

Quick Introduction to R For Vanderbilt PhD students

Richard Blissett

The following document is meant to be a quick introduction to R for Vanderbilt LPO Stata users. The examples in this document use the plans.dta dataset, which is included.

Contents

1	The What, Why, and How	3
1.1	What is R?	3
1.2	Why use R?	3
1.3	How do I get and use R?	3
1.3.1	Downloading R	3
1.3.2	RStudio Interface	3
1.3.3	What Do I Click?	4
1.4	R Resources	4
2	Conceptual Shifts	4
2.1	Object-Oriented Thinking	5
2.2	Multiple Datasets	5
2.3	Function Syntax	5
2.4	Libraries	5
3	Objects and Basic R	6
3.1	Assignment	6
3.2	Basic Math	7
3.3	Lists	7
4	Control Structures	8
4.1	If	8
4.2	While	8
4.3	For	8
5	Comments	9
6	Data Manipulation	9
6.1	Importing, Viewing, and Editing Data	9
6.2	Accessing and Manipulating Variables	9
6.3	Merging and Appending	11
7	Data Analysis	11
7.1	Summary Statistics	11
7.2	Regression	12
7.3	Survey Weighting for Estimates of Variance	13
7.4	Other Data Analysis Tools	14

8	Tables and Graphs	14
8.1	Tables	14
8.2	Graphs	15
9	Other R Things	18
9.1	Sweave/knitr	18
9.2	GIS	18
9.3	Functions	18
10	Translations	20

1 The What, Why, and How

1.1 What is R?

R is a programming language and environment used primarily for statistical computing. In terms of its general use, it can be used to do most of the things that Stata does.

1.2 Why use R?

There are benefits and drawbacks to choosing any programming language over another. Being an R tutorial, the bias here is towards R. Here are some reasons why:

- R is free.
- In R, you can hold multiple datasets open at the same time. This means that, while possible, merging is not as necessary.
- R is particularly good at producing high-quality graphics (see the R Graph Gallery online for examples).
- All R functions, as well as functions in most other programming languages, have the same basic syntactical structure. Stata is fairly unusual in its structure, which has main arguments on one side of a comma and options on the other, with frequent confusion by users about what goes where.
- R is highly extensible, meaning that it is very simple for users anywhere in the world to cook up solutions for problems they face and publish them for anyone to use.
- R can be used in conjunction with other languages fairly easily, including with markup languages like LaTeX (this document was created this way).
- The popularity of and interest in R has grown substantially in recent years.

That all being said, the R vs. Stata vs. any other programming language “debate” is about as clear and useful as the Mac vs. PC debate. In the end, both are good languages to know. If anything, knowing more than one language helps you communicate with other people, and understanding how multiple languages work can help you be more creative in the work you do.

1.3 How do I get and use R?

1.3.1 Downloading R

R can be downloaded from www.rproject.org, where you should click on “CRAN” on the side to download. Follow the on-screen instructions to install.

What you are downloading is the R language as well as basic software that you can use (like Stata’s graphical user interface). There are, however, other options (many, in fact). The one recommended for this tutorial and the one that is probably the most widely used is RStudio, available from www.rstudio.com. You should still download R from the CRAN first, and then download RStudio.

1.3.2 RStudio Interface

When you open RStudio, you should have four basic areas on your screen:

1. The top-left is where you type source code, like the script editor in Stata.
2. The bottom-left is the console, where you see what R is doing. Like in Stata, you can just type in the console next to the > character.
3. The upper-right is where you see your list of R objects and things in your workspace (explained in the next section). History is also there.
4. The bottom right is where you can see your file structure, graphical output, available packages, and help documents.

1.3.3 What Do I Click?

Console To work in the console, just click anywhere in the console and type away. For example, type “1” and press enter. You should see something like the following.

```
> 1
[1] 1
```

Very exciting, yes. For now, though, note that you did not, as you do in Stata, have to type `display` 1 to get that back. As long as you are working in this R environment, you won't have to type `display` to show things (this is not the case when running R from other scripts, but that is beyond the scope of this document).

Type “1+1” and see what happens. More excitement.

Also, remember that in Stata, to scroll through previous commands, you use the “Page Up” and “Page Down” keys. In the RStudio environment, the up and down arrow keys will do the same.

