

# **LIS Self Teaching Package 2020**

**R version**

## **Part II**

**Gender, employment, and wages**





## Part II: Gender, employment, and wages

### **Overall Plan and Structure of the Exercise**

The exercises in Part I demonstrated the use of household income data along with useful programming techniques for working with the LIS data. Part II emphasises the use of person-level data, including wages, demographics, and labour market information. Whereas Part I consisted entirely of calculating descriptive statistics, Part II introduces regression modelling in the final exercises.

The program that was written in the first set of exercises is now completed and can be set aside. Starting with the next exercise, you will begin the process of building up an entirely new program for Part II. Many of the techniques shown in the previous part will be useful again. In addition, users will learn how to combine LIS datasets by merging household and person files, and by concatenating multiple country-year datasets into a single file.

The general approach of the exercises is the same as in Part I. After beginning a new program in the initial exercise, each subsequent exercise will add new code to the existing program. Within each exercise, results will be produced that help to illuminate the central research themes of this section.

### **Research Questions**

The analysis of poverty and inequality using household income, which was covered in Part I, has always been central to research using LIS data. Over the years, however, there has been an increasing volume of work that examines individual outcomes in the labour market. The richness of the labour market data available in LIS has increased over time, and today it is possible to address many types of questions about wages and employment.

Labour market outcomes for women are one especially popular area of research. Women's rate and intensity of work shows much wider cross-country variation than men's. At the same time, on average women consistently earn lower wages.

For the exercises, we will examine three countries, using data from LIS Wave VI: The United States, Germany, and Greece. As we will see, labour market outcomes for women show distinctly different patterns in each of these countries. Looking at persons of prime working age (which we will define as ages 25-54), our central questions will be:

- How does the percentage of prime-aged women employed in paid work vary across these three countries?
- Among those who are employed, how does the rate of *part-time* employment among women vary across the countries?
- How does employment vary by partnership and family status?
- How do wage differentials between men and women vary between across countries, across levels of educational attainment, and between immigrants and non-immigrants?

In the exercises, we will begin by producing tabulations of employment and wages for various population subgroups. In the concluding exercises, we will use linear regression to study multiple determinants of wages simultaneously, in order to better understand how family structure, education, and immigrant status are related to wages for men and for women.

### **Before you begin**

Before you begin the exercises, take a look at the [2019 Template LIS User Guide](#), which can be accessed through *LIS Website* → *Our Data* → *LIS Database*. The User Guide provides an overview of the structure of LIS data and some data management practices, such as missing values policy and aggregation rules, which will be useful for working with LISSY.

In addition to this, an overview of the datasets and variables is provided through the METadata Information System ([METIS](#)) without having to login to LISSY. You can access METIS via the *LIS Website* → *METIS* → *Enter METIS* → *LIS*. After selecting the datasets and variables, consult the *Results tab* for information on variables and definitions, dataset-specific information and variable availability across datasets.

## Contents

### **1. Merging person and household data, selecting a sample**

- Selecting a sample of prime-age workers
- Homeownership rates by country

### **2. Stacking data, employment rates by gender**

- Employment rates by country and gender
- Part-time employment rates by country and gender

### **3. Family structure and employment**

- Family/partnership status and employment rates

### **4. Dependent employment and hourly wages**

- Non-dependent employment as a proportion of employment, by country
- Gender wage gaps among dependent employees, by country

### **5. Hourly wages, education, and country-specific variables**

- Harmonised vs. country-specific coding
- Gender wage gaps by country and educational attainment

### **6. Immigration and wages, understanding harmonisation**

- Understanding harmonization choices
- Gender wage gaps by country and immigrant status

### **7. Wage regressions**

- Regressing log wages on demographics, separately by gender and country

### **8. Pooled regressions and normalised weights**

- PPP adjustments for income
- Pooled regression with multiple countries, using normalised weights

**Additional guideline on saving your temporary data files at the LIS directory**



# 1. Merging person and household data, selecting a sample

## Goal

While the exercises in Part I only used data at the household level, Part II uses data from both the household and person level files. In this exercise, we will combine the person and household files in order to create a single dataset for each country in which household data is appended to each person record.

This exercise also selects the universe of persons that we will be studying in the subsequent exercises. Since we are interested in labour market outcomes, we will restrict our attention to people of *prime age*: those who are likely to be old enough to have completed schooling and young enough to not yet be retired. In these exercises, we will define prime age persons as those between 25 and 54 years old. This is a commonly used range in statistics from the United States government and other sources, but other definitions are also possible.

Some of the variables we will be using are not always available for household members other than the head and spouse. For that reason, we will further restrict our universe to heads and partners only.

All of the variables that will be needed in the subsequent exercises are introduced here. However, for now we will only analyse one: we will summarize home ownership, which is a household-level variable, in order to measure the rate of homeownership in each of the three countries under analysis. Homeownership will be included in our later multivariate analysis, because it serves as a rough proxy for wealth, which we otherwise have no information about.

## Activity

Go to *LIS website* → [Login LISSY](#) tab with your LISSY account. Write a program to loop through three datasets: United States 2004 (**us04**), Germany 2004 (**de04**) and Greece 2004 (**gr04**). In each country, merge the person file to the household file and keep the following variables:

- Household: Unique household identifier (**hid**) and owned/rented housing indicator (**own**)
- Person: Unique household identifier (**hid**), dataset name (**dname**), normalized person weight (**pwgt**), inflated person weight (**ppopwgt**), relationship to the household head (**relation**), partnership status (**partner**), age of youngest own child living in the household

(**ageyoch**), age (**age**), sex (**sex**), immigrant indicator (**immigr**), 3-category recoded educational attainment (**educ**), country-specific educational attainment (**educ\_c**), indicator for employment (**emp**), status in employment (**status1**), indicator for part-time employment (**ptime1**), and gross hourly wage in the first job (**gross1**).

Keep only those cases that are in the prime age range (between 25 and 54), and which are defined as either household heads or spouses in the variable **relation**.

Create an indicator variable equal to 1 if a person owns their house (with or without a mortgage), and 0 otherwise. Summarize this new variable to find the homeownership rate among the prime-aged persons for each country, and complete the following table:

Dataset	Homeownership %
DE04	
GR04	
US04	

### Questions

- 1.1. In which country do the largest percentage of persons of prime working age own their houses, and in which country are homeownership rates lowest?
- 1.2. Can we consider all non-homeowners as tenants? If not, what other housing tenure status are possible?

### **Guidelines**

- In order to make your code easier to read, it may be helpful to store the list of variables you will be using in global macros. You will need two global macros, one for household level variables and one for person level variables. You can also store the list of datasets in a macro, and refer to it when constructing your loop.
- In order to make your code easier to read, it may be helpful to store the list of variables you will be using in separate vectors, which you can refer to when you open the data file. You will need two vectors, one for household-level variables and one for person level variables. You can also store the list of datasets in a vector, and refer to it when constructing your loop.
- When loading data sets, you can use the parameter "labels" of the



**read.LIS** function to tell R whether to load the data with the LIS-supplied labels, or with numeric codes. For this and subsequent exercises, we will use **labels = FALSE**.

- The simplest way to merge datasets in R is to use the **merge** command. First, load the person and household files into separate data frames, and then combine them. If you have stored the person and household files in data frames called **dp** and **dh**, you can use:

```
df <- merge(dh, dp, by = c('hid'))
```

- Remember that when recoding variables, you can find a listing of the possible values of the original variable in METIS. In this case, go to the LIS Database information (*Enter METIS → LIS*). Select DE04, GR04 and US04. Select the variable *own*. Go to *Results → Crossed-compare* and click on the variable name to see the statistics and labels of the variables.
- Like the household file, the person file contains weight variables. These variables can be used to weight by person, as an alternative to the method of multiplying household weights by number of household members that was used in the Part I. Although home ownership is a household-level variable, you will want to use the person weight to determine the proportion of *persons* who live in owner-occupied dwellings. For now, use the variable **ppopwgt**, which inflates to the total population size.

