

Computer assignment #3

Regression Analysis Assignment/ Regression Procedures for Final Paper

SOLELY FOR THOSE USING LINEAR REGRESSION (OLS)

If you are using “logistic regression” for your final paper, consult the other assignment posted.

At this point in the term, you should have a pretty good sense as to which variables and what dataset you will be working with in your final research paper. You should also have a reasonable sense as to the principal hypotheses that you are putting to empirical test.

The purpose of the current assignment is twofold.

First, it will provide you with a simple strategy for dealing with sample weights when conducting multivariate analyses (sample weights relate to the sample design: I will talk further about them in class). Second, it will provide you with a strategy for conducting your final regression analysis for your final paper. The results from this assignment can be directly plugged into your final paper.

Part A: Weighting

What do I mean by “un-weighted” as opposed to “weighted” data?

The principle behind estimation in a probability sample such as the GSS, CCHS or NLSCY or the public use sample from the census (3%) is that each person in the sample "represents", besides himself or herself, several other persons not in the sample. For example, in a 2% simple random sample of the population, each person in the sample represents 50 persons in the population. The weighting phase is a step which calculates, for each record, what this number is (i.e., the number of individuals in the population represented by this record). The relevant weights appear on the respective datasets as variables. **In the NLSCY microdata file the weight is called “AWTCW01”. With the Census, the name of the relevant weight is “WEIGHT”. In the CCHS, the relevant weight is called “WTS_M”. In the GSS, the relevant variable is called “WGHT_PER”.**

These should be considered when deriving meaningful estimates from the survey.

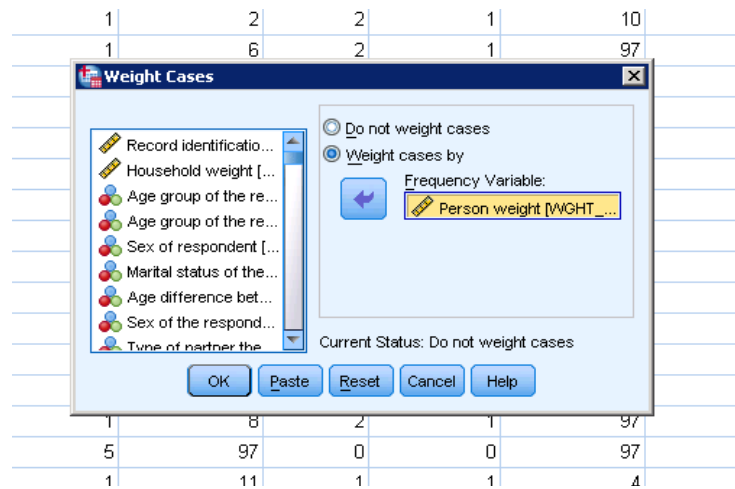
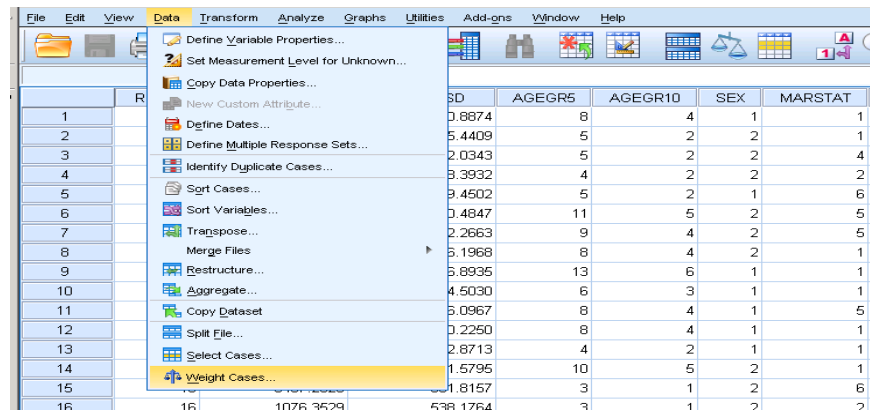
In a probability sample, the sample design itself determines weights which must be used to produce unbiased estimates of the population. Each record (i.e. each case in your sample) must be weighted by the inverse of the probability of selecting the person to whom the record refers. In the example of a 2% simple random sample, this probability would be .02 for each person and the records must be weighted by $1/.02=50$. The importance of working with weights is obvious when you start to examine the complex character of many of Statistics Canada's sample designs. In other words, Statistics Canada works with very complex sampling designs whereby it is clear that not all individuals in the targeted population have an equal likelihood of being selected.

