to be answered in short time. That's why this case is not covered. The test under normal errors is:

$$\frac{(SSR_r - SSR_u)/q}{SSR_u/(n - k - 1)} \sim F(q, n - k - 1),$$

where q is the number of restrictions being tested, SSR_u is the unrestricted SSR and SSR_r is the restricted SSR. If we drop the normality assumption, we use the following approximation based on the C.L.T.:

$$\frac{(SSR_r - SSR_u)}{SSR_u/(n - k - 1)} \approx \chi_q^2$$

We do not present solutions based on the restricted and unrestricted R^2 , because its validity depends on left-hand side of the regression being the same in the restricted and unrestricted models. However, the

solution of the joint significance of all coefficients but the intercept is based on the R^2 of the unrestricted model. We have the following with and without the normality assumption:

$$\frac{R^2/(k-1)}{(1-R^2)/(n-k-1)} \sim F(k-1, n-k-1)$$

$$\frac{R^2}{(1-R^2)/(n-k-1)} \approx \chi_{k-1}^2$$

Consider the following models:

```
data(hprice3, package="wooldridge")
hprice3 <- subset(hprice3, year==1978)
hprice3$price <- hprice3$price/1000
fitu <- lm(price~age+I(age^2)+rooms*baths+area+land, hprice3)
fitr <- lm(price~age+rooms+baths+area+land, hprice3)
printReg(fitu, maxpl=3, digits=5)
```

$$\begin{aligned} \widehat{price} = & 39.94059 - 0.56601 \text{ age} + 0.00236 I(\text{age}^2) - 2.34911 \text{ rooms} \\ & \quad (27.53351) \quad (0.13627) \quad (0.00084) \quad (4.28837) \\ & - 14.43260 \text{ baths} + 0.02322 \text{ area} + 0.00011 \text{ land} \\ & \quad (12.78385) \quad (0.00323) \quad (0.00003) \\ & + 2.68518 \text{ rooms} : \text{baths} \\ & \quad (1.89571) \\ n = 179, \quad R^2 = 0.62298, SSR = 62948.16 \end{aligned}$$

```
printReg(fitr, maxpl=3, digits=5)
```

$$\begin{aligned} \widehat{price} = & -13.58579 - 0.18786 \text{ age} + 4.15752 \text{ rooms} + 6.15880 \text{ baths} \\ & \quad (10.42160) \quad (0.04365) \quad (1.99856) \quad (3.18275) \\ & + 0.02374 \text{ area} + 0.00012 \text{ land} \\ & \quad (0.00329) \quad (0.00003) \\ n = 179, \quad R^2 = 0.59966, SSR = 66840.94 \end{aligned}$$

We want to test the significance of all coefficients but the intercept under normality. We can either provide the `lm` object of the unrestricted model, or give the values of the R^2 , sample size and the number of coefficients:

```
testRegF(fitu, dround=5)
```

Testing joint hypotheses using the F-test

$$test = \frac{R^2/7}{(1-R^2)/171} = \frac{0.623/7}{(1-0.623)/171} = 40.3653 \sim F(7, 171)$$

Since $40.3653 > 2.0635$ (the 95% quantile of the $F(7, 171)$), we reject H_0 .

```
testRegF(Rsq=0.62298, n=179, ncoef=8)
```

Testing joint hypotheses using the F-test

$$test = \frac{R^2/7}{(1-R^2)/171} = \frac{0.623/7}{(1-0.623)/171} = 40.3653 \sim F(7, 171)$$

Since $40.3653 > 2.0635$ (the 95% quantile of the $F(7, 171)$), we reject H_0 .