## **Program**

```
setups <- function(ccyy) {  
  # READ DATASETS  
  varh <- c('hid', 'dname', 'own')  
  varp <- c('hid', 'dname', 'ppopwgt', 'age', 'sex', 'relation', 'emp', 'ptime1', 'ageyoch',  
'partner', 'status1', 'gross1', 'educ', 'immigr')  
  subset <- 'age >= 25 & age <= 54 & relation <= 2200'  
  dh <- read.LIS(paste(ccyy, 'h', sep = ''), labels = FALSE, vars = varh)  
  dp <- read.LIS(paste(ccyy, 'p', sep = ''), labels = FALSE, vars = varp, subset =  
subset)  
  df <- merge(dh, dp, by = c('hid'))  
  # MAP NEW VARIABLES  
  df$home <- ifelse(df$own %in% 100:199, 1, ifelse(df$own %in% 200:299, 0, NA))  
  return(df)  
}  
#----- RUN SCRIPTS -----  
datasets <- c('us04', 'de04', 'gr04')  
for (ccyy in datasets) {  
  df <- setups(paste(ccyy, sep = ''))  
  res <- round(with(df[!is.na(df$home),], sum(home*ppopwgt) /  
sum(ppopwgt[!is.na(home)])) * 100, digits=2)  
  print(c(ccyy, res))  
}
```

## **Results**

Dataset	Homeownership %
DE04	42.7%
GR04	67.6%
US04	71.5%

### *Solutions*

- 1.1. In which country do the largest percentage of persons of prime working age own their houses, and in which country are homeownership rates lowest?
  - Homeownership rates are highest in the United States at 71.5%, and lowest in Germany at only 42.7%.
- 1.2. Can all non-homeowners be considered as tenants? If not, what other housing tenure status are possible?
  - No, in these three datasets there is also a category labelled as "free housing". According to the LIS variable definitions available in [METIS](#) this category can include "housing provided by employer, government or others, or illegal occupation".

### *Comments*

- You will notice that in this exercise the merge worked perfectly, i.e. all observations of the merging file were uniquely linked to one observation in the using file. This is always the case with LIS household and individual level files from the same dataset because all individuals belong to at least one and no more than one household.

## 2. Stacking data, employment rates by gender

### Goal

So far, we have performed all the analysis separately for each dataset, working with only one country at a time. For this and all subsequent analyses, however, we will create a “stacked” dataset that contains information for all three countries in a single file. This means your dataset will have as many value observations as your countries altogether have stored within one file. This may offer you substantial advantages to make use of the data.

After creating a combined dataset, we will examine rates of employment and part-time employment of women, and see how they differ among these three countries. As in the previous exercise, we will be looking only at prime-aged persons who are defined as household heads or partners of household heads.

We will be using the LIS variable **emp**, an indicator that reports whether or not a person is currently employed. This variable will contain the current main employment status (as derived from LIS variable **ifs**). If the current main employment status is not available in the original dataset, employment status during income reference time or employment status according to the ILO criteria in the current period will be used in LIS variable **ifs** and hence also **emp** (you can check dataset-specific notes in [METIS](#) for information about each dataset). By this definition, a person may be considered as employed as soon as he/she has carried out any work.

Rates of employment and full-time employment among prime-age men tend to be similar and consistently high across countries. Due to this, we will be examining differences in employment outcomes among women.

### Activity

Modify your program so that it first creates a combined data file for the United States, Germany, and Greece, and then performs any necessary recoding and produces descriptive statistics.

Create a set of cross-tabulations that shows the rates of prime-age employment of women within each country. Create another set of cross-tabulations showing the rates of part-time work of women within each country, among those who are employed. Use your results to complete the following table below.

You should write your code so that your overall program is broken down into three subroutines. The first subroutine should contain only the code

needed to create the merged, stacked dataset. The second subroutine should contain all of the data-preparation and recoding. The third subroutine should contain the code that produces the summary statistics. Your overall program can then simply call these two subroutines to make the dataset and output the results. Breaking up your code in this way will be important for making the program compact and efficient in later exercises.

Dataset	Female employment rate	Part-time employment rate among employed women
DE04		
GR04		
US04		

### Question

2.1 Contrast these countries in terms of their rates of female employment (high or low) and their rates of part-time employment among employed women (high or low).

### **Guidelines**

- You do not need to remove the code that you used to produce the descriptive in the last exercise (on homeownership), but you can comment it out to make your job run slightly faster. To comment out a line, place a hash (**#**) at the beginning of the line.
- The **read.LIS** command can automatically stack datasets if you call it with a vector of identifiers instead of a single string:

```
df <- read.LIS(c('us04h','de04h','gr04h'))
```

This will return a single data frame containing data from the United States, Germany, and Greece. Note, however, that you cannot mix person and household file identifiers in the same call, or you will get an error. Thus you may want to first create two stacked datasets (one for person and one for household data) and then merge them. If you do this, however, note that you must merge based on two variables, **dname** and **hid**, since the household identifier by itself will no longer uniquely identify cases. This is easily done by passing a vector of merging variables as the third argument of the **merge()** command.

- In base R, weighted proportions do not exist as such. One option is to use the function **tapply(X, INDEX, FUN, ...)** – that applies a function or operation on subset of the vector broken down by a given factor variable (or a discrete numeric variable).

For example, apply the following code to calculate the employment rate among women:

```
with(df[df$sex==2 & !is.na(df$emp),], tapply(emp*ppopwgt,  
list(dname), sum) / tapply(ppopwgt, list(dname), sum)) * 100
```

This code generates the weighted proportion of the **emp** variable subdivided by dataset name for women. In addition, the results are multiplied by 100 to get the percentage.

## **Program**

```
get_stack <- function(datasets, varp, varh, subset) {  
  # READ DATASETS  
  
  pp <- read.LIS(paste(datasets, 'p', sep = ""), labels = FALSE, vars = varp, subset =  
subset)  
  hh <- read.LIS(paste(datasets, 'h', sep = ""), labels = FALSE, vars = varh)  
  df <- merge(pp, hh, by = c("dname", "hid"))  
  return(df)  
}  
  
#----- RUN SCRIPTS -----  
varh <- c('hid', 'dname', 'own')  
varp <- c('hid', 'dname', 'ppopwgt', 'age', 'sex', 'relation', 'emp', 'ptime1')  
subset <- 'age >= 25 & age <= 54 & relation <= 2200'  
datasets <- c('us04', 'de04', 'gr04')  
df <- get_stack(datasets, varp, varh, subset)  
  
'Female Employment Rate'  
round(with(df[df$sex==2 & !is.na(df$emp)], tapply(emp*ppopwgt, list(dname), sum) /  
tapply(ppopwgt, list(dname), sum)) * 100, digits=2)  
  
'Partime Employment rate among employed women'  
round(with(df[df$sex==2 & df$emp==1 & !is.na(df$ptime1)], tapply(ptime1*ppopwgt,  
list(dname), sum) / tapply(ppopwgt, list(dname), sum)) * 100, digits=2)
```

## **Results**

Dataset	Female employment rate	Part-time employment rate among employed women
DE04	72.2%	42.6%
GR04	57.6%	21.3%
US04	72.0%	17.7%

### *Solution*

- 2.1. Contrast these countries in terms of their rates of female employment (high or low) and their rates of part-time employment among employed women (high or low).
- Employment rates among prime age women are relatively high in the United States and in Germany, and lower in Greece. In the United States, most employed women work full time, while nearly half of employed German women work part time. Greece combines low employment rates with high rates of full time employment among those women who are employed.