Therefore, in the absence of simple random sampling (whereby every case would have an equal weight), it is necessary to rely upon sample weights in obtaining descriptive statistics that are not biased by the sample design. Note: this is not much of an issue with the Census public use file as it is merely a simple random sample of the full census database and not a complex sample design (i.e. each person, with a few exceptions, represents roughly the same number persons in the full population).

Step 1.

Run a frequency distribution on “Province of Residence” from your dataset (use the dataset that you have selected for your term paper). **Print-up the syntax and output file.** This gives you an “un-weighted” distribution of your sample on province of residence. Use the TITLE procedure to properly document all steps in this assignment. No write up required here.

Step 2. Weight your data and run this frequency distribution again. Using SPSS it is easy to “weight” your dataset. Go to Data>Weight Cases, and then specify the weight appropriate to your dataset (**AWTCW01, WEIGHT, WTS_M or WGHT_PER**).



Paste this procedure into a syntax file, and then below it place the exact same frequency command as specified above (on province of residence).

When you run this syntax file, this should give you the “weighted” frequency distribution on province of residence. **Print up the respective syntax and output files (properly titled). Briefly comment on how it is different relative to the frequency distribution obtained in Step 1 above.** The weighted results are meant to be “more” representative of the Canadian population.

Do you now have a feel for what is meant by weighting your data? You take the sample, which might have a somewhat complex sampling design, and use the sample weights to make our statistics more representative of the larger population. For example, most of Stats Canada’s samples are stratified by province. Typically they “over-sample” the smaller provinces (e.g. PEI, Newfoundland), and “under sample” the larger ones (Quebec or Ontario). This is done for purely statistical reasons, to obtain reasonable estimates for all provinces. It is easy to adjust the estimates with the appropriate sample weights after the fact. Consider the distribution across provinces in the weighted and un-weighted distributions. The “weighted” estimates are obviously more representative than the “non-weighted” estimates (PEI has a population of only about 150,000 persons or less than ½ of a percent of Canada’s total: what do “un-weighted” estimates suggest in terms of percentage living in this province? (I will discuss further in class).

Step 3. Important point: In Step 2 you "weighted the data". If you are asked to work with the "unweighted data" again (as is the case below) you must first "deactivate" the weights (by going back to your weight command, and pasting into your syntax file the "do not use weights" command). Alternatively, if you close and re-open your datafile, your data will also no longer be "weighted".

A problem when running the weight feature on SPSS is that it unfortunately it can lead to *erroneous inferences with regard to statistical significance*. This is particularly the case with **regression procedures**. The reason for this is SPSS continues to treat your data like it were a sample, even though your calculations now involve the weighted numbers (for example, your sample might be only 19000 cases, but the weighted results make reference to tens of millions of Canadians). Remember that if you work with very large samples, then even small differences or small associations may be significant.

A simple way to deal with this problem is as follows: Obtain the “mean” of the sample weight corresponding to the "unweighted" data that you are using in your analysis (**AWTCW01, WEIGHT, WTS_M or WGHT_PER**). You can obtain the mean using the “Descriptives” command in SPSS (go to “Analyze”>“Descriptive Statistics”> “Descriptives” and select your “Weight” variable with this procedure. The output should give you the mean. No need to provide a syntax file or output file with this step (merely record the “mean” on this weight variable for your own records).

Next, with this mean value, compute a new set of weights. The new weights should be equal to the initial weights divided by this mean score. For example, when working with the NLSCY, assume that you found using the Descriptives procedure a mean on AWTCW01 of 200.3342 (note: the mean for the NLSCY weight does not equal 200.3342, but I am merely using this number for demonstrative purposes). This would imply that on average each person in your sample represented about 200 persons. Your revised weights for your analysis could be calculated by using the following command:

Syntax for new weights: COMPUTE NWEIGHT= AWTCW01 / 200.3342

Create a new syntax file with this command, and run your syntax file. After doing so, you have created a new set of weights (let us refer to NWEIGHT in the above example as our new “weight variable”).