Scripts R code is written in the source window (upper-left) and saved as a `.r` file. (Note that RStudio can also edit other types of files as well.) Type and save code here as you would have in Stata.

To run a particular line or highlighted block of lines, click the “Run” button at the top of the window. To run the entire script, click “Source.” To run the entire script *and* see output as you do so, click “Source with Echo.”

1.4 R Resources

Getting help in R is easy. To get information about the syntax of a function, just type a “?” before the function name as follows.

```
> ?functionname
```

To search the help files for a certain word, just use double question marks.

```
> ??regression
```

In addition, the R community is huge, and there are many great resources that can give you ideas and help with how to use R. Some particularly useful ones are:

- R search engine (rseek.org)
- Quick R website (statmethods.net)
- O'Reilly *R Cookbook*
- R Graph Gallery (gallery.r-enthusiasts.com)
- R help mailing list (stat.ethz.ch/mailman/listinfo/r-help)
- *Complex Surveys: A Guide to Analysis Using R* by Thomas Lumley
- *Data Analysis Using Regression and Multilevel/Hierarchical Models* by Andrew Gelman and Jennifer Hill

2 Conceptual Shifts

Though Stata and R are similar in many ways, there are several important differences that should be understood before starting with Stata.

2.1 Object-Oriented Thinking

In R, everything is a “thing.” This may seem obvious, but it has important implications for how to program in R. This is very different from Stata, where you think in terms of actions instead. This is reflected in the languages themselves. For example, the `tabulate` command in Stata and the `summary()` command in R do roughly the same thing for variables in datasets, but note how the command is a verb in Stata and a noun in R. In fact, when you run `summary()` in R, you are actually creating a summary “object.”

In Stata, we tend to think of the environment being the data itself, with everything else in Stata being things we “do” to the data. Anything else you store is kept in locals, estimates, and scalars, which are typically temporary and are cleared when your code exits.

Imagine working in R like everything in front of you is meant to be stored in a local (except that in contrast to Stata, the local is not temporary). In R, the data is only a “thing” in the environment. If you want to run a regression, R essentially creates a regression object, which you can store as another “thing” in your workspace.

With this, though, this does mean that using R requires a real understanding of how information is stored. When regression estimates are stored in Stata, for example, we are often not particularly concerned about how the data is stored, just as long as we can recall it when necessary. Working creatively in R will require consideration of how this information is stored (an illustration of this is in the section on regression).

2.2 Multiple Datasets

It is because of the object-oriented nature of R that you are able to keep multiple datasets open at the same time. When you open a dataset, you store it in a data frame object. You can have as many objects as your memory allows. As noted before, this can occasionally preclude the need for merging.

This does mean, however, that you have to think of variables as belonging to different datasets, and most of the time, when using variables, you have to reference which dataset they come from using the `$` symbol. (This can be avoided by using the `attach()` function, which will be discussed later.) For example, the full name of the math score variable within the `plans` dataset would be `plans$byne1s2m`.

2.3 Function Syntax

All functions (commands) are written the same in R:

```
> functionname(part1, part2, part3)
```

On other words, every function, in following with our object-oriented thinking, is made up of “parts,” which include data and options. For example, a regression is made up of data, a formula, weights, etc. This is somewhat different than functional syntax in Stata, where “main parts” and “options” are separated by a comma, and how the parameters are included differs from function to function.

2.4 Libraries

R comes with a base set of functions, but there are many others that have been written by people to perform more advanced tasks. Like Stata’s `.ado` files, these are written by users and distributed for public use. In R, however, you must indicate when you are using an external package. This is done as follows:

```
> library(libraryname)
```

There are a number of packages that come with your R installation. You can download others by using the `install.packages()` command or by going to `Tools > Install Packages...` in RStudio. It is customary to place all `library()` commands at the top of the script.

3 Objects and Basic R

An R object is just information, and it is how information is stored in R. A data table is an object. A regression is an object.

Objects are created when the output from functions are stored within them. The rules for object names are similar to the rules for variable names in Stata. To remove an object, use the `rm()` command. The following command will clear all objects.

```
> rm(list=ls())
```

Above, the `rm()` function is being instructed to remove a list of objects. That list is obtained from the `ls()` command, which returns a list of all objects currently in your workspace.

