

Exercise 2

```
library(rio)
library(tidyverse)
```

Educational opportunity in the NLSY

For this exercise, you will be working with an extract from the 1997 *National Longitudinal Survey of Youth* (NLSY97),¹ a major panel study tracking about 9,000 youth who were between 12 and 15 on December 31, 1996. NLSY97 is intended to follow young adults both as they complete their education and over their working lives. It's also a pretty messy data set to work with.

1. Create your exercise directory

- On your computer, create a folder for this seminar exercise.
- Download the NLSY97 data and extract the files into your seminar exercise folder.

2. Perform the included processing of the NLSY data file

- The NLSY provides some preliminary formatting to the raw data file (*nsl97.dat*) via a R script file (here *nsl97.R*). Open this script file in RStudio.
- To run the script, you will first need to uncomment the `setwd()` function and provide the path to your exercise directory.
- At the end of the script, uncomment the necessary lines to “rename variables using Qnames instead of Reference Numbers.”
- Export the data frame *new_data* as *nlsy97.rds*.
- Save the *nsl97.R* and begin working in a new working R script file, wherein you will complete the remainder of the exercise.

3. Load the data and view the structure

- In the new script file you've created, import *nlsy97.rds* dataframe as *nlsy97*.

```
nlsy97 <- import("m2_nlsy97.rds")
```

- Look at the structure of the data frame using the `glimpse()` and `str()` functions.

```
glimpse(nlsy97)
str(nlsy97)
```

- The list of column names is a bit hard to read in either approach - and is even truncated with `str()`. To make it easier to scan the column names, save them to a new data object formatted as a **tibble**.

```
cols_nlsy97 <- colnames(nlsy97) %>% as_tibble()
```

```
## Warning: Calling `as_tibble()` on a vector is discouraged, because the behavior is likely to cha
## This warning is displayed once per session.
```

- “Tidy” the dataset:** A key principle of data analysis is that data should follow the “tidy” data structure, wherein each row of a data set corresponds to an observation, and each column a variable.

¹Available at <https://www.nlsinfo.org/investigator/>. Keep in mind that this exercise is strictly pedagogical - we ignore things like missing value codes in the data, hence there are obvious errors.

- a. Notice that there are both time-varying and time-invariant variables. Rename the variables that are time-invariant as follows (to make your life easier in the next step):

Old Name	New Name
PUBID_1997	personid
KEY_SEX_1997	sex
KEY_RACE_ETHNICITY_1997	race
KEY_BDATE_Y_1997	birthyr
CV_HGC_RES_MOM_1997	motheredys

- b. Begin to tidy the data by using the **gather()** function to gather all variables that are time-varying into a key and value column. To do that, use the **gather()** function on all columns except those mentioned in step (d), with *variable* as the name of the key column and *value* as the name of the value column.

```
nlsy97 <- nlsy97 %>%
gather(variable, value, -c(personid,sex,race,birthyr,motheredys))
```

- c. All time-varying variables should now appear in a single column and all corresponding values in another. But notice that we have another issue – each year of a variable has a different name. To fix this:

- (i) Create a *year* column, with values that are the numbers from the *variable* column. You can do so as follows:

```
nlsy97$year <- str_extract(nlsy97$variable, "[0-9]+")
```

- (ii) Notice that some of the year values are written in the wrong format. Recode the following values of *year*:

Old Value	New Value
9293	1993
9394	1994
9495	1995
9596	1996
9697	1997
9798	1998
9899	1999
9900	2000
0001	2001
0102	2002
0203	2003
0304	2004

- (iii) Replace the numbers in the *variable* column with an empty string. You can do so as follows:

```
nlsy97$variable <- str_replace(nlsy97$variable, "[0-9]+","")
```

- d. Now that the variable names and years are consistent, finish “tidying” the data by **spreading** each of the variables in *variable* to its own column.

```
nlsy97 <- nlsy97 %>% spread(key = variable, value=value)
```

5. Rename and select variables

- a. Rename the time-varying variables as follows:

