

# A Knitr Demo

Charles J. Geyer

February 10, 2017

## 1 Licence

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License <http://creativecommons.org/licenses/by-sa/4.0/>.

## 2 R

The version of R used to make this document is 3.3.2. The version of the `knitr` package used to make this document is 1.15.1. The version of the `xtable` package used to make this document is 1.8.2.

## 3 Introduction

This is a demo for using the R package `knitr`. To get started make a regular  $\LaTeX$  file (like this one) but give it the suffix `.Rnw` instead of `.tex` and then turn it into a PDF file (`bar.pdf`) with the (unix) commands

```
R CMD Sweave bar.Rnw
pdflatex bar
```

(as always with  $\LaTeX$  one may have to rerun the latter command multiple times until it stops complaining about needing to be rerun). For this to work R needs to have package `knitr` installed.

This works because of the comment

```
%\VignetteEngine{knitr::knitr}
```

(otherwise R CMD Sweave would expect Sweave rather than `knitr`).

Now include R in our document. Here's a simple example

```
2 + 2
## [1] 4
```

This is not  $\text{\LaTeX}$ . It is a “code chunk” processed by `knitr`. When `knitr` hits such a thing, it processes it, runs R to get the results, and stuffs (by default) the output in the  $\text{\LaTeX}$  file it is creating. The  $\text{\LaTeX}$  between code chunks is copied verbatim (except for `Sexpr`, about which see below). Hence to create an Rnw document you just write plain old  $\text{\LaTeX}$  interspersed with “code chunks” which are plain old R.

## 4 Plots

### 4.1 Make Up Data

Plots get a little more complicated. First we make something to plot (simulate regression data).

```
n <- 50
x <- seq(1, n)
a.true <- 3
b.true <- 1.5
y.true <- a.true + b.true * x
s.true <- 17.3
y <- y.true + s.true * rnorm(n)
out1 <- lm(y ~ x)
summary(out1)

##
## Call:
## lm(formula = y ~ x)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -40.191 -13.264   1.795  12.272  34.201
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)    4.402     5.509   0.799   0.428
## x              1.367     0.188   7.272 2.82e-09 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 19.18 on 48 degrees of freedom
## Multiple R-squared:  0.5242, Adjusted R-squared:  0.5143
## F-statistic: 52.88 on 1 and 48 DF,  p-value: 2.823e-09
```




Figure 1: Scatter Plot with Regression Line

## 4.2 Figure with Code to Make It Shown

Figure 1 (p. 3) is produced by the following code

```
plot(x, y)
abline(out1)
```

Note that `x`, `y`, and `out1` are remembered from the preceding code chunk. We don't have to regenerate them. All code chunks are part of one R "session".

This was a little tricky. We did this with two code chunks, one visible and one invisible. See the source file `bar.Rnw` to see how it is done. The first code chunk just shows the code (so we can see it above). It has the optional argument `eval=FALSE` which makes it not actually do anything. Then the second code

chunk (which we do not see) actually makes the plot. It has optional argument `echo=FALSE` which says it should not be shown (because it would be shown in the figure, which is not what we want).

The second code chunk does not repeat the code of the first chunk. Instead it “quotes” the first chunk, just reuses its code. This follows the DRY/SPOT rule (*don't repeat yourself* or *single point of truth*) so we only have one bit of code for generating the plot. What the reader sees is guaranteed to be the code that made the plot. If we had used cut-and-paste, just repeating the code, the duplicated code might get out of sync after edits.

So making a figure is a bit more complicated in some ways but much simpler in others. Note the following virtues

- The figure is guaranteed to be the one described by the text (at least by the R in the text).
- No messing around with sizing or rotations. It just works!