3.1 Assignment

Assignment isn't such a big deal in Stata unless you work exclusively in locals, estimate storage, and scalars (which all are assigned differently). In R, however, assignment is crucial.

In R, type the number 2.

```
> 2
[1] 2
```

Note the [1] before the number. You can ignore that for now. It just indicates the line of output at which you are looking.

If you wanted that 2 to be stored, you would assign it to an object, say, `a`.

```
> a <- 2
```

You can also use the “=” sign to do assignment, but it is generally not recommended in R because of its limitations (not overly critical for early R users).

If you want to see what is in `a`, just type it.

```
> a
[1] 2
```

The same applies for functional output. For example, the `sqrt()` function takes the square root of a number. So...

```
> sqrt(a)
[1] 1.414214
> a
[1] 2
```

Note that the value stored in `a` did not change. The `sqrt()` function just did the operation and gave you the output on the screen. If you wanted to save that, you would do...

```
> b <- sqrt(a)
> b
[1] 1.414214
```

Here, the square root was stored in the `b` object, and not sent to the screen. Note that you also could have overwritten what was in `a` by assigning `sqrt(a)` to `a` itself. R will evaluate the expression to the right of `<-` and store the output in the object to the left.

If you wanted to see the output as well, just put parentheses around the command as follows.

```
> (b <- sqrt(a))
[1] 1.414214
```

3.2 Basic Math

All basic math is the same in R as it is in Stata (using the `+`, `-`, `*`, and `/` operators). You can use objects in math as well.

```
> a <- 2
> b <- 4
> 2+4
[1] 6
> a+b
[1] 6
```

Math follows the normal order of operations. Functions like `log()` and `sqrt()` also exist.

3.3 Lists

In R, to create a list of things, you use the `c()` function, which stands for concatenate.

```
> c(1, 2, 3, 4, 5)
[1] 1 2 3 4 5
```

This created a list, which you can see. But you didn't store it, so you can't call it back if you need that. You can store it and then see it the same way we just did above. List items can be accessed using bracket notation.¹

Also, note that the sequence operator is `:` (as opposed to `/` in Stata).

```
> list <- 1:5
> list
[1] 1 2 3 4 5
> list[1]
[1] 1
```

Note again that when you assign something to an object, it doesn't show any output. If you want to see the output as well, just put parentheses around the line.

```
> (list <- c(1, 2, 3, 4, 5))
[1] 1 2 3 4 5
```

You can create lists of objects as well, which may be lists themselves.

```
> a <- c(1, 2, 3, 4, 5)
> b <- 12
> combine1 <- c(a, b)
> combine1
[1] 1 2 3 4 5 12
> c <- "fish"
> combine2 <- c(combine1, c)
> combine2
[1] "1" "2" "3" "4" "5" "12" "fish"
```

Note that when you combine numeric types and string types, they all get changed to strings.

¹If you have previous programming experience, note that unlike many other languages, list indices start at 1 rather than 0.

4 Control Structures

The function of R's control structures (e.g. if, while, foreach) are the same as those in Stata.

4.1 If

Below is an example of an if statement.

```
> a <- 2
> if(a==2) {
+   print("HEY")
+ } else if(a==3) {
+   print("HELLO")
+ }
[1] "HEY"
```

4.2 While

Below is an example of a while statement.

```
> a <- 1
> while(a <= 5) {
+   print("Fish")
+   a <- a + 1
+ }
[1] "Fish"
[1] "Fish"
[1] "Fish"
[1] "Fish"
[1] "Fish"
```

As is the case in any other language, if using a while loop, always be sure that you've set up the loops such that its condition will one day be untrue. In other words, *make sure that the loop will eventually end.*

4.3 For

Below if an example of a for loop.

```
> for(i in 1:5) {
+   print(i)
+ }
[1] 1
[1] 2
[1] 3
[1] 4
[1] 5
```

For loops in R are somewhat simpler, as there is only one command - `for()`. That being said, R does lack the ability to simply loop through variable names. In Stata, we would type:

```
local vars "var1 var2 var3"
foreach var in vars {
    sum `var'
}
```


In R, you would instead do:

```
vars <- c("data$var1", "data$var2", "data$var3")
for(var in vars) {
  v <- get(var)
  summary(v)
}
```

Because R doesn't have the 'variable' notation, you must instead get the data from the variable using the `get()` command and then apply your functions to the extracted data. If you're actually modifying the original variables, then you have to overwrite the original variable data by putting something like `assign(var, newdata)` at the end of the loop.