### 3. Family structure and employment

#### Goal

In the previous exercise we examined cross-national differences in women's employment. In this exercise, we will examine the variation in employment rates among women, based on their partnership and family status. We will contrast partnered and single women. Within each of those two categories, we will contrast women without children in the household, women with young children, and women with older children. The variables created in this exercise will be useful later, when we combine family structure with other personal characteristics in a multivariate analysis of wages.

#### Activity

Since you already created the merged and stacked dataset in the previous exercise, you do not need to create it again. Modify your code so that the subroutine that merges and stacks the data is commented out, and add a line that simply loads the merged and stacked file at the beginning of the program.

Create a variable **achildcat**, to indicate the age of the youngest own child living in the household. This variable should be equal to 0 if there are no children under 18, equal to 1 if the youngest child is under 6 years old, and equal to 2 if the youngest child is between 6 and 17. You can create this variable based on the information in the variable **ageyoch**.

Produce summary statistics to complete the table below using **achildcat**, an indicator of whether a person is currently living with a partner (=1) or not (=0) (**partner**) and the employment indicator you used in the previous exercise (**emp**).

#### **Employment Rates**

Dataset	All women	Single			Partnered		
		No children under 18	Child under 6	Child 6-17	No children under 18	Child under 6	Child 6-17
DE04							
GR04							
US04							

## Question

3.1 Within each country, which subpopulation of prime age women has the lowest employment rates?

## Guidelines

- You will again need to recode LIS variables, as in exercise 2.1. See that exercise for more information. In this exercises, you will be creating variables **achildcat** and **partner**.
- As in the last exercise, you can use the **tapply** function to calculate weighted proportions over subsets of the data. You can extend this command by adding additional variables to the list of categorical variables, which will allow you to analyze subsets within subsets. For example:

```
with(df[df$sex == 1 & !is.na(df$emp), ], tapply(emp * ppopwgt, list(dname, achildcat, partner), sum) / tapply(ppopwgt, list(dname, achildcat, partner), sum)) * 100
```

This code will produce a tabulation of employment rates among women, which is separated by country, by age of the youngest child in the family, and by partnership status.

## **Program**

```
get_stack <- function(datasets, varp, varh, subset) {
  # READ DATASETS

  pp <- read.LIS(paste(datasets, 'p', sep = ''), labels = FALSE, vars = varp, subset =
subset)

  hh <- read.LIS(paste(datasets, 'h', sep = ''), labels = FALSE, vars = varh)
  df <- merge(pp, hh, by = c("dname", "hid"))
  # MAP NEW VARIABLES

df$achildcat <- ifelse(df$ageyoch < 6, 1, ifelse( df$ageyoch > 5 & df$ageyoch < 18, 2,
0))
df$achildcat [is.na(df$achildcat)] <- 0

  return(df)
}
#----- RUN SCRIPTS -----
---

datasets <- c('us04', 'de04', 'gr04')
varh   <- c('hid', 'dname', 'own')
varp   <- c('hid', 'dname', 'ppopwgt', 'age', 'sex', 'relation', 'emp', 'ptime1', 'ageyoch',
'partner', 'status1')
subset <- 'age >= 25 & age <= 54 & relation <= 2200'
df     <- get_stack(datasets, varp, varh, subset)

round(with(df[df$sex == 2 & !is.na(df$emp), ], tapply(emp * ppopwgt, list(dname,
achildcat, partner), sum) / tapply(ppopwgt, list(dname, achildcat, partner), sum)) *
100, digits = 2)
```

## **Results**

### **Employment Rates**

Dataset	All women	Single			Partnered		
		No children under 18	Child under 6	Child 6-17	No children under 18	Child under 6	Child 6-17
DE04	72.2%	82.3%	43.5%	72.8%	80.8%	50.8%	72.0%
GR04	57.6%	68.6%	64.2%	81.4%	49.7%	57.7%	60.5%
US04	72.0%	80.1%	66.9%	78.7%	76.2%	58.4%	71.7%

### *Solutions*

**Question:** Within each country, which subpopulation of prime age women has the lowest employment rates?

- In the United States, partnered women with young children have the lowest employment rates. In Greece, partnered women without children have the lowest employment rates. In Germany, however, single mothers with children have lower employment rates. This may be because of more generous child policy in Germany that makes it easier for mothers of young children to support themselves without paid employment.

### *Comments*

- There is no clear-cut definition of a single-mother household. In this exercise, we allow other adult members to be present (as long as they are not defined as her partner). An alternative approach would be to limit the sample to households composed of a single female adult and her children. Another possibility is to limit single mother households to those with children under a specified age limit.
- When subdividing subsets of the population as has been done here, pay attention to sample sizes. In small datasets, estimates for narrowly defined groups may become very small, making estimates less reliable. The estimate for single Greek women with young children in this exercise, for example, is based on only 17 cases!

## 4. Dependent employment and hourly wages

### Goal

In the next several exercises, we will shift from considering employment to analysing the earnings of those who are employed. We will focus our analysis on the hourly wages and thus restrict our sample to those in dependent employment only — that is, those who are employees. The self-employed, along with several other small categories of workers, are excluded.

In this exercise we will first determine how many workers are excluded from the analysis when the sample is restricted to those in dependent employment. We will then measure the gap in hourly wages for men and women, in each of the three countries in our study.

We will be using a measure of gross hourly wages, which is available in the three datasets we are using. In other datasets, however, it could be that only annual wages were available. In such cases, researchers must account for variations in employment over the year, perhaps by restricting the sample to full-year, full-time workers.

In part I, we have introduced bottom- and top-coding as a technique to deal with extreme values. In this exercise we will also introduce a different technique that deals with extreme values: ‘trimming’ the distribution (i.e. deleting rather than recoding the extreme values). This technique is especially useful when calculating measures that are not defined for non-positive values (such as logarithmic measures). In later exercises we will convert the hourly wages into logs, and thus we need to make sure that the sample that we analyse at this stage is the same that we will keep for our final analysis.

### Activity

Recode the variable **status1** to create a new variable **depemp** that indicates whether a person is in dependent employment. Using this variable, produce summary statistics reporting the proportion of dependent employment among prime-age male workers and among prime-age female workers, and complete the following table.

## Employment Rates

Dataset	Men		Women	
	Non-dependent Employment (%)	Dependent employment (%)	Non-dependent Employment (%)	Dependent employment (%)
DE04				
GR04				
US04				

Next, use the LIS hourly wage variable **gross1** to construct a new hourly wage variable **hourwage**, where the bottom and the top of the distribution are corrected as follows:

- at the top, we will carry out the same top-code as used in Part I: hourly wages that are greater than ten times the national median are set to ten times the national median (make sure you calculate the median and apply the top-code *separately* for each country);
- at the bottom, we will 'trim' the distribution so that observations with negative and zero hourly wages are set to missing.

Using this new hourly wage variable, calculate the *gender wage gap* for dependent employees in each country. The gender wage gap is defined here as the ratio of the median wages of women to the median wages of men. Use your results to complete the table below.

	Gender wage gap for dependent employees
DE04	
GR04	
US04	

### Questions

- 4.1 Does the percentage of workers not in dependent employment differ substantially across countries? Does it differ between men and women?
- 4.2 Which country has the most wage inequality between men and women, among dependent employees?