If we want to test the second model against the first, which corresponds to $H_0 : \beta_2 = \beta_7 = 0$, we proceed as follows (`dround` is set to 2, because only two decimals of the SSR's are printed). To see what happens if we do not assume normality, we set `assume` to `"nonNormal"`.

```
testRegF(fitu, fitr, dround=2, assume="nonNormal")
```

Testing joint hypotheses using the F-test

$$test = \frac{(SSR_r - SSR_{ur})}{SSR_{ur}/171} = \frac{(66840.94 - 62948.16)}{62948.16/171} = 10.5748 \approx \chi_2^2$$

Since $10.5748 > 5.9915$ (the 95% quantile of the χ_2^2), we reject H_0 . (The χ_2^2 is an approximation based on the C.L.T. because the distribution of the error term is unknown.

It is also possible to just provide the values of the SSR's and change the size:

```
testRegF(assume="Normal", size=0.1, SSRu=62948.16, SSRr=66840.94,
        n=179, nrest=2, ncoef=8)
```

Testing joint hypotheses using the F-test

$$test = \frac{(SSR_r - SSR_{ur})/2}{SSR_{ur}/171} = \frac{(66840.94 - 62948.16)/2}{62948.16/171} = 5.2874 \sim F(2, 171)$$

Since $5.2874 > 2.3339$ (the 90% quantile of the $F(2, 171)$), we reject H_0 .

4 Miscellaneous

In this section, I present some functions I used in the course as tools for teaching or to generate random assignments.

4.1 Probability using a normal table

Let X be distributed as $N(\mu, \sigma^2)$. In intro to statistics courses, students have to compute $Pr(X \leq a)$, $Pr(X \geq a)$ or $Pr(a \leq X \leq b)$ using a standardized normal table. The function `probFromTable` explains using a latex equation, how to compute the probability. The main arguments are `a`, `b` (if we want the last probability), `type` (`"less"` or `"greater"` for the first two probabilities), `mu` ($E(X)$), `sig2` ($Var(X)$), `digits` (the number of digits shown in the normal table), `varName` (the name of the random variable to be printed) and `approx` (is the normal an approximation). By default, the function assumes that the normal table shows $Pr(0 \leq Z \leq z)$. Here is a few examples:

- $Pr(X \leq 8)$ for $X \sim N(10, 2)$

```
probFromTable(a=8, mu=10, sig2=2)
```

$$\begin{aligned} Pr(X \leq 8) &= Pr\left(Z \leq \frac{8 - 10}{\sqrt{2}}\right) \\ &\approx Pr(Z \leq -1.41) \\ &= 0.5 - Pr(-1.41 \leq Z \leq 0) \\ &= 0.5 - Pr(0 \leq Z \leq 1.41) \\ &\approx 0.5 - 0.421 \\ &= 0.079 \end{aligned}$$

When generating a solution from exams or problem sets, we may be interested in the numeric solution without the equation. This is done by setting the argument `print` to `FALSE`:

```
probFromTable(a=8, mu=10, sig2=2, print=FALSE)
```

```
## [1] 0.079
```

- $Pr(X \geq 11)$ for $X \sim N(10, 2)$

```
probFromTable(a=11, type="greater", mu=10, sig2=2)
```

$$\begin{aligned} Pr(X \geq 11) &= Pr\left(Z \geq \frac{11-10}{\sqrt{2}}\right) \\ &\approx Pr(Z \geq 0.71) \\ &= 0.5 - Pr(0 \leq Z \leq 0.71) \\ &\approx 0.5 - 0.261 \\ &= 0.239 \end{aligned}$$

- $Pr(40 \leq Y \leq 48)$ for $Y \sim Binomial(100, 0.5)$, which can be approximated by a $N(50, 25)$. Note that `a` does not have to be the smallest number.

```
probFromTable(a=48, b=40, mu=50, sig2=25, approx=TRUE, varName="Y")
```

$$\begin{aligned} Pr(40 \leq Y \leq 48) &\approx Pr\left(\frac{40-50}{\sqrt{25}} \leq Z \leq \frac{48-50}{\sqrt{25}}\right) \\ &= Pr(-2 \leq Z \leq -0.4) \\ &= Pr(-2 \leq Z \leq 0) - Pr(-0.4 \leq Z \leq 0) \\ &= Pr(0 \leq Z \leq 2) - Pr(0 \leq Z \leq 0.4) \\ &\approx 0.477 - 0.155 \\ &= 0.322 \end{aligned}$$