(Generally, though, people tend not to use for loops in R and instead prefer the `apply()` function family for efficiency reasons.)

5 Comments

Comments in R are preceded by the `#` symbol. As in Stata, comments are used for documentation reasons. Note that in RStudio, you can comment/uncomment entire blocks of code by highlighting them and going to `Code > Comment/Uncomment Lines` from the menu bar.

```
> # This is a comment
```

6 Data Manipulation

6.1 Importing, Viewing, and Editing Data

To import data, you must first find data. Like in Stata, you must set your working directory. The R equivalent of Stata's `cd` command is `setwd()`.

```
> # The same as Stata "cd" command
> setwd("/Users/Richard/Dropbox/R Tutorial")
```

Data, when imported into R, is stored in a data frame object. Data can be read in using the `read.table()` function, for most data types. There is also a `read.csv()` function that can be used for convenience. In addition, you can read Stata `.dta` files using the `read.dta()` function in the `foreign` package.

```
> # Load foreign package
> library(foreign)
> # Read in the Stata file
> plans <- read.dta("plans.dta")
```

The data above is now stored in the `plans` object. You can view the data using the `View()` command. In RStudio, you can also just click on the `plans` in the Workspace viewer in the top right.

It is not recommended, but you can also edit data by using the `edit()` command.

6.2 Accessing and Manipulating Variables

Use the `names()` command to see a list of all variables in the dataset.

```
> # See list of variable names
> names(plans)
```

```

[1] "stu_id" "sch_id" "strat_id" "psu" "f1sch_id" "f1univ1"
[7] "f1univ2a" "f1univ2b" "g10cohrt" "g12cohrt" "bystuwt" "bysex"
[13] "byrace" "bydob_p" "bypared" "bymoethed" "byfathed" "byincome"
[19] "byses1" "byses2" "bystexp" "bynels2m" "bynels2r" "f1qwt"
[25] "f1pnlwt" "f1psepln"

```

Variables are accessed using the \$ symbol. For example, to reference the `bynels2m` variable in the `plans` dataset, you would type `plans$bynels2m`. Thus, getting summary statistics for that variable would look as follows:

```

> # Get summary statistics
> summary(plans$bynels2m)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-8.00  34.37   45.46   44.45  55.76   79.27

```

To add a new variable, just assign something to it. For example, if you wanted a new variable `total` that was a sum of `bynels2m` and `bynels2r`, you would type the following:

```

> # Create new variable
> plans$total <- plans$bynels2m + plans$bynels2r
> # Show new list of variable names
> names(plans)

[1] "stu_id" "sch_id" "strat_id" "psu" "f1sch_id" "f1univ1"
[7] "f1univ2a" "f1univ2b" "g10cohrt" "g12cohrt" "bystuwt" "bysex"
[13] "byrace" "bydob_p" "bypared" "bymoethed" "byfathed" "byincome"
[19] "byses1" "byses2" "bystexp" "bynels2m" "bynels2r" "f1qwt"
[25] "f1pnlwt" "f1psepln" "total"

```

To drop a variable, just assign `NULL` to it.

```

> # Drop a variable
> plans$total <- NULL
> # Show new list of variable names
> names(plans)

[1] "stu_id" "sch_id" "strat_id" "psu" "f1sch_id" "f1univ1"
[7] "f1univ2a" "f1univ2b" "g10cohrt" "g12cohrt" "bystuwt" "bysex"
[13] "byrace" "bydob_p" "bypared" "bymoethed" "byfathed" "byincome"
[19] "byses1" "byses2" "bystexp" "bynels2m" "bynels2r" "f1qwt"
[25] "f1pnlwt" "f1psepln"

```

Having to type the name of the data object for every variable may seem cumbersome, and it will be functionally unnecessary when you are only using one dataset. For these situations, R's `attach` command will allow you to use the dataset as if it were the only dataset.

```

> # Attach dataset to workspace
> attach(plans)
> # Show summary, no need to reference dataset
> summary(bynels2m)

  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-8.00  34.37   45.46   44.45  55.76   79.27

```

If you make any modifications to any of the variables, you will need to re-attach the data before moving on.