## **Guidelines**

- Create a new variable for dependent employment using the recoding techniques from the earlier exercises.
- You can use the **wNtile(var, wgt, split)** function used in the part I of the self-teaching package to calculate the weighted median for the purpose of top-coding.

```
topline <- with(df[!is.na(df$hrwg) & df$dname==datasets[i],], 10 *  
wNtile(hrwg, ppopwgt, 0.5))
```

- As in the last exercise, you can use the **tapply** function to calculate the employment rates.
- Use the **wNtile(var, wgt, split)** function to calculate median wages for women and men by country with a for loop technique.

```
for (i in (1:length(datasets))) {  
  for (j in 1:length(unique(df$sex))) {  
    mat[j, i] <- ...  
  }  
}
```

- Then allocate the results to a (2LX3C) matrix and calculate the ratio of the values per line:

```
round(mat[2, ] / mat[1, ], digits = 2)
```

## **Program**

```
get_stack <- function(datasets, varp, varh, subset) {

# READ DATASETS
pp <- read.LIS(paste(datasets,'p',sep="),labels=FALSE, vars=varp, subset = subset)
hh <- read.LIS(paste(datasets,'h', sep="), labels=FALSE, vars=varh)
df <- merge(pp, hh, by = c("dname", "hid"))

# MAP NEW VARIABLES
df$sex <- ifelse(df$sex == 1, 0, 1)
df$dept <- ifelse(df$status1 %in% c(100:120),1,ifelse(is.na(df$status1),NA,0))
df$hrwg <- df$gross1
df$hrwg <- ifelse(df$hrwg <= 0, NA, df$hrwg)
for (i in 1:length(datasets)) {
topline <- with(df[!is.na(df$hrwg) & df$dname==datasets[i],], 10 * wNtile(hrwg,
ppopwgt, 0.5))
df$hrwg <- with(df[!is.na(df$hrwg) & df$dname==datasets[i],], ifelse(df$hrwg >
topline, topline, df$hrwg))
}
return(df)
}

wNtile <- function(var, wgt, split) {
x <- var[order(var)]
y <- wgt[order(var)]
z <- cumsum(y) / sum(y)
cop <- rep(NA,length(split))
for (i in 1:length(cop)) {
cop[i] <- x[Find(function(h) z[h] > split[i], seq_along(z))]
}
return(cop)
}

#----- RUN SCRIPTS -----
datasets <- c('us04', 'de04', 'gr04')
varh <- c('hid', 'did', 'dname', 'own')
varp <- c('hid', 'dname', 'ppopwgt', 'age', 'sex', 'relation', 'emp', 'ptime1', 'ageyoch',
'partner', 'status1', 'gross1')
subset <- 'age >= 25 & age <= 54 & relation <= 2200'
```



```

df    <- get_stack(datasets, varp, varh, subset)

# EMPLOYMENT RATES
print('Employment Rate')
round(with(df[df$emp == 1 & !is.na(df$dept), ], tapply(dept * ppopwgt, list(dname,
sex), sum) / tapply(ppopwgt, list(dname,sex), sum)) * 100, digits=2)

# GENDER WAGE GAP
mat <- matrix(NA, 2 ,3)
colnames(mat) <- datasets
for (i in (1:length(datasets))) {
  for (j in 1:length(unique(df$sex))) {
    mat[j,i] <- with(df[df$dname==datasets[i] & !is.na(df$hrwg) & df$sex== j-1, ],
wNtile(hrwg, ppopwgt, 0.5))
  }
}
'Gender Wage Gap'
round(mat[2, ] / mat[1, ], digits = 2)

```

## **Results**

### **Employment Rates**

Dataset	Men		Women	
	Non-dependent Employment (%)	Dependent employment (%)	Non-dependent Employment (%)	Dependent employment (%)
DE04	12%	88%	8%	92%
GR04	34%	66%	29%	71%
US04	14%	86%	8%	92%

### **Gender wage gaps**

	Gender wage gap for dependent employees
DE04	0.80
GR04	0.85
US04	0.75

### *Solutions*

- 4.1 Does the percentage of workers not in dependent employment differ substantially across countries? Does it differ between men and women?
- Greece has a much higher rate of non-dependent employment (which is primarily self-employment). In all the countries, women have somewhat higher rates of dependent employment than men. Keep in mind, therefore, that the results in the subsequent exercises may be somewhat unrepresentative, particularly for Greece, because they exclude a substantial number of workers.
- 4.2 Which country has the most wage inequality between men and women, among dependent employees?
- The United States shows the largest gender wage gap. Among prime age workers, the median hourly wage of women is only 75% that of men.

### *Comments*

- The wage gap calculated here is based on the median, but some researchers calculate an alternative version based on the mean, which will give slightly different results.

## 5. Hourly wages, education, and country-specific variables

### Goal

This exercise continues the analysis of gender wage gaps in hourly wages among those in dependent employment, which we started to program in the previous exercise. In this exercise, we will see how gender disparities in wages differ by educational attainment.

This exercise also demonstrates the use of two different LIS variables measuring educational attainment. One is fully standardised for cross-national compatibility, but contains few categories. The other may contain more information, but has country-specific codes, and thus requires researchers to perform their own standardisation.

The standardised variable is called **educ**, which is based on the International Standard Classification of Education (ISCED). The non-standardised version (from which **educ** is constructed) is **educ\_c**. This is one of many attributes for which LIS provides both a standardised and country-specific variable. Any variable ending in **\_c** is non-standardised, meaning that it can have different contents in different datasets. It is important to carefully examine the dataset-specific documentation before using such variables.

### Activity

Add code to your program to create a table cross-tabulating the variables **educ** and **educ\_c** for each country. This will show how the standardised variable was constructed in each case. Be sure to:

- include missing values in your table, so that you can see whether any of the cases in the original education variable could not be allocated to a category in the standardised version;
- remove the value labels from the tabulation of the **educ\_c** variable (since the value labels of the **\_c** variables are by definition dataset-specific, in a stacked dataset with observations from several LIS datasets, the value labels of those variables will be incorrect, as they can only refer to one specific LIS dataset – usually the last one that was used to construct the stacked data, see further details in the comments section of this exercise).

Using the hourly wage variable **hourwage** that you created in the last exercise, calculate the gender wage gap by education for each country, and complete the table below. The gender wage gap is defined as it was in the

previous exercise. To obtain the earnings ratio by education, simply calculate the ratio separately for individuals in each of the three categories of the standardised education variable.

### Gender wage gaps by educational attainment

	Low education	Medium education	High education
DE04			
GR04			
US04			

### Questions

- 5.1 For each of the three countries, what are the categories in the original dataset that are recoded as “high education” in the standardised education variable?
- 5.2 Are there any categories in the original **educ\_c** variable that could not be translated into the standardised form?
- 5.3 In general, which educational attainment category shows the greatest earnings inequality between genders? How do the patterns differ by country?

### Guidelines

- Because the program so far has been designed to load variables without their labels, you will have to consult [METIS](#) for variable information and codebook in order to determine what the country-specific education codes refer to. If you prefer, you could write your program to load the education variables with labels before tabulating, but this may be somewhat redundant and burdensome because of the way the program has been designed so far.
- As in the previous exercises, you can produce tabulations of the wage ratio for each country by using **tapply()** function and loop technics.