In some textbooks, the normal table shows the CDF instead of $Pr(0 \leq Z \leq z)$. If that's the case, we can set the argument `CDF` to `TRUE`:

```
probFromTable(a=48, b=40, mu=50, sig2=25, approx=TRUE, varName="Y", CDF=TRUE)
```

$$\begin{aligned} Pr(40 \leq Y \leq 48) &\approx Pr\left(\frac{40-50}{\sqrt{25}} \leq Z \leq \frac{48-50}{\sqrt{25}}\right) \\ &= Pr(-2 \leq Z \leq -0.4) \\ &= Pr(0.4 \leq Z \leq 2) \\ &= Pr(Z \leq 2) - Pr(Z \leq 0.4) \\ &\approx 0.977 - 0.655 \\ &= 0.322 \end{aligned}$$

4.2 Generate statistics tables

For exercises or exam questions that require critical values from statistics table, the function `distTable` created a Latex table using the `xtable` package. For the t-distribution, it returns a table with 5 columns: 90%, 95%, 97.5%, 99% and 99.5% quantiles. We can use the default degrees of freedom, or enter them

manually. By default, it is a floating table with a title a the label "tTable" for referencing. In the following, I set float to FALSE:

```
distTable(df=c(4:9, Inf), dist="t", float=FALSE, digits=3)
```

df	0.900	0.950	0.975	0.990	0.995
4	1.533	2.132	2.776	3.747	4.604
5	1.476	2.015	2.571	3.365	4.032
6	1.440	1.943	2.447	3.143	3.707
7	1.415	1.895	2.365	2.998	3.499
8	1.397	1.860	2.306	2.896	3.355
9	1.383	1.833	2.262	2.821	3.250
Inf	1.282	1.645	1.960	2.326	2.576

Example: $P(t \leq x) = 0.975$, where t is t_5 implies $x = 2.571$

Example: $P(t \leq x) = 0.99$, where t is t_7 implies $x = 2.998$

Example: $P(t \leq x) = 0.01$, where t is t_7 implies $x = -2.998$

For the chi-square table, the lower quantiles are also printed:

```
distTable(df=c(4:10), dist="chisq", float=FALSE, digits=3)
```

df	0.010	0.025	0.050	0.100	0.900	0.950	0.975	0.990
4	0.297	0.484	0.711	1.064	7.779	9.488	11.143	13.277
5	0.554	0.831	1.145	1.610	9.236	11.070	12.833	15.086
6	0.872	1.237	1.635	2.204	10.645	12.592	14.449	16.812
7	1.239	1.690	2.167	2.833	12.017	14.067	16.013	18.475
8	1.646	2.180	2.733	3.490	13.362	15.507	17.535	20.090
9	2.088	2.700	3.325	4.168	14.684	16.919	19.023	21.666
10	2.558	3.247	3.940	4.865	15.987	18.307	20.483	23.209

Example: x that solves $P(Q \leq x) = 0.95$, where Q is χ^2_5 , is $x = 11.07$

Example: x that solves $P(Q \leq x) = 0.10$, where Q is χ^2_7 , is $x = 2.83$

The F-distribution has two sets of degrees of freedom, so a table only contains one quantile at the time. We can change the default 95% by setting the argument size to something other than 0.05:

```
distTable(df=c(4:10), df2=1:6, dist="f", size=0.10, float=FALSE, digits=3)
```

df2\df1	4	5	6	7	8	9	10
1	55.833	57.240	58.204	58.906	59.439	59.858	60.195
2	9.243	9.293	9.326	9.349	9.367	9.381	9.392
3	5.343	5.309	5.285	5.266	5.252	5.240	5.230
4	4.107	4.051	4.010	3.979	3.955	3.936	3.920
5	3.520	3.453	3.405	3.368	3.339	3.316	3.297
6	3.181	3.108	3.055	3.014	2.983	2.958	2.937

