

14.310x: Data Analysis for Social Scientists

Joint, Marginal, and Conditional distributions & Functions of Random Variables

Welcome to your third homework assignment! We encourage you to get an early start, particularly if you still feel you need more experience using R. We have provided this PDF copy of the assignment so that you can print and work through the assignment offline.

Good luck!

In this problem set we will guide you through different ways of accessing real data sets and how to summarize and describe it properly. First we will go through some of the data that is collected by the World Bank. We will do some cleaning on the data before we start analyzing it. Then, we will try to do a simple web scraping exercise where we will analyze the data as well.

Let's start with this dataset from the World Bank.

This dataset is from the World Bank Gender Statistics website. You may visit the website and explore other datasets they offer. **However, for the purposes of this assignment, please use the file in this set of instructions**, since the dataset on the World Bank website may have been updated in the time since this problem set and answer key was posted.

NOTE: It is important to work in the same directory that the files are or to use the whole path when you specify opening a data set. To know which directory you are currently working in, you can use the command `getwd()`. Similarly, in order to set a different directory, you can use the

For the purpose of analyzing the data, we are going to use the packages **utils** and **tidyverse**. Once you have uploaded the data to R you are going to see there are multiple indicators of gender, countries and years in the data. In this case we are just interested in analyzing the data for one indicator that is the *Adolescent Fertility Rate*, in the data the indicator code for this variable is called `SP.ADO.TFRT`. This indicator measures the annual number of births to women 15 to 19 years of age per 1,000 women in that age group. It represents the risk of childbearing among adolescent women 15 to 19 years of age. It is also referred to as the age-specific fertility rate for women aged 15-19. Once you have completed this problem set you'll have more information of how this rate has evolved over time and how it varies across different groups of countries.

Take a look at the following lines of code, whose main purpose is to upload the data in a data frame and to choose the proper indicator. Please, try to understand the code and then run it in your computer. Remember to set the directory accordingly to the folder where you saved the files.

```

#Preliminaries
rm(list=ls())
library("utils")

setwd("-----") #fill in with your working directory!

#Getting the data
gender_data <- as_tibble(read.csv("Gender_StatsData.csv"))

```

Using the code, answer the following questions:

Question 1

What is the purpose of the line `rm(list = ls())`?

- To remove all current existing objects in R
- To change the current directory path
- To list all the files in the current directory
- To look in the web for the World Bank dataset

Question 2

The first thing you want to figure out when you look at a new dataset is how it is organized. If your dataset is stored as a tibble, you can simply print the object and it will print in a nice-looking format.

Alternatively, you can also use the built-in R commands such as `str()`, which allows you to see the structure of an object in R. Likewise, the commands `head()` and `tail()` will allow you to see the first six and last six observations of your data frame respectively. Another useful function is the function `dim()`, which will give you the number of rows and columns in your dataset. Take the time to explore the data using these commands and others.

Which of the following statements best describes how your data is organized? Note that we use “indicator” to mean something that is being measured (e.g. fertility rate, enrollment rate, etc.).

- Each unit of observation (row) is a country/region for a given year.
- Each unit of observation (row) corresponds to a country/region and an indicator
- Each unit of observation (row) corresponds to an indicator for a given year.
- Each unit of observation (row) is a country/region’s indicator for a given year.

Question 3

Now, generate a tibble called “`teenage_fr`”, which contains only the adolescent fertility rate indicator for each country-year. Please fill in the blank with the correct code.

```

teenager_fr <- _____(gender_data, Indicator.
Code=="SP.ADO.TFRT")

```

- o Filter
- o Select

What is the equivalent base-R function? (Select one)

- o The equivalent base-R function is `match()`.
- o The equivalent base-R function is `which()`.
- o The equivalent base-R function is `subset()`.

Question 4

Since we are not interested in any other variables and the `gender_data` dataset is quite large, you might want to get rid of it instead of asking R to keep it stored in memory.

Which of the following is the correct code for getting rid of the object `gender_data`?