6.3 Merging and Appending

Merging is fairly simple in R. Just use the `merge()` function and specify the variable on which you are merging with the `by` option. (You can even specify that the variable you're merging on is named differently in the two datasets by using the `by.x` and `by.y` options.)

```
> newdata <- merge(data1, data2, by="ID")
```

To append data, notice that a data frame is basically a matrix with column labels. So to append two datasets on top of each other (they must have the same number of columns), use the `rbind()` command, which is a matrix command that glues two matrices on top of each other. (The `cbind()` command does the same, but glues the matrix side-by-side.)

```
> newdata <- rbind(data1, data2)
```

7 Data Analysis

There are hundreds of data analysis packages and functions out there, but the basic ones are covered here. Note that most of these do have options to include analytic weights, similar to Stata.

7.1 Summary Statistics

Most summary statistics can be retrieved using the `summary()` command. You can store the information in an object.

```
> # Store summary information in an object
> msum <- summary(plans$bynels2m)
```

The `msum` object above is a summary object, and it contains information. The data is stored as a labeled list. It can be accessed in many ways. See below for examples.

```
> # Show full summary
> msum
  Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
-8.00  34.37   45.46   44.45  55.76   79.27

> # Get first item in list by index
> msum[1]
Min.
-8

> # Get first item in list by name
> msum["Min."]
Min.
-8
```

The `summary()` command can also be used for categorical variables.

```
> # Show summary
> summary(plans$byrace)
 {survey component legitimate skip/na}
                                276
 {nonrespondent}
                                651
amer. indian/alaska native, non-hispanic
```

```

131
asian, hawaii/pac. islander,non-hispanic
1465
black or african american, non-hispanic
2027
hispanic, no race specified
998
hispanic, race specified
1229
multiracial, non-hispanic
740
white, non-hispanic
8735

```

7.2 Regression

Standard linear regression uses the `lm()` function, which stands for “linear model.” See `?lm` for the specific syntax, but the most general form is `lm(formula, dataobj)`.

The formula proceeds in the same order as it does in Stata, with additional formatting. For example, a regression modeling `bynels2m` on `bynels2r` and `byses1` would look as follows:

```

> # Run the model
> model1 <- lm(bynels2m ~ bynels2r + byses1, data=plans)

```

The results from this regression are now stored in `model1`, which is an object itself. To see the results or do anything with the results, you would just type `model1`. For linear model objects, to see the full results, type `summary(model1)`.

```

> # See abbreviated results
> model1
Call:
lm(formula = bynels2m ~ bynels2r + byses1, data = plans)

```

```

Coefficients:
(Intercept)    bynels2r    byses1
      13.980         1.062         1.287

```

```

> # See more detailed results
> summary(model1)

```

```

Call:
lm(formula = bynels2m ~ bynels2r + byses1, data = plans)

```

```

Residuals:
    Min     1Q   Median     3Q    Max
-32.478  -5.716  -0.290   5.744  46.584

```

```

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept) 13.980313   0.239390   58.40  <2e-16 ***
bynels2r     1.062190   0.007649  138.88  <2e-16 ***
byses1       1.287077   0.054390   23.66  <2e-16 ***
---

```

```

Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1 1

```

```
Residual standard error: 8.812 on 16249 degrees of freedom
Multiple R-squared: 0.6589, Adjusted R-squared: 0.6588
F-statistic: 1.569e+04 on 2 and 16249 DF, p-value: < 2.2e-16
```

The `model1` object is the first really complex object we've seen (it is of class `lm`). In fact, there are 12 different components of the object, which you can see if you use the `attributes()` command.

```
> # See attributes of model1 object
> attributes(model1)

$names
 [1] "coefficients" "residuals" "effects" "rank"
 [5] "fitted.values" "assign" "qr" "df.residual"
 [9] "xlevels" "call" "terms" "model"

$class
 [1] "lm"
```