Example: The 90% quantile of F where F is $F(4, 3)$ is 5.34

Example: The 90% quantile of F where F is $F(6, 4)$ is 4.01

The last option is the Normal table. By default, it shows $Pr(0 \leq Z \leq z)$ for $Z \sim N(0, 1)$. To have the CDF, we set CDF to TRUE.

```
distTable(df=c(4:10), df2=1:6, dist="normal", CDF=TRUE, float=FALSE, digits=3)
```

z	0	1	2	3	4	5	6	7	8	9
0.0	0.500	0.504	0.508	0.512	0.516	0.520	0.524	0.528	0.532	0.536
0.1	0.540	0.544	0.548	0.552	0.556	0.560	0.564	0.567	0.571	0.575
0.2	0.579	0.583	0.587	0.591	0.595	0.599	0.603	0.606	0.610	0.614
0.3	0.618	0.622	0.626	0.629	0.633	0.637	0.641	0.644	0.648	0.652
0.4	0.655	0.659	0.663	0.666	0.670	0.674	0.677	0.681	0.684	0.688
0.5	0.691	0.695	0.698	0.702	0.705	0.709	0.712	0.716	0.719	0.722
0.6	0.726	0.729	0.732	0.736	0.739	0.742	0.745	0.749	0.752	0.755
0.7	0.758	0.761	0.764	0.767	0.770	0.773	0.776	0.779	0.782	0.785
0.8	0.788	0.791	0.794	0.797	0.800	0.802	0.805	0.808	0.811	0.813
0.9	0.816	0.819	0.821	0.824	0.826	0.829	0.831	0.834	0.836	0.839
1.0	0.841	0.844	0.846	0.848	0.851	0.853	0.855	0.858	0.860	0.862
1.1	0.864	0.867	0.869	0.871	0.873	0.875	0.877	0.879	0.881	0.883
1.2	0.885	0.887	0.889	0.891	0.893	0.894	0.896	0.898	0.900	0.901
1.3	0.903	0.905	0.907	0.908	0.910	0.911	0.913	0.915	0.916	0.918
1.4	0.919	0.921	0.922	0.924	0.925	0.926	0.928	0.929	0.931	0.932
1.5	0.933	0.934	0.936	0.937	0.938	0.939	0.941	0.942	0.943	0.944
1.6	0.945	0.946	0.947	0.948	0.949	0.951	0.952	0.953	0.954	0.954
1.7	0.955	0.956	0.957	0.958	0.959	0.960	0.961	0.962	0.962	0.963
1.8	0.964	0.965	0.966	0.966	0.967	0.968	0.969	0.969	0.970	0.971
1.9	0.971	0.972	0.973	0.973	0.974	0.974	0.975	0.976	0.976	0.977
2.0	0.977	0.978	0.978	0.979	0.979	0.980	0.980	0.981	0.981	0.982
2.1	0.982	0.983	0.983	0.983	0.984	0.984	0.985	0.985	0.985	0.986
2.2	0.986	0.986	0.987	0.987	0.987	0.988	0.988	0.988	0.989	0.989
2.3	0.989	0.990	0.990	0.990	0.990	0.991	0.991	0.991	0.991	0.992
2.4	0.992	0.992	0.992	0.992	0.993	0.993	0.993	0.993	0.993	0.994
2.5	0.994	0.994	0.994	0.994	0.994	0.995	0.995	0.995	0.995	0.995
2.6	0.995	0.995	0.996	0.996	0.996	0.996	0.996	0.996	0.996	0.996
2.7	0.997	0.997	0.997	0.997	0.997	0.997	0.997	0.997	0.997	0.997
2.8	0.997	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.998
2.9	0.998	0.998	0.998	0.998	0.998	0.998	0.998	0.999	0.999	0.999
3.0	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999	0.999