With this new weight variable, you must re-weight your data. This is very easy to do. Again, use the Data>Weight cases procedure (selecting this new weight variable that you've just created which should be at the very end of your list of variables) and paste this in your syntax file.

Alternatively, you could alternatively merely place under your COMPUTE COMMAND:

WEIGHT BY NWEIGHT.

When you run this weighting procedure, you have just reweighted your data again. The results should be “unbiased” by sample design but now the TOTAL reported should now be the same size as your initial sample (allowing for meaningful inference tests).

To assure that you have done this properly, again run a “Frequency distribution on the same province of residence variable, yet this time with the “re-weighted” data. Place the frequency command in the same syntax file, yet make sure it is below the aforementioned COMPUTE and WEIGHT commands. Run this procedure and comment.

Note: If you have done this properly, you should have a distribution from Step 3 that has the same 'Total' as your initial “un-weighted sample” in Step 1 yet have the same “percentage distribution” as your weighted result obtained in Step 2.

Print up your syntax and output files from Part 3, properly titled, and include a brief comment on results.

Part B. Target population

If applicable, using SPSS, carefully select the subsample that you will be focusing on in your final paper (example, all children aged 4-11 or all adults aged 18+). As in previous assignments, use the “SELECT CASES” procedure in delineating the sample for your analysis. Again, this relates to the “target population” of your analysis.

If your variables that you select for your research are only applicable to certain subsamples within the GSS, CHHS, NLSCY or Census public use file (e.g. an age appropriate behavioral scale) then you must explicitly take this into consideration when you select your subsample for your multiple regression. Also note that it doesn't make sense to select subsamples if there is no strong theoretical or technical reason to do so. **You need not hand anything in for Step 1, merely report the dataset and your subsample (if applicable) or state that you are working with the full sample.**

PART C. Model Building

From now on, we shall be using “weighted” data in running **ALL REMAINING REGRESSIONS in this assignment.** This will allow for “unbiased results” that take into consideration your sample design while also allowing for more meaningful significance tests.

Although there are various strategies possible when building models (stepwise regression; forward selection procedures, among other largely automated procedures), the best of all possible strategies is let your “theory” guide you in terms of introducing variables into multiple regression. This is what I am recommending for the current assignment and for your final paper, whether you are working with linear or logistic regression. This is also why you are currently conducting and writing up a “literature review”, i.e. you are trying to situate your empirical analysis in terms of a broader research literature. Remember: if your results are non-significant, don’t worry... that’s okay. Merely report it in your final paper.

Working with Weighted Data:

So far in Assignments 1 and 2, you have worked exclusively with “unweighted” data. For the current exercise, you will also conduct several regressions using “weighted” data. **Again, the results from Part C can be used directly in your final paper.**

For all the following steps, your data MUST be properly weighted. If you use the full sample (no subsample) for your analysis, you can merely use the “NEW” weights that you created in Part A above (i.e. make sure that all the syntax files that you use in Part C are weighted by the weight that you created in Part A). If you are selecting a sub-sample for your analysis, it is necessary to quickly repeat what you did in Part A, yet this time exclusively for your selected sub-sample:

- I) select your subsample
- II) obtain the mean of the initial weight (**AWTCW01, WEIGHT, WTS_M or WGHT_PER**). using descriptives.
- III) create a new weight using this mean and the initial weight (using compute as in Part A)
- IV) make sure that all of your syntax files (Steps 1 – 3 below) used from here on in are weighted by this new weight

Step 1-3 below must all be properly documented (TITLE) and weighted (WEIGHT)_
Merely place your appropriate WEIGHT command directly below your TITLE command. All other commands should be placed below these commands.

Step 1.

I had suggested that you come up with 1 (or preferably 2) primary hypothesis(es) for your final paper (i.e. involving 1 or 2, independent variables). Run a separate simple regression with this independent variable, or two separate simple regressions with each independent variable if you have two primary hypotheses.