As you can see, there are many things stored in this model. Unlike in Stata, where predicted values and residuals need to be calculated separately, R calculates them automatically when running the model and stores them in the `fitted.values` and `residuals` components, respectively.

```
> # Create a new variable and assign residuals to that variables
> plans$resid <- model1$residuals
```

In fact, the summary of the model has components of its own.

```
> # See attributes of summary of model1 object
> attributes(summary(model1))

$names
 [1] "call" "terms" "residuals" "coefficients"
 [5] "aliased" "sigma" "df" "r.squared"
 [9] "adj.r.squared" "fstatistic" "cov.unscaled"

$class
 [1] "summary.lm"

> # Pull out just coefficient information
> summary(model1)$coefficients


```

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	13.980313	0.239389805	58.39978	0.00000e+00
byne1s2r	1.062190	0.007648541	138.87483	0.00000e+00
byses1	1.287077	0.054389852	23.66391	9.65036e-122

Note that the coefficient information above is simply a data table, with variable names in the rows and statistics in the columns. Realizing things like this helps with coming up with creative ways to manipulate this information for your own use.

Other regression types (e.g. binomial, poisson) can be run with the `glm()` function, where the `family` option specifies which to use. Multinomial logit can be run with `multinom()` in the `nnet` package. Other types can be found online.

7.3 Survey Weighting for Estimates of Variance

Unfortunately, using complex survey weighting is just as complicated in R as it is in Stata. The main suite of functions used to do complex survey weighting in R is in the `survey` package, created by Thomas Lumley from the University of Washington. There are plenty of resources online, but there is also a very helpful

accompanying book: *Complex Surveys: A Guide to Analysis Using R*. The limitations of complex survey weighting in Stata are similar to those in R.

Using complex survey weights in R involves first creating a `survey` object and then using the packaged functions to perform data analysis. Below is the syntax, for example, for creating a survey object using PIRLS data, which uses jackknife replicate weights.

```
> # Create svy object
> svy <- svrepdesign(repweights="jkjr_[0-9]+",
+                 weights=~totwgt,
+                 data=pirldata,
+                 type="other")
```

From here, you use the various functions in the `survey` package to do data analysis. Some sample functions:

```
> # To get the mean of the math variable
> svymean(~math, svy, na.rm=TRUE)
> # To create a regression of math on ses and reading
> svyglm(math ~ ses + reading, design=svy)
```

The structures of the objects created by these functions mostly mirror those of the non-survey functions (`mean()` and `lm()`, in this case).

7.4 Other Data Analysis Tools

For almost all data analysis techniques that you can do in Stata, there is an equivalent function in R. For example, multilevel modeling uses `xtmixed` in Stata, and `lmer()` in R. The internet is a fast way to figure out what these functions are.

8 Tables and Graphs

8.1 Tables

In our field, there are two basic types of tables that people make: descriptive tables and regression tables. These can be approached in many different ways, but here are two techniques that work fairly well.

Descriptive Tables Descriptive tables can be looked at as data tables in and of themselves, with the different statistics across one dimension and the variables across the other. As such, to output descriptives tables, store the descriptives as a data frame and use the `xtable()` and `print.xtable()` functions to output the tables to files. These can be Word, HTML, LaTeX, etc. The output from the following code is in Table 1.

```
> # Store summaries of score variables
> mathsum <- summary(bynels2m)
> readsum <- summary(bynels2r)
> # Glue them together
> sumtable <- cbind(mathsum, readsum)
> sumtable
```

	mathsum	readsum
Min.	-8.00	-8.00
1st Qu.	34.37	21.94
Median	45.46	29.79
Mean	44.45	29.00
3rd Qu.	55.76	36.85
Max.	79.27	50.57

```

> # Change the column names
> colnames(sumtable) <- c("Math Summary", "Read Summary")
> sumtable

      Math Summary Read Summary
Min.          -8.00         -8.00
1st Qu.       34.37         21.94
Median        45.46         29.79
Mean          44.45         29.00
3rd Qu.       55.76         36.85
Max.          79.27         50.57

> # Flip the table, just for kicks, using the transpose function t()
> (sumtable <- t(sumtable))

      Min. 1st Qu. Median Mean 3rd Qu. Max.
Math Summary  -8   34.37  45.46 44.45  55.76 79.27
Read Summary  -8   21.94  29.79 29.00  36.85 50.57

> # Load xtable library
> library(xtable)
> # Create xtable object
> textble <- xtable(sumtable, caption="Math and Reading Summary")
> # Output to file
> print.xtable(textble, type="latex", file="sumtable.tex",
+               append=FALSE, caption.placement="top")