Example 1: If $Z \sim N(0, 1)$, then $P(0 \leq Z \leq 1.84) = 0.467$

Example 2: $P(-1.54 \leq Z \leq 0) = P(0 \leq Z \leq 1.54) = 0.438$

4.3 Generate time series


The function `genTS` generates an `mts` object with two time series, `X` and `Y`. One of the series is trend stationary (TS) and the second one is a unit root process with possibly a constant or a linear drift. The purpose of this dataset is to identify the process for each series. The following is an example with all possible arguments, which are all optional.

```
dat <- genTS(n=200, frequency=4, start=1970, seed=112345)
plot(dat[, "X"], ylab="X")
plot(dat[, "Y"], ylab="Y")
```



Since we know the possible processes, it is not too hard to determine the process of each series:

```
fitX <- lm(X~time(X) + I(time(X)^2), dat)
Xdt <- ts(residuals(fitX), start=start(dat), frequency=frequency(dat))
fitY <- lm(Y~time(Y) + I(time(Y)^2), dat)
Ydt <- ts(residuals(fitY), start=start(dat), frequency=frequency(dat))
acf(Xdt, main="ACF of the detrended X")
acf(Ydt, main="ACF of the detrended Y")
```



We would conclude that since the detrended Y is highly persistent, it is probably a unit root process, and X would be a trend stationary series. The true processes are saved in the attribute "which":

```
attr(dat, "which")
```

```
## [1] "X is a TS"          "Y is a Unit Root"
```

4.4 Generate supply and demand functions

This tool is to illustrate the identification problem when estimating demand and supply equations. The main function is `genDemSup`, which generates a dataset that contains price, quantity, exogenous variables and errors. The system is:

$$\begin{aligned} Q^d &= d_1 + d_2 P + d_3 X + \eta \\ Q^s &= a_1 + a_2 P + a_3 Z + e \end{aligned}$$

The function has the following arguments:

- **n**: The sample size
- **a**: The vector of supply coefficients.
- **d**: The vector of demand coefficients.
- **sige**, **sigeta**, **sigX**, **sigZ**: The variances of e , η , X and Z .

It returns an object of class `DemSup`, with a `print` method that generates the equations in Latex format, and a `plot` method that plots the data points and a few demand and supply curves with their equilibrium points. The coefficients must be such that an equilibrium exists, which is the case of the default values. The following shows the system of equation with by default:

```
set.seed(1223345)
sys <- genDemSup(200)
print(sys)
```

$$\begin{aligned} Q^d &= 100 - 1P + 2X + \eta \\ Q^s &= 20 + 1P - 2Z + e \end{aligned}$$

The element `dat` of the object is a `data.frame` that contains Q , P , X , Z , e and η (as `eta`). We can therefore try to estimate the demand function by OLS:

```
fit <- lm(Q~P+X, sys$dat)
printReg(fit)
```


$$\begin{aligned} \hat{Q} &= 67.2620 - 0.3868 P + 1.5411 X \\ &\quad (3.1138) \quad (0.0599) \quad (0.1114) \\ n &= 200, \quad R^2 = 0.4955, SSR = 1804.158 \end{aligned}$$

We could then plot some demand and supply with the fitted line to illustrate the identification problem. The arguments of `plot` are

- **x**, the `DemSup` object
- **nCurves**: The number of demand and supply curves to plot. The default is all curves.
- **nPoints**: The number of data points to plot. The default is all points.
- **fac**: This is a vector of 2. It determines the range for each demand and supply curve plotted. The price range is $P^* \times \text{fac}[1]$ to $P^* \times \text{fac}[2]$, where P^* is the equilibrium price. By default it goes from $0.9P^*$ to $1.1P^*$.