### 4.3 Figure with Code to Make It Not Shown

For this example we do a cubic regression on the same data.

```
out3 <- lm(y ~ x + I(x^2) + I(x^3))
summary(out3)

##
## Call:
## lm(formula = y ~ x + I(x^2) + I(x^3))
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -44.657 -12.344   2.284  12.990  30.406
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept) -14.398732  11.539098  -1.248   0.2184
## x             4.610886   1.940104   2.377   0.0217 *
## I(x^2)       -0.127628   0.087926  -1.452   0.1534
## I(x^3)         0.001404   0.001134   1.238   0.2219
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 18.89 on 46 degrees of freedom
## Multiple R-squared:  0.558, Adjusted R-squared:  0.5291
## F-statistic: 19.35 on 3 and 46 DF,  p-value: 2.925e-08
```

If you don't care to show the R code to make a figure, it is simpler still. Figure 2 (p. 5) shows the plot of the (cubic) regression function for these data.




Figure 2: Scatter Plot with Cubic Regression Curve

It is made by one code chunk with optional argument `echo=FALSE` so it makes a figure and does not show the code.

Also note that every time we rerun `knitr` Figures 1 and 2 change because the simulated data are random. Everything just works. This should tell you the main virtue of `knitr`. It's always correct. There is never a problem with stale cut-and-paste.

## 5 R in Text

This section illustrates the command `\Sexpr`, which allows running text to be computed by R.

We show some numbers calculated by R interspersed with text. The quadratic and cubic regression coefficients in the preceding regression (page 4) were  $\beta_2 = -0.1276$  and  $\beta_3 = 0.0014$ . Magic!

In order for your document to be truly reproducible, you must never cut-and-paste anything R computes. Always have R recompute it every time the document is processed, either in a code chunk or with `Sexpr`.

## 6 Tables

The same goes for tables. The `xtable` command is used to make tables. Here is a “table” of sorts in some R printout.

```
out2 <- lm(y ~ x + I(x^2))
anova(out1, out2, out3)

## Analysis of Variance Table
##
## Model 1: y ~ x
## Model 2: y ~ x + I(x^2)
## Model 3: y ~ x + I(x^2) + I(x^3)
##   Res.Df  RSS Df Sum of Sq    F Pr(>F)
## 1      48 17666
## 2      47 16958  1    707.66 1.9835 0.1657
## 3      46 16411  1    547.05 1.5334 0.2219
```

We want to turn that into a  $\text{\LaTeX}$  table. First we have to figure out what the output of the R function `anova` is and capture it so we can use it.

```
foo <- anova(out1, out2, out3)
class(foo)

## [1] "anova"      "data.frame"
```

So now we are ready to turn the data frame `foo` into Table 1 using the R functions `xtable` and `print.xtable` (that is, the `xtable` method of the generic

Table 1: ANOVA Table

	Res.Df	RSS	Df	Sum of Sq	F	Pr(>F)
1	48	17666.07				
2	47	16958.41	1	707.66	1.984	0.166
3	46	16411.36	1	547.05	1.533	0.222

function `print`), which are found in the R package `xtable`, which is a CRAN package.

The code chunk that makes Table 1 has optional arguments `echo=FALSE` and `results=tex` so we do not see it and the result is not put back in the code chunk but it just treated as ordinary L<sup>A</sup>T<sub>E</sub>X.

The reason why we had to use both `xtable` and `print.xtable` is that we needed to use arguments for both. We used the `caption`, `label`, and `digits` arguments of the former and the `table.placement` and `caption.placement` arguments of the latter.

## 7 Summary

`knitr` is terrific, so important that we cannot get along without it or its older competitor `Sweave` or its newer competitor `Rmarkdown`.

Its virtues are

- The numbers and graphics you report are actually what they are claimed to be.
- Your analysis is reproducible. Even years later, when you've completely forgotten what you did, the whole write-up, every single number or pixel in a plot is reproducible.
- Your analysis actually works—at least in this particular instance. The code you show actually executes without error.
- Toward the end of your work, with the write-up almost done you discover an error. Months of rework to do? No! Just fix the error and rerun `knitr` and `pdflatex`. One single problem like this and you will have all the time invested in `knitr` repaid.
- This methodology provides discipline. There's nothing that will make you clean up your code like the prospect of actually revealing it to the world.

Whether we're talking about homework, a consulting report, a textbook, or a research paper. If they involve computing and statistics, this is the way to do it.