```

Table 1: Math and Reading Summary

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Math Summary	-8.00	34.37	45.46	44.45	55.76	79.27
Read Summary	-8.00	21.94	29.79	29.00	36.85	50.57

Regression Tables The `texreg()` function is fairly adept at producing tables in TeX format. Word tables can be produced in a few different ways, which can be found online. The output from the code below is in Table 2.

```

> # Two example regressions
> model1 <- lm(bynels2m ~ byses1 + bynels2r, data=plans)
> model2 <- lm(bynels2r ~ byses1 + bynels2m, data=plans)
> # Load texreg library
> library(texreg)
> # Output to file
> # Note that in R, "T" is often used as an abbreviation for "TRUE" (and "F" for "FALSE")
> regtable <- texreg(list(model1, model2), file="regtable.tex",
+                    caption.above=T, caption="Regression Table")

The table was written to the file 'regtable.tex'.

```

For both of these situations, you can also just use the `write.table()` or `write.csv()` functions to output the tables to a spreadsheet that you can modify manually in Excel.

8.2 Graphs

Among other things, R is particularly well-known for being able to produce fantastic graphics (see the R Graph Gallery for examples). That being said, plots beyond very simple scatter and bar plots do take a

Table 2: Regression Table

	Model 1	Model 2
(Intercept)	13.98*** (0.24)	6.54*** (0.18)
byses1	1.29*** (0.05)	1.00*** (0.04)
bynels2r	1.06*** (0.01)	
bynels2m		0.51*** (0.00)
R ²	0.66	0.66
Adj. R ²	0.66	0.66
Num. obs.	16252	16252


*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

little bit of practice. Luckily, there are code examples all over the web that you can reference to produce graphs.

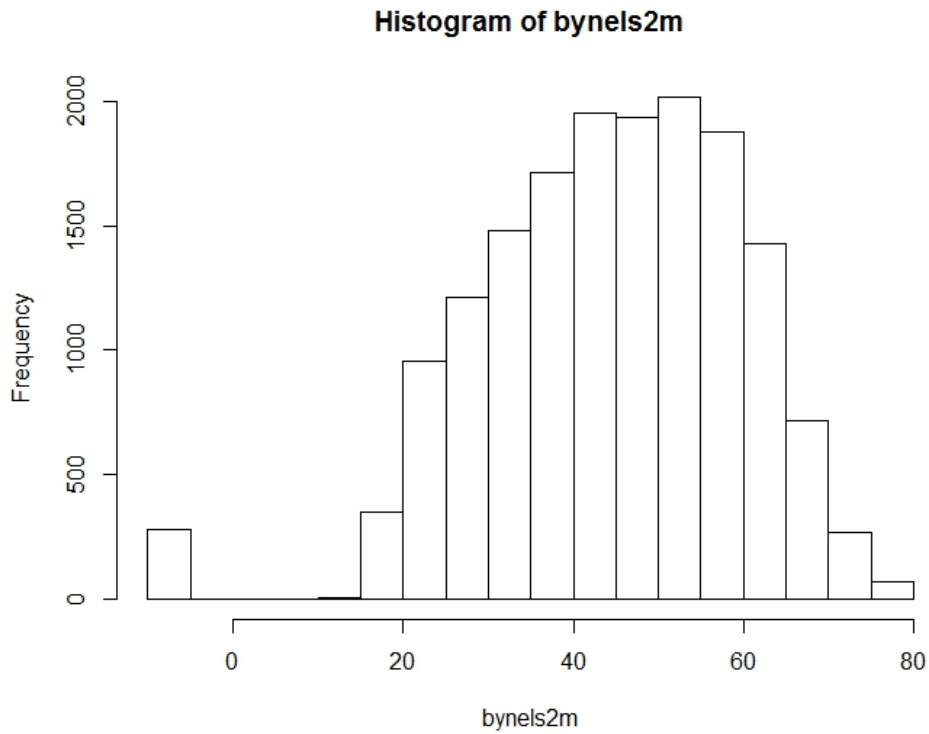
All plots will show up in the plot window in the bottom-right of the RStudio window. They can be downloaded from there. Alternatively, you can output from the code by using functions like `pdf()`.