In the following, we plot the fitted line when $X = \bar{X}$:

```
plot(sys, nCurves=10)
pr <- predict(fit, newdata=data.frame(P=50:70, X=mean(sys$dat$X)))
lines(50:70, pr, col=2, lty=2, lwd=2)
```



We can use the data to analyze the exogeneity condition, since we are given the vector of true errors. For example, if we want to estimate the demand, we need to show that Z is exogenous. It can be done by regressing Z on both errors and check the significance of the f-test:

```
fit2 <- lm(Z~e+eta, sys$dat)
printReg(fit2)
```

$$\hat{Z} = \underset{(0.1491)}{10.0558} + \underset{(0.0358)}{0.0042} e + \underset{(0.0392)}{0.0665} \eta$$

$$n = 200, \quad R^2 = 0.0145, \quad SSR = 847.6275$$

```
testRegF(fit2)
```

Testing joint hypotheses using the F-test

$$test = \frac{R^2/2}{(1-R^2)/197} = \frac{0.0145/2}{(1-0.0145)/197} = 1.4446 \sim F(2, 197)$$

Since $1.4446 < 3.0418$ (the 95% quantile of the $F(2, 197)$), we do not reject H_0 .

We therefore, conclude that Z is exogenous, because it cannot be explained by the error terms. We can also look at the relevance of the instrument by estimating the reduced form equation and test the significance of the coefficient of Z :

```
fit2 <- lm(P~Z+X, sys$dat)
printReg(fit2, stars=TRUE)
```

$$\hat{P} = \underset{(1.3012)^{***}}{38.2509} + \underset{(0.0945)^{***}}{1.1007} Z + \underset{(0.0849)^{***}}{1.0587} X$$

$$n = 200, \quad R^2 = 0.5904, \quad SSR = 1512.813$$

*pv < 0.1; **pv < 0.05; ***pv < 0.01


Since the coefficient of Z is highly significant, we conclude that the instrument is relevant. We can then estimate the slope by two-stage least squares and plot the fitted line:

```
sys$dat$Phat <- fitted(fit2)
fit3 <- lm(Q~Phat+X, sys$dat)
printReg(fit3)
```

$$\hat{Q} = 87.6695 - \frac{0.7988}{(4.3537)} Phat + \frac{1.9669}{(0.1231)} X$$


$n = 200, R^2 = 0.5745, SSR = 1521.729$

```
plot(sys, nCurves=10)
lines(50:70, pr, col=2, lty=2, lwd=2)
pr2 <- predict(fit3, newdata=data.frame(Phat=50:70, X=mean(sys$dat$X)))
lines(50:70, pr2, col="darkgreen", lty=2, lwd=2)
legend("topleft", c("OLS", "TSLS"), col=c(2, "darkgreen"), lty=2, lwd=2, bty='n',
      cex=.7)
```



The parameter values can be used to illustrate different aspects of the simultaneity problem. For example, we can show that the demand is identified only if the supply shifts ($\text{sigeta}=0$ and $\text{sigX}=0$)

```
set.seed(1122345)
sys <- genDemSup(200, sigeta=0, sigX=0)
plot(sys, nCurves=10)
```



We can also illustrate the weak instrument problems. The `tsls` function is from the `momentfit` package (Chausse (2025)) and the `texreg` function from the `texreg` package (Leifeld (2013)). The `extract` method needed to run `texreg` on `tsls` objects is included in the `metricsUW` package.

```
library(texreg)
library(momentfit)
sysS <- genDemSup(200, a=c(20,1,-2))
sysW <- genDemSup(200, a=c(20,1,0.01))
```

```
fitS <- tsls(Q~P+X, ~X+Z, data=sysS$dat)
fitW <- tsls(Q~P+X, ~X+Z, data=sysW$dat)
texreg(list(Strong=fitS, Weak=fitW), digits=4, table=FALSE)
```

	Strong	Weak
(Intercept)	101.3751*** (6.7180)	98.0278 (53.2214)
P	-1.0647*** (0.1336)	-0.8910 (1.3073)
X	2.2572*** (0.1879)	1.6657 (1.2205)
J-test Statistics		
J-test p-value		
First Stage F-stats(P)	113.2684	0.9667

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$


We see that the weakness of the instrument (represented here by $a_3 = 0.01$ instead of $a_3 = -2$) increase the variance of the coefficients. The weakness is also clear from the first stage F-test. Similarly, we can see what happens if the instrument has a small variance by changing the value of `sigZ`.