Note: This all relates to your literature review: which independent variable(s) in your proposed research do you consider most important or most interesting in explaining your dependent variable. What do you want to build your analysis around? What are you trying to highlight in your research paper?

Highlight with a yellow marker on your output: (i) the standardized slope(s), (ii) the unstandardized slope(s), (iii) whether or not your variable(s) have a significant effect or not (at the .05 level), (iv) the y-intercept, and (v) the overall R^2 of each regression.

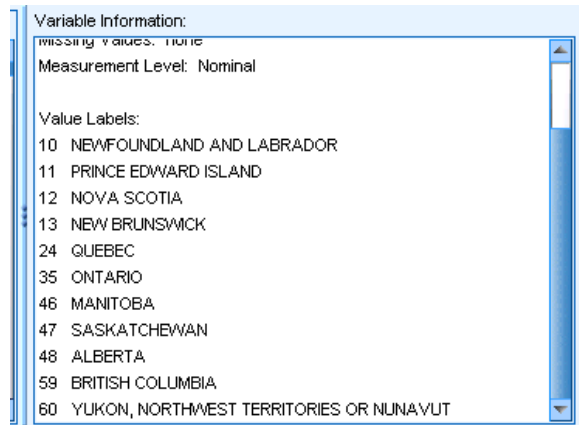
Note:

If your primary hypotheses involve any variable(s) measured at the “nominal” level or a crudely categorized “ordinal variable”, you should create “dummy variables” when working with OLS

linear regression. If you forget how to do this, take another look at assignment 2. Recall the following syntax in creating Dummy variables:

Dummy Variable Syntax

Assume that we have the following variable in our data set with 11 categories (GEOGPRV)



If we were to create a series of dummy variables, we can use the following syntax in our SPSS syntax file (in this example, I will be creating 5 different dummies meant to measure region of country). With some variables it is often worthwhile collapsing categories, so that you don't have too many dummy variables in your final model.

```
RECODE GEOGPRV (10 thru 13=1) (24 thru 60=0) (ELSE=SYSMIS) INTO ATLANTIC.  
VARIABLE LABELS ATLANTIC 'Atlantic Canada'.  
EXECUTE.
```

```
RECODE GEOGPRV (10 thru 13=0) (24=1) (35 thru 60=0) (ELSE=SYSMIS) INTO QUEBEC.  
VARIABLE LABELS QUEBEC 'Prov of Quebec'.  
EXECUTE.
```

```
RECODE GEOGPRV (10 thru 24=0) (35=1) (46 thru 60=0) (ELSE=SYSMIS) INTO ONTARIO.  
VARIABLE LABELS ONTARIO 'Prov of Ontario'.  
EXECUTE.
```

```
RECODE GEOGPRV (10 thru 35=0) (46 thru 48=1) (59 thru 60=0) (ELSE=SYSMIS) INTO  
PRAIRIES.  
VARIABLE LABELS PRAIRIES 'Prairie Provinces'.  
EXECUTE.
```

```
RECODE GEOGPRV (10 thru 48=0) (59 thru 60=1) (ELSE=SYSMIS) INTO BC.  
VARIABLE LABELS BC 'BC and North'.  
EXECUTE.
```

In this example, you have created 5 dummy variables. Note that potentially 11 dummy variables could have been created, but I decided to use only 5 to simplify my analysis. In a subsequent regression, it is necessary that you include only 4 of these dummy variables (the dummy variable that is not included acts as your reference category in regression). For example, if you were to

exclude Quebec, the subsequent slopes in the regression always make implicit reference to this reference category. For example, in reading the slope associated with BC, it would give you the effect of living in BC on your dependent variable relative to living in Quebec (which acts as your reference category).

Usually (but not always) it is recommended to work with “dummy variables” when introducing crudely categorized ordinal variables into regression (i.e. for example, ordinal variables that have only three or four response categories). For the purpose of the current assignment, I recommend that you do so if working with such variables).

The only stuff to hand in with Step 2 are your syntax and output files.