## **Program**

```
get_stack <- function(datasets, varp, varh, subset) {
  # READ DATASETS
  pp <- read.LIS(paste(datasets, 'p', sep = ''), labels = FALSE, vars = varp, subset =
subset)
  hh <- read.LIS(paste(datasets, 'h', sep = ''), labels = FALSE, vars = varh)
  df <- merge(pp, hh, by = c("dname", "hid"))
  # MAP NEW VARIABLES
  df$sex <- ifelse(df$sex == 1, 0, 1)
  df$hrwg <- df$gross1
  df$hrwg <- ifelse(df$hrwg <= 0, NA, df$hrwg)
  for (i in 1:length(datasets)) {
    topline <- with(df[!is.na(df$hrwg) & df$dname == datasets[i], ], 10 * wNtile(hrwg,
ppopwgt, 0.5))
    df$hrwg <- with(df[!is.na(df$hrwg) & df$dname == datasets[i], ], ifelse(df$hrwg >
topline, topline, df$hrwg))
  }
  return(df)
}
wNtile <- function(var, wgt, split) {
  x <- var[order(var)]
  y <- wgt[order(var)]
  z <- cumsum(y) / sum(y)
  cop <- rep(NA, length(split))
  for (i in 1:length(cop)) {
    cop[i] <- x[Find(function(h) z[h] > split[i], seq_along(z))]
  }
  return(cop)
}
#----- RUN SCRIPTS -----
datasets <- c('us04', 'de04', 'gr04')
varh <- c('hid', 'dname', 'own')
varp <- c('hid', 'dname', 'ppopwgt', 'age', 'sex', 'relation', 'emp', 'ptime1', 'ageyoch',
'partner', 'status1', 'gross1', 'educ')
subset <- 'age >= 25 & age <= 54 & relation <= 2200'
df <- get_stack(datasets, varp, varh, subset)
```

```

# GENDER WAGE GAP
ctry_list <- list()
for (k in (1:length(datasets))) {
  mat <-
  matrix(NA,length(unique(df$sex[!is.na(df$sex)]),length(unique(df$educ[!is.na(df$educ)
])))
  colnames(mat) <- c('Low', 'Medium', 'high')
  for (j in (1:length(unique(df$educ[!is.na(df$educ)])))) {
    for (i in (1:length(unique(df$sex[!is.na(df$sex)])))) {
      mat[i,j] <- with(df[df$dname == datasets[k] & !is.na(df$hrwg) & !is.na(df$sex) &
!is.na(df$educ) & df$sex == i - 1 & df$educ == j, ], wNtile(hrwg, ppopwgt, 0.5))
    }
  }
  ctry_list[[datasets[k]]] <- round(mat[2, ] / mat[1, ], digits = 2)
}

```

'Gender Wage Gap'

ctry\_list

## Results

### Gender wage gaps by educational attainment

	Low education	Medium education	High education
DE04	0.78	0.83	0.81
GR04	0.67	0.78	0.96
US04	0.71	0.71	0.78

### *Solutions*

- 5.1. For each of the three countries, what are the categories in the original dataset that are recoded as “high education” in the standardised education variable?
- In the United States, high education combines those with associate degrees, bachelor's degrees, and advanced degrees (masters, professional school, or doctorate).
  - In Germany, high education combines those classified as having “higher vocational” (of any kind) or “higher education”.
  - In Greece, high education includes those with tertiary graduate level education, postgraduate level education or Ph.D.
- 5.2. Are there any categories in the original **educ\_c** variable that could not be translated into the standardised form?
- In the United States, all values of **educ\_c** receive a value in **educ**. In Germany, a small number of persons categorised as “inadequately, other diploma” or “still in education” are set to missing. In Greece, a small number of persons listed as “still in education” are set to missing.
- 5.3. In general, which educational attainment category shows the greatest earnings inequality between genders? How do the patterns differ by country?
- In all three countries, there is a smaller gender wage gap among highly educated workers. This is particularly notable in Greece. In that country, wage inequality is greater among the low-educated than in the United States and Germany, but there is near equality among the highly educated.

### *Comments*

- You may have some doubts and questions why we did advise you in the exercise to not show the labels of **educ\_c** for the cross-tabulation of

**educ** and **educ\_c**. As you are aware we did append values for several countries for each variable to get the stacked file. In your stacked file, for standardised variables these values have all the same meaning, as the values and labels are completely standardised. However, it is more complicated for non-standardized values and labels of **\_c** variables, as each dataset has its own values and own meaning, as indicated by the labels attached to the data.

- Be aware that while appending the data, your programming software will very likely overwrite the label automatically. Thus we do in general advise you to drop the labels from the variables **\_c** within your code. You always have the full information on the labels in the codebooks. However, if you prefer to keep the full labels somewhere in the data there are several solutions.
- A simple solution is to tabulate each country separately (see for example exercise 2.1 of part II) before you generate the stacked version. Alternatively, you can also easily rename the variables **\_c** to the specific **\_`ccyy'** of each dataset - this way you will append a separate variable **\_`ccyy'** for each of your datasets, which does have only observations for the specific **`ccyy'** with the country specific labels attached. Be aware that you then need to tabulate for each **`ccyy'** separately to get the right percentage of missing values!



## 6. Immigration and wages, understanding harmonisation

### Goal

Each of the countries we are examining has a significant immigrant population, and their labour market outcomes are often very different from those of the non-immigrant population. In this exercise we will compare the wages of immigrants and non-immigrant men and women, just as we compared individuals of different educational levels in the last exercise.

LIS provides a variable indicating whether someone is an immigrant, which we will be using in this exercise. However, the choices that go into constructing this variable are complex, because the information available to construct it varies widely from country to country. It is important to understand the assumptions behind variables such as this one, because in some cases researchers may prefer to develop their own standardisation procedures based on their particular needs.

### Activity

Using the top and bottom-coded hourly wage variable, calculate the gender earnings ratio by immigration status for each country, and complete the table below. The gender earnings ratio is computed just as in the previous exercise, except that you will now subdivide the population into immigrants and non-immigrants, rather than by educational attainment categories.

#### **Gender earnings ratios by immigration status**

	Non-immigrants	Immigrants
DE04		
GR04		
US04		

### *Questions*

- 6.1 What information is used to construct the **immigr** variable? If you wanted to determine how the indicator was constructed in a particular dataset, what other variables would you need to look at?
- 6.2 Is gender earnings inequality larger among immigrants or non-immigrants? Does this differ by country?

## **Guidelines**

- The coding required for this exercise is essentially the same as the one used in the previous exercise.

## **Program**

```
get_stack <- function(datasets, varp, varh, subset) {
# READ DATASETS
pp <- read.LIS(paste(datasets, 'p', sep = ""), labels=FALSE, vars=varp, subset=subset)
hh <- read.LIS(paste(datasets, 'h', sep = ""), labels = FALSE, vars = varh)
df <- merge(pp, hh, by = c("dname", "hid"))

# MAP NEW VARIABLES
df$sex <- ifelse(df$sex == 1, 0, 1)
df$dept <- ifelse(df$status1 %in% c(100:120), 1, ifelse(is.na(df$status1) , NA, 0))
df$hrwg <- df$gross1
df$hrwg <- ifelse(df$hrwg <= 0, NA, df$hrwg)
for (i in 1:length(datasets)) {
  topline <- with(df[!is.na(df$hrwg) & df$dname==datasets[i],], 10*wNtile(hrwg,
  ppopwgt,0.5))
  df$hrwg <- with(df[!is.na(df$hrwg) & df$dname==datasets[i],], ifelse(df$hrwg >
  topline, topline, df$hrwg))
}
return(df)
}

wNtile <- function(var, wgt, split) {
  x <- var[order(var)]
  y <- wgt[order(var)]
  z <- cumsum(y) / sum(y)
  cop <- rep(NA,length(split))
  for (i in 1:length(cop)) {
    cop[i] <- x[Find(function(h) z[h] > split[i], seq_along(z))]
  }
  return(cop)
}
```

```

#----- RUN SCRIPTS -----
datasets <- c('us04', 'de04', 'gr04')
varh   <- c('hid', 'dname', 'own')
varp   <- c('hid', 'dname', 'ppopwgt', 'age', 'sex', 'relation', 'emp', 'ptime1', 'ageyoch',
'partner', 'status1', 'gross1', 'educ', 'immigr')
subset <- 'age >= 25 & age <= 54 & relation <= 2200'
df     <- get_stack(datasets, varp, varh, subset)

# GENDER WAGE GAP
ctry_list <- list()
for(k in (1:length(datasets))) {
  mat <- matrix(NA, length(unique(df$sex[!is.na(df$sex)])),
length(unique(df$immigr[!is.na(df$immigr)]))
colnames(mat) <- c('Non Immigrant', 'Immigrant')
for(j in (1:length(unique(df$immigr[!is.na(df$immigr)])))) {
  for(i in (1:length(unique(df$sex[!is.na(df$sex)])))) {
    mat[i,j] <- with(df[df$dname == datasets[k] & !is.na(df$hrwg) & !is.na(df$sex)
& !is.na(df$immigr) & df$sex== i-1 & df$immigr== j-1,],
wNtile(hrwg,ppopwgt,0.5))
  }
}
ctry_list[[datasets[k]]] <- round(mat[2, ] / mat[1, ], digits = 2)
}

'Gender Wage Gap'
ctry_list