The tool also offers the possibility of having an over-identified model. By default, we have `d=c(100, -1, 2, 0)` and `sigX=c(5,5)`, which implies that only one exogenous variable is included in the demand function. By replacing `d[4]` by a non-zero value, the demand becomes an equation with the exogenous variables X_1 and X_2 , and `sigX` contains the variance of the two variables:

```
sys2 <- genDemSup(n=200, d=c(100, -1, 2, 0.5))
print(sys2)
```

$$\begin{aligned} Q^d &= 100 - 1P + 2X_1 + 0.5X_2 + \eta \\ Q^s &= 20 + 1P - 2Z + e \end{aligned}$$

```
plot(sys2, nCurve=15)
```



We can test the strength of the instruments using the `testRegF` function:

```
fitu <- lm(P~X1+X2+Z, sys2$dat)
fitr <- lm(P~Z, sys2$dat)
printReg(fitu)
```


$$\hat{P} = \underset{(1.4353)}{39.3874} + \underset{(0.0813)}{1.0948} X1 + \underset{(0.0819)}{0.3879} X2 + \underset{(0.0874)}{0.8531} Z$$

$$n = 200, \quad R^2 = 0.6249, SSR = 1373.691$$

```
printReg(fitr)
```

$$\hat{P} = \underset{(1.2864)}{53.4196} + \underset{(0.1241)}{0.9523} Z$$

$$n = 200, \quad R^2 = 0.2292, SSR = 2822.91$$

```
testRegF(fitu, fitr)
```

Testing joint hypotheses using the F-test

$$test = \frac{(SSR_r - SSR_{ur})/2}{SSR_{ur}/196} = \frac{(2822.9104 - 1373.6915)/2}{1373.6915/196} = 103.3882 \sim F(2, 196)$$

Since $103.3882 > 3.042$ (the 95% quantile of the $F(2, 196)$), we reject H_0 .

We therefore have strong instruments. We can also test the exogeneity of the instruments, by using the `specTest` method of `momentfit` or the `printReg` method for objects of class `tsls`:

```
fit <- tsls(Q~P+Z, ~X1+X2+Z, data=sys2$dat)
printReg(fit, strength=TRUE, Jtest=TRUE, Exo=TRUE)
```

$$\hat{Q} = \underset{(5.8247)}{17.7942} + \underset{(0.1060)}{1.0022} P - \underset{(0.1666)}{1.8216} Z$$

TSLS: n = 200, Jtest = 0.0082(pv = 0.9279)

First Stage: P : [F(2, 196) = 103.3882, pv = 0]

Exogenous: $X1, X2, Z$

```
specTest(fit)
```

```
##
## J-Test
##           Statistics  df  pvalue
## Test E(g)=0:      0.0081975  1  0.92786
```

References

- Chausse, Pierre. 2025. *Momentfit: Methods of Moments*. University of Waterloo. <https://CRAN.R-project.org/package=momentfit>.
- Kleibner, Christian, and Achim Zeileis. 2008. *Applied Econometrics with R*. New York: Springer-Verlag. <https://doi.org/10.1007/978-0-387-77318-6>.
- Leifeld, Philip. 2013. “texreg: Conversion of Statistical Model Output in R to LaTeX and HTML Tables.” *Journal of Statistical Software* 55 (8): 1–24. <https://doi.org/10.18637/jss.v055.i08>.
- R Core Team. 2025. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. <https://www.R-project.org/>.
- Shea, Justin M. 2024. *Wooldridge: 115 Data Sets from "Introductory Econometrics: A Modern Approach, 7e" by Jeffrey m. Wooldridge*. <https://doi.org/10.32614/CRAN.package.wooldridge>.
- Zeileis, Achim. 2004. “Econometric Computing with HC and HAC Covariance Matrix Estimators.” *Journal of Statistical Software* 11 (10): 1–17. <https://doi.org/10.18637/jss.v011.i10>.