Step 2. Move beyond simple regression, to include all independent variables simultaneously into a multiple regression. In your research proposal, I had suggested that you also consider potential control variables in your multivariate model (three of four). With the results from this regression, highlight with a yellow marker on your output (i) the standardized slopes, (ii) the unstandardized slopes, (iii) whether or not your variables have a significant effect or not (at the .05 level), (iv) Y-intercept, and (v) the overall R^2 of your regression.

If you anticipate the possibility of an interaction effect, it is here that you will have to introduce an interaction term into your model (as will be discussed in class). This usually involves computing a new variable which is the product of the two independent variables hypothesized as interacting. For example, assume that you hypothesize an interaction between VAR1 and VAR2 in explaining your dependent variable VAR3.

You first need compute an interaction term, as:

```
COMPUTE NEWVAR = VAR1 * VAR2.
```

You can then include all three variables in your regression (i.e. VAR1, VAR2 and NEWVAR). If NEWVAR has a significant effect, then this is supportive of the hypothesis of interaction. In including this interaction term, it doesn't matter if VAR1 and VAR2 continue to have a significant effect, but it is important that you continue to include them in the model along with the interaction term (i.e. the final regression must include VAR1, VAR2 and NEWVAR).

Note: You need not “hypothesize” or “test” for interaction effects in this assignment or in your final paper, although it can be interesting. Recall the example provided in class, i.e. the interaction effect of “immigrant status” and “education” in explaining income, i.e. the pay-back for education tends to be much lower for immigrants than it is for other Canadians. Introducing an interaction term which is the product of “immigration status” and “education” could explicitly test for this.

Only stuff to hand in are your syntax and output files.

Step 3

After completing these regressions, I would like you to present the results from the regressions in steps 1-2 in two separate tables (one table which includes standardized slopes and the other table which includes non-standardized slopes). Each table should include the regression results obtained from steps 1 and 2.

In so doing, I suggest that provide me with tables that resemble the format as specified below (Tables 1 and 2). Note: when we present the standardized results, we don't bother including the constant (i.e. the y intercept), but only include the constant in presenting the "unstandardized results".

A nice software for creating these tables is "Excel" (it is very easy to use), or alternatively merely use the Table function available in "Word".

Present the results from Step 1-2 as separate models. In the first model, present the results from the simple regression for one of your independent variables (i.e. unless of course you had to introduce dummy variables in measuring the effect of a given concept whereby you would include all the appropriate dummy variables measuring a given concept). In the second model, include the results from the second simple regression with the other independent variable (if applicable). Regardless of whether you work with one or two simple regressions, you must also also present an additional model (model 3 in our example) with all the variables in your model, including all relevant control variables (this is your full model). The simpler models are called "nested" models, relative to this complete model with all relevant explanatory variables.

More specifically, I suggest that you create one table that includes "unstandardized" slopes (call it Table 1) and another that includes "standardized" slopes (call it Table 2). See below for an example of each:

Table 1. Individual Income Regressed on Selected Independent Variables, 2001 Census, Unstandardized Coefficients.

Covariates	Model 1	Model 2	Model 3
Constant	18033	17500	16003
Generation			
1st generation	-.2322 ***		-.5044 ***
2nd generation	1500 ***		2324 ***
Education		3242 ***	4444 ***
Age			343 *
Region			
Atlantic Canada			-.3245 ***
Quebec			.223
Ontario			3232 ***
Prairies			202
Hrs worked			4323 ***
Sex			22342 ***
N	738333	738333	722444
R ²	0.02	0.033	0.121

P-values ; * P-value < .05; ** P-value < .01 *** P-value < .001

Reference Category: 3+ generation; BC and Female.

The second table that summarizes all the results while using "standardized" slopes is particularly important in determining which of your independent variables has the strongest effect on your dependent variable. (Table 2 does not include the reported intercept (y intercepts) – when

reporting “standardized” slopes). Again, as mentioned above, it is recommended that you present “R²”, a preferred indicator of overall model performance in both Tables.