```

## **Results**

### **Gender earnings ratios by immigration status**

	Non-immigrants	Immigrants
DE04	0.80	0.77
GR04	0.89	0.82
US04	0.76	0.82

### *Solutions*

- 6.1 Question: What information is used to construct the **immigr** variable? If you wanted to determine how the indicator was constructed in a particular dataset, what other variables would you need to look at?
- As indicated in the Variable Definition of **immigr** using [METIS](#) documentation tool, Immigrants are defined by LIS as all persons who have that country as country of usual residence and (in order of priority):
    - whom the data provider defined as immigrants;
    - who self-define them-selves as immigrants;
    - who are the citizen/national of another country;
    - who were born in another country.
  - The definition of immigrant used in **immigr** may differ substantially from dataset to dataset. The variables that may be used in its construction include **citizen**, **ctrybrth**, **ysresid**, **ethnic\_c** and **immigr\_c** (you can check dataset-specific notes in [METIS](#) for information about each dataset).
- 6.2 Question: Is gender earnings inequality larger among immigrants or non-immigrants? Does this differ by country?
- In the United States, the gender wage gap is greater among non-immigrants, but in Germany and Greece it is greater among immigrants.

## 7. Wage regressions

### Goal

We have seen how employment varies by gender and family structure, and how gender wage gaps vary by education and immigration status. In this exercise, we will investigate the impact of all these variables on wages, using a multivariate regression.

Wages are generally not normally distributed. We will therefore apply a logarithmic transformation in order to create an outcome variable that is approximately normal, which is more suitable for regression modelling.

In addition to the variables we have already seen, we will also control for age, which has a strong relationship with earnings. Since the relationship between age and income is not necessarily linear, we will also add a term for age-squared.

### Activity

If you have followed the instructions up to this point, you should not need to create any additional variables to run regression models predicting **logwage**.

You should run the regressions separately in each country. In addition, within each country you should run a separate model for men and women. Please use the following model:

**logwage = f (age agesq meduc heduc immigr partn ychild ochild ptime1 homeowner)**

Produce a table of the six resulting models, with coefficients, standard errors, sample sizes, and r-squared values.

### *Questions*

- 7.1 Who receives a higher wage premium from being highly educated, men or women?
- 7.2 When controlling for other individual characteristics, what is the relationship between immigrant status and wages?
- 7.3 When controlling for other individual characteristics, do women with young children make more or less than women without children?

### Guidelines

- Linear regression in R is done with the **glm()** function. Since this

command accepts weights, you can use it instead of the `svyglm()` function from the **survey** package, which is necessary for more complex sample designs.

- R will automatically detect categorical variables if they are coded as factors, meaning that you do not need to manually create dummy variables. You can also include mathematical operations directly in the definition of your regression formula if you enclose them in the `I()` function, so you do not need to manually create log wages or age squared. For example, the following estimates the log wage regression for men in the United States:

```
glm(I(log(hrwg))~age+I(age^2)+educ+immigr+partner+achildcat+ptime1  
+homeowner, data=df, weights=df$ppopwgt, subset=sex=="Male" &  
dname=="us04")
```

- When performing several regression models in a single program, one strategy is to again write a loop that estimates each regression in turn, and then prints a summary of the model

```
model <- formula(I(log(hrwg))~age +  
I(age^2)+meduc+heduc+immigr+partn+ychild+ochild+ptime1+homeowne  
r)  
for(s in ...) for(d in datasets) {  
  res <- glm(model, data = df, weight = df$ppopweight, subset = sex &  
dname = d)  
  print(summary(res))  
}
```

Here the definition of the model has been done before calling the loop, which makes the call to `glm()` itself somewhat easier to read. The summary function will output coefficients, standard errors and R-squared statistics.

## **Program**

```
get_stack <- function(datasets, varp, varh, subset) {
# READ DATASETS
  pp <- read.LIS(paste(datasets, 'p', sep=""), labels=FALSE, vars=varp, subset=subset)
  hh <- read.LIS(paste(datasets, 'h', sep = ""), labels = FALSE, vars = varh)
  df <- merge(pp, hh, by = c("dname", "hid"))

# MAP NEW VARIABLES
df$homeowner <- ifelse(df$own %in% 100:199, 1, ifelse(df$own %in% 200:299, 0,
NA))
df$achildcat <- ifelse(df$ageyoch < 6, 1, ifelse( df$ageyoch > 5 & df$ageyoch < 18, 2,
0))
df$achildcat [is.na(df$achildcat)] <- 0
df$ychild <- ifelse(df$achildcat==1, 1, ifelse(is.na(df$achildcat), NA, 0))
df$ochild <- ifelse(df$achildcat==2, 1, ifelse(is.na(df$achildcat), NA, 0))
df$meduc <- ifelse(df$educ==2 , 1, ifelse(is.na(df$educ), NA, 0))
df$heduc <- ifelse(df$educ==3 , 1, ifelse(is.na(df$educ), NA, 0))
df$hrwg <- df$gross1
df$hrwg <- ifelse(df$hrwg <= 0, NA, df$hrwg)
for (i in 1:length(datasets)) {
  topline <- with(df[!is.na(df$hrwg) & df$dname==datasets[i],],
  10*wNtile(hrwg,ppopwgt,0.5))
  df$hrwg <- with(df[!is.na(df$hrwg) & df$dname==datasets[i],],
  ifelse(df$hrwg>topline, topline, df$hrwg))
}
  return(df)
}

wNtile <- function(var, wgt, split) {
  x <- var[order(var)]
  y <- wgt[order(var)]
  z <- cumsum(y) / sum(y)
  cop <- rep(NA,length(split))
  for (i in 1:length(cop)) {
    cop[i] <- x[Find(function(h) z[h] > split[i], seq_along(z))]
  }
  return(cop)
}
```

```

}
#----- RUN SCRIPTS -----
options(scipen = 2)
datasets <- c('us04', 'de04', 'gr04')
varh <- c('hid', 'dname', 'own')
varp <- c('hid', 'dname', 'ppopwgt', 'age', 'sex', 'relation', 'emp', 'ptime1',
         'ageyoch', 'partner', 'status1', 'gross1', 'educ', 'immigr')
subset <- 'age >= 25 & age <= 54 & relation <= 2200'
df <- get_stack(datasets, varp, varh, subset)

# REGRESSION MODEL
gender <- c('Male', 'Female')
model <- formula(I(log(hrwg))~age + I(age^2) +
meduc+heduc+immigr+partner+ychild+ochild+ptime1+homeowner)
for(s in (1:length(unique(df$sex[!is.na(df$sex)])))) for(d in datasets) {
  res <- glm(model, data = df, weights = df$ppopwgt, subset = sex == s & dname ==
d)
  print('-----')
  print(paste(gender[s], d, sep = " : "))
  print(summary(res), digits=1)
}

```



## Results

```
[1] "-----"  
[1] "Male : us04"
```

Coefficients:

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	1.14388	0.08409	14	<2e-16 ***
age	0.04854	0.00435	11	<2e-16 ***
I(age^2)	-0.00049	0.00005	-9	<2e-16 ***
meduc	0.27898	0.01235	23	<2e-16 ***
heduc	0.66048	0.01245	53	<2e-16 ***
immigr	-0.06060	0.00920	-7	5e-11 ***
partner	0.04106	0.00972	4	2e-05 ***
ychild	0.05658	0.00976	6	7e-09 ***
ochild	0.06317	0.00878	7	6e-13 ***
ptime1	-0.30601	0.01799	-17	<2e-16 ***
homeowner	0.24114	0.00850	28	<2e-16 ***