There are simple graph functions such as `plot()` function, which can cover your basic scatter plots. There is a `barplot()` function for bar plot, and a `hist()` function for histograms.

```
> # Attach dataset for convenience
> attach(plans)
> # Basic scatterplot
> plot(bynels2m, bynels2r)
```




```
> # Basic histogram
> hist(bynels2m)
```



All of these functions have options for specifying labels and such.

While these functions may be sufficient for your use, exploring other plotting packages (in particular, the `ggplot2` package) is recommended for the future. An example is below, for comparison.

```
> # Load ggplot2 library
> library(ggplot2)
> # Create ggplot graph
> ggplot(plans, aes(x=bynels2m, y=bynels2r, color=byrace)) +
+   geom_point() +
+   ggtitle("Math vs. Reading") +
+   xlab("Math") +
+   ylab("Reading")
```



9 Other R Things

9.1 Sweave/knitr

One benefit of using R, for LaTeX users, is being able to embed R code using either the **Sweave** or **knitr** packages. (This document was created using **Sweave**). As such, when compiling documents, the R code can be run during compilation such that code output and input need not be manually copied and pasted from R into LaTeX.

9.2 GIS

R has a variety of GIS functions and packages that can be used to create GIS maps while running data analysis fairly easily. See the **maps**, **mapproj**, and **maptools** packages.

9.3 Functions

As in Stata, you can write your own functions in R. R's syntax for doing so, however, (1) matches the syntax of most other popular programming languages and (2) is a bit more intuitive. For example, see the following function that moves a column in the data frame to the beginning of the data.

```
> # Define a function
> to.first <- function(name, data) {
+   # Get the index of that variable
+   idx <- grep(name, names(data))
+   # Overwrite the data with that variable first
+   data <- data[,c(idx, (1:ncol(data))[-idx])]
+   # Return the data
}
```

```
+ return(data)
+ }
```

This creates the function `to.first()`, which takes in the name of the variable and the data and returns the re-ordered data.

```
> names(plans)
 [1] "stu_id"  "sch_id"  "strat_id" "psu"      "f1sch_id" "f1univ1"
 [7] "f1univ2a" "f1univ2b" "g10cohort" "g12cohort" "bystuwt"  "bysex"
[13] "byrace"   "bydob_p"  "bypared"  "bymoethed" "byfathed" "byincome"
[19] "byses1"   "byses2"   "bystexp"  "bynels2m"  "bynels2r" "f1qwt"
[25] "f1pnlwt"  "f1psepln" "resid"

> plans <- to.first("bynels2m", plans)
> names(plans)
 [1] "bynels2m" "stu_id"  "sch_id"  "strat_id" "psu"      "f1sch_id"
 [7] "f1univ1"  "f1univ2a" "f1univ2b" "g10cohort" "g12cohort" "bystuwt"
[13] "bysex"    "byrace"   "bydob_p"  "bypared"  "bymoethed" "byfathed"
[19] "byincome" "byses1"   "byses2"   "bystexp"  "bynels2r"  "f1qwt"
[25] "f1pnlwt"  "f1psepln" "resid"
```

The variable is now moved. You can use this function now whenever you want to move a variable to the beginning.

10 Translations

Below are some rough translations between Stata and R, for more non-obvious translations.

Stata	R
help function	?function
cd	setwd()
drop var	var <- NULL
=	<- or =
/*, ///, *	#
display	print()
keep <i>varlist</i>	data <- data[<i>varlist</i>]
keep if	data <- data[which()]
a + b (strings)	paste0(a, b)
'v'Name = data	assign(paste0(v, "Name"), data)
save, outsheet	write.table()
use, insheet	read.table()
., ""	NA
ds	names(data)
svyset	svrepdesign()
svy: sum	svymean()
svy: reg	svyglm()
if var==.	is.na(var)
estout	xtable(), print.xtable()
reg	lm()
esttab	texreg()
xtmixed	lmer()
tab	table()