These tables should also specify the sample sizes of all of your regressions. To locate the sample size of your regression in your corresponding output files, the degrees of freedom (df) as associated with your TSS (total sum of squares) on the ANOVA table gives you a direct indication of the total number of cases involved in each regression. This is not necessarily the same number of cases across the different regressions, due to the issue of missing cases (see note below).

Table 2. Individual Income Regressed on Selected Independent Variables, 2001 Census, Standardized Coefficients.

Covariates	Model 1	Model 2	Model 3
Generation			
1st generation	-0.221 ***		-0.331 ***
2nd generation	0.141 ***		0.145 ***
Education		0.211 ***	0.255 ***
Age			0.002 *
Region			
Atlantic Canada			0.144 ***
Quebec			-0.001
Ontario			0.323 ***
Prairies			0.001
Hrs worked			0.122 ***
Sex			0.366 ***
N	738333	738333	722444
R²	0.02	0.033	0.121
P-values ; * P-value < .05; ** P-value < .01 *** P-value < .001			

Reference Category: 3+ generation, BC and female

As output from Part C all I am asking for is your syntax files and corresponding output files, organized as steps 1 through 3. I am also asking for these two tables, without any interpretation. The results from the “weighted regressions” in Part C can be used in your final paper.

In your final paper, you will have to make sense of these two tables. What do they tell us about the initial relationships as hypothesized. Are they significant? Do you find support for your initial hypotheses? Do they remain, even after including all your relevant control variables??

Note: if you wish, you can depart a bit from the above recommended steps, by including more than 3 models (perhaps 4 alternatives, or even 5 models). You might for example, in addition to what is recommended above, include an additional model with the two independent variables together, prior to presenting all of the variables together in your full model). Alternatively, you might have a separate model that tests for an “interaction” effect. I leave you with some flexibility in developing these “alternative models”.

NOTE: For the current assignment, there is no need for interpretation (merely put together the two tables). THAT'S IT!!! Interpret these tables in your final paper.

A few additional observations:

In our empirical research this year we have yet to explicitly address the problem of “missing values”. In a database, “missing values” can surface as a direct result of a variety of factors, including the possibilities of “don’t know”, “refusal” or “non-applicable” in gathering information through a survey. Typically, non-applicable is not that great an issue as we can avoid this type of missing value by properly defining our subsample for analysis. On the other hand, “refusals” and/or “don’t know” can potentially be problems.

The default procedure in regression when working with SPSS is to merely delete “missing cases” on a “LISTWISE” basis. Listwise deletion uses only cases from your data set that have valid (non-missing) values for all variables included in the regression. While this is the default procedure in regression, it is not always the preferred option. The problem with “Listwise” deletion of missing cases is that you can potentially lose quite a few cases (and this can seriously reduce your sample size). If you find that you are losing a lot of cases, you should be clear on why this is occurring. Can it be avoided? If not, do note it when reporting your final results in your final paper.

NOTE: If you are working with the CENSUS, missing values is not an issue (i.e. there are not any in this dataset, as Stats Canada imputes all missing information). The three regressions should be identical when working with the Census, unless you have made some sort of serious mistake in terms of “not applicable”.

Also, for future reference:

It is useful to examine initially your data with un-weighted counts. This provides you with some sense as to potential problems with small numbers and missing cases.

Yet if you require descriptive estimates on how the variables likely look in the population (example, frequency distributions), then run all your procedures with the correct sample weights. If your sample design is not simple random sample, then these results can be quite different from what was initially observed. On the other hand, if your survey is based on a “simple random sample”, whereby all persons in the sample had an equal probability of selection, then your results from the “un-weighted” procedure will likely be near identical to what you obtain with the above procedure (note: the 2006 public use file is a “simple random sample” of the complete census).

One final observation:

When you move into your final multivariate analysis, the only way to get meaningful significance tests is to revise the weight as specified above (this is not perfect, but it is certainly better than doing nothing). In using this method, the subsequent estimates should not be seriously biased and the significance tests should be reasonably accurate. This is a preferred option over not “weighting” your analysis or merely doing the analysis with the initial weights. There are more complex procedures possible, beyond the scope of the current course, that are preferred to the recommended procedure as specified above (save it for graduate school).