---

Signif. codes: 0 '\*\*\*' 0.001 '\*\*' 0.01 '\*' 0.05 '.' 0.1 ' ' 1  
Null deviance: 17260992 on 29985 degrees of freedom

```
[1] "-----"  
[1] "Male : de04"
```

Coefficients:

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	1.2101	0.2053	6	4e-09 ***
age	0.0545	0.0104	5	2e-07 ***
I(age^2)	-0.0006	0.0001	-5	3e-06 ***
meduc	0.1230	0.0300	4	4e-05 ***
heduc	0.3997	0.0313	13	<2e-16 ***
immigr	-0.0669	0.0237	-3	0.005 **
partner	0.0240	0.0206	1	0.243
ychild	0.0691	0.0238	3	0.004 **
ochild	0.0465	0.0207	2	0.025 *
ptime1	-0.3008	0.0370	-8	6e-16 ***
homeowner	0.1067	0.0163	6	7e-11 ***

---

Signif. codes: 0 '\*\*\*' 0.001 '\*\*' 0.01 '\*' 0.05 '.' 0.1 ' ' 1  
Null deviance: 2558852 on 3651 degrees of freedom

```
[1] "-----"  
[1] "Male : gr04"
```

Coefficients:

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	0.8121	0.3071	2.6	0.008 **
age	0.0328	0.0158	2.1	0.038 *
I(age^2)	-0.0002	0.0002	-1.2	0.234
meduc	0.1676	0.0249	6.7	3e-11 ***
heduc	0.4311	0.0277	15.6	<2e-16 ***
immigr	-0.2475	0.0320	-7.7	2e-14 ***
partner	-0.0139	0.0367	-0.4	0.705
ychild	0.1195	0.0323	3.7	2e-04 ***
ochild	0.1140	0.0292	3.9	1e-04 ***
ptime1	0.3258	0.0502	6.5	1e-10 ***
homeowner	0.0349	0.0230	1.5	0.130

---  
Signif. codes: 0 '\*\*\*' 0.001 '\*\*' 0.01 '\*' 0.05 '.' 0.1 ' ' 1  
Null deviance: 178954 on 1174 degrees of freedom

```
[1] "-----"  
[1] "Female : us04"
```

Coefficients:

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	1.12401	0.08909	13	<2e-16 ***
age	0.04156	0.00462	9	<2e-16 ***
I(age^2)	-0.00044	0.00006	-8	3e-14 ***
meduc	0.32083	0.01515	21	<2e-16 ***
heduc	0.74893	0.01521	49	<2e-16 ***
immigr	-0.01975	0.01050	-2	0.06 .
partner	-0.01027	0.00845	-1	0.22
ychild	0.03743	0.01083	4	5e-04 ***
ochild	-0.04907	0.00851	-6	8e-09 ***
ptime1	-0.17588	0.00929	-19	<2e-16 ***
homeowner	0.18423	0.00911	20	<2e-16 ***

---  
Signif. codes: 0 '\*\*\*' 0.001 '\*\*' 0.01 '\*' 0.05 '.' 0.1 ' ' 1  
Null deviance: 16539528 on 28596 degrees of freedom

```
[1] "-----"  
[1] "Female : de04"
```

Coefficients:

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	1.2569	0.2335	5.4	8e-08 ***
age	0.0507	0.0120	4.2	3e-05 ***
I(age^2)	-0.0006	0.0002	-3.8	2e-04 ***
meduc	0.0824	0.0301	2.7	0.006 **
heduc	0.3811	0.0327	11.7	<2e-16 ***
immigr	-0.0698	0.0251	-2.8	0.005 **
partner	-0.0162	0.0206	-0.8	0.432
ychild	-0.0030	0.0297	-0.1	0.919
ochild	-0.1173	0.0217	-5.4	6e-08 ***
ptime1	-0.1055	0.0190	-5.5	3e-08 ***
homeowner	0.0707	0.0188	3.8	2e-04 ***

```
---  
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1  
Null deviance: 2988987 on 3574 degrees of freedom
```

```
[1] "-----"  
[1] "Female : gr04"
```

Coefficients:

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	-0.2595	0.3414	-0.8	0.447
age	0.0707	0.0180	3.9	9e-05 ***
I(age^2)	-0.0007	0.0002	-2.9	0.004 **
meduc	0.3472	0.0337	10.3	<2e-16 ***
heduc	0.7634	0.0350	21.8	<2e-16 ***
immigr	-0.2595	0.0408	-6.4	3e-10 ***
partner	-0.0675	0.0351	-1.9	0.055 .
ychild	0.0577	0.0381	1.5	0.130
ochild	0.0256	0.0332	0.8	0.440
ptime1	0.2257	0.0310	7.3	7e-13 ***
homeowner	-0.0083	0.0275	-0.3	0.761

```
---  
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1  
Null deviance: 205347 on 965 degrees of freedom
```

## *Solutions*

- 7.1 Who receives a higher wage premium from being highly educated, men and women?
- In the US and Greece, the coefficient for high education is higher for women, indicating a larger wage premium from having high educational attainment, while in Germany the premium is slightly higher for men.
- 7.2 When controlling for other individual characteristics, what is the relationship between immigrant status and wages?
- The association between immigrant status and wages is negative in all three countries. In general, the negative association appears stronger in Germany than in the United States, and stronger in Greece than Germany.
- 7.3 When controlling for other individual characteristics, do women with young children make more or less than women without children?
- In general, women with young children have higher wages than women without children in the United States and Greece, but there is no association in Germany. Higher wages for women with young children could be due to a selection effect, where mothers of young children are more likely to enter the labour market if they have higher earning power.

## *Comments*

- As we have seen, employment rates, particularly among women, vary substantially across countries. The wage regressions shown here do not account for this differential selection into employment. For this reason, many studies of wages apply a technique such as a Heckman correction, which attempts to correct for this selection bias.

## 8. Pooled regressions and normalised weights

### Goal

In the previous exercise, we ran parallel, separate regressions for each country. In this exercise, we see an alternative approach, in which all countries are pooled together in a single model. We will continue to use classical OLS regression, but the approach shown here can easily be extended for more complex multilevel estimation approaches.

The income variables in these datasets use different currencies. To compare them, we need to convert them to a common scale. We will apply *Purchasing Power Parity* (PPP) deflators, which are intended to ensure that equal quantities of income correspond to equivalent purchasing power across currencies and national economies. The PPP deflators provided here are taken from the World Bank Development Indicators and are constantly updated. In addition, LIS provides the LIS PPPs, which combine CPI and PPP deflators and can be used to compare real amounts across countries and over time. Since May 2020, with the release of the results from the 2017 cycle by the International Comparison Program ([ICP](#)), two sets of PPP deflators are available by LIS via LISSY: 2017 PPPs and the revised 2011 PPPs. In order to convert LIS monetary values into 2017 USD PPPs, amounts expressed in nominal currency should be divided by the LIS PPP of the corresponding year. For more information on the CPI and PPP deflators, please see [here](#), and check our tutorial videos on [Price deflators in LIS](#), and [Price deflators in LISSY using Stata](#).

Up to this point, we have been using the weight variable **ppopwgt**, which inflates to the total population. If we use this weight in a pooled regression, every household will receive equal weight. However, this would mean that Greece — which has a much smaller population than the United States or Germany — will not have much influence on the results. In order to give each country equal weight in the model, we will use the alternative normalized weight variable **pwgt**, which always sums to 10,000 within each dataset.

### Activity

Adjust wages for Purchasing Power Parities by dividing **hrwage** by the following deflators before taking the natural logarithm (see <https://www.lisdatacenter.org/resources/ppp-deflators/>).

dataset	PPP deflator
US04	1
DE04	0.79
GR04	0.65

Estimate the following model, for men and women separately:

**logwage = f (age agesq mededuc hieduc immigr partner youngchild oldchild ptime1 homeowner germany greece)**

The model is the same as in the last exercise, except that it includes an indicator for country. This time, however, make sure to use *normalised*, not inflated weights.