- o `data$gender_data <- NULL`
- o `rm(gender_data)`
- o `delete(gender_data)`
- o `gender_data <- 0`

Now that you have loaded the data we want to analyze and have familiarized yourself with the structure, it is time to get our hands dirty!

A second exploratory thing to do once we have organized a data set is to get basic summary statistics of the data. Now let's do this! To print summary statistics directly in your console, you can use any of the basic summary functions in R (`mean()`, `sd()`, `min()`, `max()`, `sum()`...). The basic summary functions take vectors as an input, and output a single value.

For example, if you were interested in obtaining the sample mean of the Adolescent Fertility Rate in 1975, one way of doing this is as follows:

```
mean(teenager_fr$X1975, na.rm=TRUE)
```

Question 5

Why is it necessary to add the option `na.rm=TRUE` to the above command? (Select all that apply)

- The default option of `na.rm` is set to `FALSE`. Therefore, if we don't specify this, R will try to calculate the mean using all the observations in the data.
- This part is necessary since otherwise R would duplicate some of the observations in the data set when it calculates the sample mean. In particular, the observations with missing values would have higher weights than the observations without missing values.
- It is not necessary to add this option to the command to obtain the mean of this variable.
- Otherwise we will obtain a missing value (`na` in R) since not all the countries in the data have information on the adolescent fertility rate in 1975.

To calculate summary statistics for a group of variables, there are a few different commands. The command `mean()` is just one example of the different options available. Now, we ask you to go through the R documentation and explore some of the other commands by yourself.

If you want to store the output as values in your dataset, or if you want to do something more complicated (ex. Generate these by group, or use one of the `dplyr` summary functions (ex. `n_distinct()`), you can use any of the basic summary functions as well as others, in combination with `mutate()` and `summarise()` to generate variables in your dataset containing summary values.

Now that you've learned how to look at and generate summary statistics, answer the following questions.

Question 6

What is the sample mean and standard deviation of the adolescent fertility rate in 1960?

Sample mean:

Standard deviation:

Question 7

What is the sample mean and standard deviation of the adolescent fertility rate in 2000?

Sample mean:

Standard deviation:

Question 8

True or False? From the values that you have calculated above we can conclude that the Adolescent Fertility Rate has had a permanent decreasing trend (i.e. only decreases and never increases during this period) from 1960-2000, and that the dispersion of this variable has decreased over time.

- True
- False

Now, we are interested in plotting the evolution of the Adolescent Fertility Rate from 1960 to 2015. In addition, we are interested in having different information in the same plot. First, we want to plot the sample mean of all the data set, but also we want to add more information such as the rate for low, middle and high income countries (an indicator for country code is stored in the variable "Country.Code").

Inspect this variable to get a sense of what it contains. Note that it includes indicators for both countries, regions, and income group. Since we are only interested in the trends by income group, we want to filter the data to contain only the fertility rate for high, middle, and low income countries as well as the world average.

Question 9

Use the `dplyr filter()` command and the logical `%in%` to keep only the relevant `Country.Code` observations in `teenage_fr`. Make sure you name the new dataset "byincomelevel". Choose the line of code below:

- o `byincomelevel <- filter(teenager_fr, Country.Code%in%c(LIC, MIC, HIC, WLD))`
- o `teenage_fr <- filter(byincomelevel, Country.Code%in%(LIC, MIC, HIC, WLD))`
- o `byincomelevel <- filter(Country.Code%in%c("LIC", "MIC", "HIC", teenager_fr))`
- o `byincomelevel <- filter(teenager_fr, Country.Code%in%c("LIC", "MIC", "HIC", "WLD"))`

Notice, there are still two problems with the resulting data:

1. It contains additional variables that we don't need or are meaningless at this level of aggregation.
2. It is not organized in a very intuitive way. A more natural way to organize this data, and prepare it for plotting, is to have each observation represent either a year or a country group-year, and each of the columns represent either the fertility rate for a given group, or if the data is at the country-group year level, then just the fertility rate.

Question 10

Suppose you decide you prefer to have one observation per income group and year. The `dplyr` command `gather()` can help you achieve this. Look up the command in the help files. Select the set of arguments that belong in the blanks below.