Old Name	New Name
CV_CENSUS_REGION__	region
CV_INCOME_GROSS_YR__	parentincome
CV_SCHOOL_TYPE__	schooltype
CV_HGC_EVER_EDT__	highestgrade
TRANS_AB_AY_HSTR	absences
TRANS_CRD_GPA_YR__HSTR	gpa

b. Keep only the variables you've renamed in 4(a) and 5(a) as well as the year variable.

```
nlsy97 <- nlsy97 %>% select(
  personid, year, sex, race, birthyr, region, parentincome,
  schooltype, motheredyrs, highestgrade, absences, gpa)
```

6. Finalize the data set

a. Ensure that each variable has the correct data type. This means you'll probably have to change the data type for the following variables: *personid*, *year*, *sex*, *race*, *region*, and *schooltype*.

```
export(nlsy97, "nlsy97.rds")

nlsy97$personid <- as.factor(nlsy97$personid)
nlsy97$year <- as.numeric(nlsy97$year)
nlsy97$sex <- as.factor(nlsy97$sex)
nlsy97$race <- as.factor(nlsy97$race)
nlsy97$region <- as.factor(nlsy97$region)
nlsy97$schooltype <- as.factor(nlsy97$schooltype)
```

b. Create the following new variables:

- $age = year - birthyr$
- $gradesback = (age - 7) - highestgrade$
- A female indicator variable
- A non-white indicator variable

```
nlsy97 <- nlsy97 %>% mutate(
  age = year - birthyr,
  gradesback = (age - 7) - highestgrade,
  female = (sex == 2),
  nonwhite = (race != 4)
)
```

c. Rescale GPA: Redefine GPA so that it is the original value divided by 100.

```
nlsy97 <- nlsy97 %>% mutate(
  gpa = gpa / 100
)
```

d. Filter the data set to only include the years 1997-2007, for individuals from the South-East (region 4), with a highest grade completed (*highestgrade*) between 8 and 17.

```
nlsy97 <- nlsy97 %>% filter(
  year %in% 1997:2007,
  region == 4,
  highestgrade %in% 8:17
)
```

7. Create and view summary statistics

a. By age, create and display summary statistics with average values of:²

- Highest grade completed (*highestgrade*)
- Number of grades below expected level (*gradesback*)

```
summarybyage <- nlsy97 %>% group_by(age) %>%
  summarize(
    avg_gradecompleted = mean(highestgrade, na.rm = TRUE),
    avg_gradesback = mean(gradesback, na.rm = TRUE))
summarybyage
```

age	avg_gradecompleted	avg_gradesback
14	8.015267	-1.0152672
15	8.662760	-0.6627604
16	9.570429	-0.5704287
17	10.473154	-0.4731544
18	11.304080	-0.3040805
19	11.857471	0.1425287
20	12.261765	0.7382353
21	12.578632	1.4213685
22	12.921179	2.0788207
23	13.112805	2.8871951
24	13.164405	3.8355946
25	13.216009	4.7839912
26	13.043253	5.9567474
27	13.003636	6.9963636

b. By highest grade completed, create and display summary statistics with average values of:

- Parental income (*parentincome*)
- Mother's education (*motheredyrs*)
- GPA (*gpa*)
- Share women
- Share nonwhite

```
summarybygrade <- nlsy97 %>% group_by(highestgrade) %>%
  summarize(
    avg_parincome = mean(parentincome, na.rm = TRUE),
    avg_mothersed = mean(motheredyrs, na.rm = TRUE),
    avg_gpa = mean(gpa, na.rm = TRUE),
    share_female = mean(female, na.rm = TRUE),
    share_nonwhite = mean(nonwhite, na.rm = TRUE))
summarybyage
```

age	avg_gradecompleted	avg_gradesback
14	8.015267	-1.0152672
15	8.662760	-0.6627604
16	9.570429	-0.5704287
17	10.473154	-0.4731544
18	11.304080	-0.3040805
19	11.857471	0.1425287
20	12.261765	0.7382353

²For both summary statistics tables, you may need to use the `na.rm = TRUE` to remove missing values when calculating the averages.

age	avg_gradecompleted	avg_gradesback
21	12.578632	1.4213685
22	12.921179	2.0788207
23	13.112805	2.8871951
24	13.164405	3.8355946
25	13.216009	4.7839912
26	13.043253	5.9567474
27	13.003636	6.9963636

Bonus! If you aren't ready to stop having fun, the steps needed to really make this analysis complete would be to consult the codebook (*.cdb*) file to recode the factor variables into readable values and to appropriately recode missing values using `na_if()`.