Produce a table of the two resulting models, with coefficients, standard errors, sample sizes, and r-squared values.

### *Questions*

- 8.1. How can you interpret the meaning of the coefficients for the dummy variables for Germany and Greece?
- 8.2. In this pooled model, which carries a higher wage penalty: being an immigrant, or working part time?

### **Guidelines**

- Run your regressions as you did in the previous exercise. This time, only two models need to be produced, so you may not want to use a loop.

## **Program**

```
get_stack <- function(datasets, varp, varh, subset) {
  # READ DATASETS
  pp <- read.LIS(paste(datasets, 'p', sep = ''), labels = FALSE, vars = varp, subset =
subset)
  hh <- read.LIS(paste(datasets, 'h', sep = ''), labels = FALSE, vars = varh)
  df <- merge(pp, hh, by = c("dname", "hid"))
  # MAP NEW VARIABLES
  df$homeowner <- ifelse(df$own %in% 100:199, 1, ifelse(df$own %in% 200:299, 0,
NA))
  df$achildcat <- ifelse(df$ageyoch < 6, 1, ifelse( df$ageyoch > 5 & df$ageyoch < 18,
2, 0))
  df$achildcat [is.na(df$achildcat)] <- 0
  df$ychild <- ifelse(df$achildcat == 1, 1, ifelse(is.na(df$achildcat), NA, 0))
  df$ochild <- ifelse(df$achildcat == 2, 1, ifelse(is.na(df$achildcat), NA, 0))
  df$meduc <- ifelse(df$educ == 2, 1, ifelse(is.na(df$educ), NA, 0))
  df$heduc <- ifelse(df$educ == 3, 1, ifelse(is.na(df$educ), NA, 0))
  df$ppp <- ifelse(df$dname == 'de04', 0.79, ifelse(df$dname == 'gr04', 0.65, 1))
  df$hrwg <- df$gross1
  df$hrwg <- ifelse(df$hrwg <= 0, NA, df$hrwg)
  df$germany <- ifelse (df$dname == 'de04', 1, 0)
  df$greece <- ifelse (df$dname == 'gr04', 1, 0)
  for(i in 1:length(datasets)) {
    df$hrwg <- df$hrwg / df$ppp
    topline <- with(df[!is.na(df$hrwg) & df$dname==datasets[i],], 10 *
wNtile(hrwg,ppopwgt,0.5))
    df$hrwg <- with(df[!is.na(df$hrwg) & df$dname==datasets[i],],
ifelse(df$hrwg>topline,topline,df$hrwg))
  }
  return(df)
}

wNtile <- function(var, wgt, split) {
  x <- var[order(var)]
  y <- wgt[order(var)]
  z <- cumsum(y) / sum(y)
  cop <- rep(NA,length(split))
  for(i in 1:length(cop)) {
```

```

    cop[i] <- x[Find(function(h) z[h] > split[i], seq_along(z))]
  }
  return(cop)
}
#----- RUN SCRIPTS -----
options(scipen = 2)
datasets <- c('us04', 'de04', 'gr04')
varh <- c('hid', 'dname', 'own')
varp <- c('hid', 'dname', 'pwgt', 'ppopwgt', 'age', 'sex', 'relation', 'emp', 'ptime1', 'ageyoch',
         'partner', 'status1', 'gross1', 'educ', 'immigr')
subset <- 'age >= 25 & age <= 54 & relation <= 2200'
df <- get_stack(datasets, varp, varh, subset)

# REGRESSION MODEL
gender <- c('Male', 'Female')
model <- formula(I(log(hrwg))~age + I(age^2) + meduc + heduc + immigr + partner
+ ychild + ochild + ptime1 + homeowner + germany + greece)
for(s in (1:length(unique(df$sex[!is.na(df$sex)])))) {
  res <- glm(model, data = df, weights = df$pwgt, subset = sex == s)
  print('-----')
  print(gender[s])
  print(summary(res), digits=1)
}

```



## Results

[1] "-----"

[1] "Male"

Coefficients:

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	1.25668	0.07287	17	<2e-16 ***
age	0.05192	0.00376	14	<2e-16 ***
I(age^2)	-0.00052	0.00005	-11	<2e-16 ***
meduc	0.18727	0.00867	22	<2e-16 ***
heduc	0.52872	0.00903	59	<2e-16 ***
immigr	-0.11822	0.00800	-15	<2e-16 ***
partner	0.03734	0.00805	5	4e-06 ***
ychild	0.07608	0.00830	9	<2e-16 ***
ochild	0.07070	0.00742	10	<2e-16 ***
ptime1	-0.12782	0.01369	-9	<2e-16 ***
homeowner	0.13917	0.00630	22	<2e-16 ***
germany	0.51068	0.00664	77	<2e-16 ***
greece	0.45812	0.00727	63	<2e-16 ***

---

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

Null deviance: 1292.80 on 34822 degrees of freedom

[1] "-----"

[1] "Female"

Coefficients:

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	0.96161	0.07801	12	<2e-16 ***
age	0.05609	0.00407	14	<2e-16 ***
I(age^2)	-0.00060	0.00005	-12	<2e-16 ***
meduc	0.24292	0.01019	24	<2e-16 ***
heduc	0.65078	0.01058	62	<2e-16 ***
immigr	-0.07796	0.00888	-9	<2e-16 ***
partner	-0.01296	0.00730	-2	0.08 .
ychild	0.01700	0.00934	2	0.07 .
ochild	-0.06603	0.00740	-9	<2e-16 ***
ptime1	-0.05801	0.00704	-8	<2e-16 ***
homeowner	0.08530	0.00691	12	<2e-16 ***
germany	0.60560	0.00742	82	<2e-16 ***
greece	0.61338	0.00810	76	<2e-16 ***

---

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

Null deviance: 1347.89 on 33145 degrees of freedom

## *Solutions*

- 8.1 How can we interpret the meaning of the coefficients for the dummy variables for Germany and Greece?
- These coefficients represent the overall national level of PPP-adjusted wages, when controlling for the other variables. The small positive value for Germany reflects the fact that Germany is a slightly richer country than the United States, while the negative value for Greece reflects the fact that Greece is poorer.
- 8.2 In this pooled model, which carries a higher wage penalty: being an immigrant, or working part time?
- For women, the wage penalty for working part time is smaller than that for being an immigrant. The reverse is true for men, although it should be noted that part time work is very uncommon among prime age men.

## Additional guideline on saving your temporary data files at the LIS directory

When working on your datasets in R, you may need to save some data files that you created from LIS/LWS files for a latter (re-)analysis. You can save these data files via LISSY in the LIS directory (see the code below). In order to ensure that your data file(s) are not overwritten by other users, you should save your file(s) with a unique name (e.g. you could include your LIS username in the filename). You can also ask the LIS User Support to create a specific folder in the LIS directory where you will be able to save your data files (created from LIS/LWS data) via LISSY.

### **Program**

```
require(foreign)
require(readstata13)
#read LIS data
data <- read.LIS('lu10h')
#save the above data in your folder that LIS created for you*
save.dta13(data,paste(USR_DIR,"/name_of_your_folder/name_of_your_data_file_that_you_want_to_save.dta", sep = ""), version = 117,convert.factors=FALSE)
#open the above data that you've just saved
data1 <-read.dta13(paste(USR_DIR,
"/name_of_your_folder/name_of_your_data_file_that_you_saved_in_your_folder.dta",
sep=""),convert.factors=FALSE)
```

Note: in order to have your specific folder, please write to the LIS user support at [usersupport@lisdatacenter.org](mailto:usersupport@lisdatacenter.org) .