Note: Depending on your operating system, you may encounter an encoding error when trying to reference the variable "Country.Name", if that is the case, add the following line before running the code:

```
byincomelevel <- colnames(byincomelevel)[1]="Country.Names")
```

OR

```
colnames(byincomelevel)[1]<- "Country.Name"
```

```
plotdata_bygroupyear <- gather(byincomelevel, ____ (1) ____,
____ (2) ____,
____ (3) ____ ) %>%
select(Year, Country.Name, Country.Code, FertilityRate)
```

1. Select from Country.Code, Year, FertilityRate, X1960:X2015
2. Select from Country.Code, Year, FertilityRate, X1960:X2015
3. Select from Country.Code, Year, FertilityRate, X1960:X2015

Question 11

Suppose you take a look at the data and change your mind. You decided you prefer to look at the data at the year level and have the fertility rates for each income-group as separate variables. The `dplyr` command `spread()` can help you achieve this. Look up the command in the help files. Select the set of arguments that belong in the blanks below:

```
plotdata_byyear <- select(plotdata_bygroupyear, Country.Code,
Year, FertilityRate) %>%
spread(____(1)____, ____ (2)____)
```

1. Select from Year, Country.Code, FertilityRate
2. Select from Year, Country.Code, FertilityRate

Question 12

True or False? The `select` statement in the code for question 11 is redundant, since we already selected the variables we wanted in generating `plotdata_byyear`.

- True
- False

Question 13

Good news. We are finally ready to plot the data! Let's begin by plotting the fertility rate over time, separately for each income level. To do this, we can use the basic `ggplot` syntax Prof. Duflo explained in lecture.

Let's start by trying to generate this plot using the `plotdata_bygroupyear` tibble we generated earlier. Here is the code to generate this plot. Select the set of arguments that belong in the blanks to generate the desired plot.

```
ggplot(plotdata_bygroupyear, aes(x=____(1)____, y=____(2)____,
group=____(3)____ ))
+ geom_line()
```

1. Select from Year, Country.Code, FertilityRate
2. Select from Year, Country.Code, FertilityRate
3. Select from Year, Country.Code, FertilityRate

Question 14

It would be nicer if the different plot lines had different colors. You can add the argument `color=Country.Code` to the code you generated in question 13. Where do you need to specify this argument? Select one of the Roman numeral blanks in the code below to replace with `,color=Country.Code` or `color=Country.Code` in order for each of the lines to have a different color.

```
ggplot(plotdata_bygroupyear, aes(x=____, y=____, group=____ I)
II) + III
geom_line(IV)V
```

- I
- II
- III
- IV
- V

Question 15

It is good practice to include titles in your plot. To do this, look up the `ggplot` (`labs`). Select one of the Roman numeral blanks in the code below to replace with (possibly preceded by `.` or `+`) `labs(title='Fertility Rate by Country-Income-Level over Time')`.

```
ggplot(plotdata_bygroupyear, aes(x=____, y=____, group=____ I)
II) +
  geom_line(III) IV
  ○ I
  ○ II
  ○ III
  ○ IV
```

Question 16

One more thing we could improve in this plot is the x-axis labels. First, we can remove the leading “X”. Second, by storing them as numeric, `ggplot` can use its “optimal” scaling to make a prettier plot instead of having a label for each year. To do this, we can transform the `Year` variable using `dplyr`’s `mutate` function and a combination of the functions `as.numeric()` and the `stringr` package. Try to figure out a few ways to do this.

Which of the following statements are equivalent and can be used in the blank below to complete it? Select all that apply

```
plotdata_bygroupyear <- mutate(plotdata_bygroupyear,
Year=as.numeric(____))
 str_sub(year, -4)
 str_sub(Year, 2, 5)
 str_replace(Year, "X", "")
```

Question 17

Which of the following statements can you conclude from the plot?

- The rate of the world average is always below the rate of high and middle income countries.
- While the rate for high income countries presents a decreasing trend in the period, the rate for low income countries fluctuates until the mid-nineties when it becomes to decrease significantly.
- The gap between high and middle income countries is lower in 2014 than in 1960, while the gap between low and middle income countries is actually larger in 2014 than it is in 1960.
- Since the mid-nineties, the rate for low income countries has decreased more than it has for high and middle income countries.

Now, we are not going to consider the trends of the different categories over the years. Instead, we are going to compare how the distribution of the Adolescent Fertility Rate is different between 1960 and 2000.

We have provided the R script here to help you answer the next set of questions.

First, we want to generate `histdata_twoyears`.

```
histdata_twoyears <- select(teenager_fr, Country.Name,
  Country.Code, Indicator.Name,
  Indicator.Code, X1960, X2000)


histdata_twoyears <- gather(teenager_fr, Year, FertilityRate,
  X1960, X2000) %>%
  select(Year, Country.Name, Country.Code, FertilityRate)

histdata_twoyears <-
  filter(histdata_twoyears, !is.na(FertilityRate))
```

We want to plot a histogram of the two variables. The following code in R plots the histogram in the same graph. Please take a look at the code and try to understand what it is doing.

```
ggplot(histdata_twoyears, aes(x=FertilityRate)) +
  geom_histogram(data=subset(histdata_twoyears,
  Year=="X1960"),
  color="darkred", fill="red", alpha=0.2) +
  geom_histogram(data=subset(histdata_twoyears,
  Year=="X2000"),
  color="darkblue", fill="blue", alpha=0.2)
```

Here is the figure that this code has produced:



Question 18

What does the argument `alpha` dictate?

- The width of the bins.
- The width of the outline of the bins.
- The extent to which the plot colors are different.
- The level of transparency in the color of bins.

Question 19


As you can see, we have certain number of bins in the figure. Go to the R documentation and look for the option in the command `geom_histogram()` that will allow you to change the number of bins in the figure. Select all that apply.

- `binwidth`
- `position`
- `breaks`
- `bins`
- `center`

Now, we are going to add some kernels to the histogram. The kernels were done using the command `density`, and all the default options in R. Again, take a look at the code, run it on your

```
ggplot(histdata_twoyears, aes(x=FertilityRate, group=Year, color=Year,
alpha=0.2)) + geom_histogram(aes(y=..density..)) +
  geom_density(data=subset(histdata_twoyears, Year=="X1960"),
color="darkred", fill="red", alpha=0.2, bw=5)+
  geom_density(data=subset(histdata_twoyears, Year=="X2000"),
color="darkblue", fill="blue", alpha=0.2, bw=5)
```

The below figure is produced by running the code above.



Question 20

As it was stated before, the plot was done using the default options in R. For the kernel, the default option is to use `gaussian`. There are other options that the user can state when running

not bell-shaped? In other words, which one doesn't underweight observations at the boundaries of each bin.


- `gaussian`
- `epanechnikov`
- `rectangular`
- `triangular`
- `biweight`
- `cosine`
- `optcosine`

Question 21

The following plots were made by changing the bandwidth of the kernel function in R. Which one of them was made with the largest bandwidth?


- It is not possible to tell just by looking at the figure.

Change in the distribution




o

Change in the distribution



o


Change in the distribution



o

One of the things that Professor Duflo also discussed in the lecture was the construction of the Empirical Cumulative Distribution (ECD). The following figures shows the ECD for the Adolescent Fertility Rate in the World in 1960 and in 2000. However, as you can see, the person who made the graph forgot to properly label it.

Empirical Cumulative Distribution



Question 22

Can you infer from the histograms that were plotted before, which one corresponds to the Adolescent Fertility Rate in 2000 and which one to the same indicator in 1960? (Select all that apply)

- Blue corresponds to 2000
- Black corresponds to 2000
- Blue corresponds to 1960
- Black corresponds to 1960
- It is not possible to tell from the plot

Question 23

Using the figure, can you determine whether the distribution used to construct the black series satisfies the First Order Stochastic Dominance property over the distribution used to construct the blue series? Assume that the blue plot is always above (has a value greater than) the black plot.

- Yes
- No

MIT OpenCourseWare
<https://ocw.mit.edu/>

14.310x Data Analysis for Social Scientists
Spring 2023

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>.