

Exercise non-probability sampling

Peter Lugtig

29-11-2020

The goal of this exercise is to practice with calibration, raking or imputation to correct for selection bias in non-probability-based surveys. In principle the methods discussed today can be applied to any data source that one knows is suffering from selection bias. E.g. experiments, social media data, or other sources of Big Data. However you should always be aware that your non-probability dataset should at least exhibit a similar type of variation as you have in the population, and that all relevant subgroups in the population are also present in your sample (what Mercer et al, call positivity). To give an example of a problematic situation regarding positivity, it will be impossible to adjust a typical psychological experiment conducted among 18-25 year old students at a Western University to the general population, because you just have young and highly educated people.

The dataset(s) you will use today are slightly adapted from data that are publicly available through PEW, which is a non-profit survey data collection organisation in the USA (www.pewresearch.org) Between June and July 2016, PEW designed a short questionnaire and then asked three organizations that rely on volunteer opt-in survey panels to administer this questionnaire in their panel. This resulted in a dataset of about 30.000 respondents. You are today getting about two-thirds of these respondents to develop an adjustment method. The remaining 10.000 respondents will serve as the hold-out sample against which I will test your adjustment method. I also excluded many variables: your dataset will consist of about 30 variables which you can use for adjustment, and 1 dependent variable called VOTESUM.

The variable VOTESUM asks for the Future voting behavior in the November 2016 Presidential Election, with a choice between Clinton, Trump and being undecided. At 1 July 2016, the aggregated polls indicated that Clinton would receive 46% of the vote, and Trump 42%, with the rest being undecided. If I correct for the bias that was present in the polls throughout the 2016 election cycle (as assessed in the Kennedy et al report that you read in week 1 of this course), my best guess for the true difference between the candidates would have been 45% for Clinton and 43% for Trump.

If you want to know more about the study and data, have a look at: <https://www.pewresearch.org/methods/2018/01/26/for-weighting-online-opt-in-samples-what-matters-most/>

Lets first load the data.

```
library(foreign)
library(survey)
library(dplyr)
library(plyr)
library(mice)

nonprob <- readRDS("PEW_nonprob_samples.RDS")
# below you see how many cases there are from each vendor
table(nonprob$vendor)
```

Variables

The covariates in this dataset consist of demographic variables but also a range of other variables for which the people at Pew hoped they would correlate to both selection bias (R) and the dependent variable (Y): here VOTESUM. Pew employed a superpopulation approach to assemble additional population level characteristics (more on this later), which were then also included in the survey as questions. Here is the list of variables:

GENDER # make, female
AGE # age in years EDUCCAT5 # educational level in 5 categories (lo to hi)
DIVISION # region in 4 categories
MARITAL_ACS # 5 cat marital status
HHSIZECAT #1,2 3+
CHILDRENCAT # 1,2,3_ children at home
CITIZEN_REC # US citizen or not
BORN_ACS # born inside, or outside USA
FAMINC5 #income in5 bracks: <20k, 20-40,40-75,75-150, >150
EMPLOYED # in employment or not
MIL_ACS_REC # never been on active military, has been in military
HOME_ACS_REC # 1. own, 2. rent 3. rent without pay
FDSTMP_CPS # 1.Do you receive foodstamps?
TENURE_ACS # did you live on current address one year ago
PUB_OFF_CPS # have you visited a public official to express your opinion in the past 12 months?
COMGRP_CPS # have you participated in a school, neighborhood or community association in the past year?
TALK_CPS # talk to family 5 cat
TRUST_CPS # do you trust the people in your neighborhood (5 cat)
TABLET_CPS # do you use a tablet?
TEXTIM_CPS # do you every send text messages?
SOCIAL_CPS # active on social media
VOLSUM # volunteering
REGISTERED # registered to vote: yes, no
VOTE14 # voted in 2014 miterm election
PARTYSALE5 #party attachment 5 categories
RELIGCAT #5 cat religious affiliation: roman catholic, evangelical, main protestant, other, unaffiliated
IDEO3 # liberal, moderate, conservative
OWNGUN_GSS # owns a gun (yes, no)
FOLGOV # do you follow the government

(super)population data

Pew collected data from the Census bureau, as well as some other Gold-standard administrative sources, and surveys for the eligible voting USA population. They also used some of the surveys to compute (aggregated) correlations between these variables, and used this to create a SuperPopulation dataset of the USA. The dataset includes only 20,000 people, but it is supposed to be representative for the USA on the variables included in this dataset. The variable names, labels, and codings in this dataset are the exact same as in the non-probability based survey dataset. This is neat, no recoding or data handling! Load the data below.

```
popdata <- readRDS("PEW_population_data.RDS")
```

What to do?

You have 30 variables to now build a model from. Please note the following: - It is nice if the covariates predict Y (voting behavior). There is no need to read the literature on voting behavior, but do think about this when you select your adjustment variables. - If you want to include 30 variables in one go into your model, perhaps with some interactions, you will encounter some issues: First of all, R may become quite slow! Second, there may be estimation issues, because these are perhaps just many variables. Maybe it is worthwhile to first try a model with just a few variables, like in the example below:

How did raking and calibration work again?

```
# poststratification example
# first, specify the unweighted svydesign object
svy.unweighted <-svydesign(ids=~rid, data=nonprob)
```

```
## Warning in svydesign.default(ids = ~rid, data = nonprob): No weights or
## probabilities supplied, assuming equal probability
```

```
# below, I use gender and whether someone voted in the 2014 midterm election
#gendervote<- as.data.frame(table(nonprob2$GENDER,nonprob2$VOTE14))
gendervote.dist <-xtabs(~GENDER+VOTE14, data=popdata) # the 2x2 table for the population
poststratdesign <- postStratify(design= svy.unweighted, strata =~GENDER+VOTE14,
                             population=gendervote.dist)
```

```
# for Raking (easier when you have many variables in your model, but be careful with continuous vars)
gender.dist <- as.data.frame(table(popdata$GENDER))
colnames(gender.dist) <- c("GENDER", "Freq")
vote.dist <- as.data.frame(table(popdata$VOTE14))
colnames(vote.dist) <- c("VOTE14", "Freq")
rakeddesign <- rake(design=svy.unweighted,sample.margins=list(~GENDER,~VOTE14),
                 population.margins=list(gender.dist,vote.dist))
```

```
# add the weights if you want to do diagnostics
nonprob$PSweights <- poststratdesign$prob
nonprob$rakeweights <- rakeddesign$prob
```

```
##### and then estimate the outcome
svyby(~VOTESUM,~vendor, design=svy.unweighted,svymean)[2] # p(Trump) across vendors.
```

```
##           VOTESUMTrump
## Vendor 1    0.4335610
## Vendor 2    0.4321534
## Vendor 3    0.4090466
```

```
svyby(~VOTESUM,~vendor, design=poststratdesign,svymean)[2]
```

```
##           VOTESUMTrump
## Vendor 1    0.4384632
## Vendor 2    0.4317756
## Vendor 3    0.4160576
```

```
svyby(~VOTESUM,~vendor, design=rakeddesign,svymean)[2]
```

```
##           VOTESUMTrump
## Vendor 1    0.4373653
## Vendor 2    0.4310338
## Vendor 3    0.4153107
```

Imputation models

Are you curious to try Mass Imputation? Great! I made your life a bit easier by creating an artificial population of the USA, and then match the survey data to this population dataset. Please note that the population dataset has about 1 million cases, which is only 1/200 of the actual population size. R gets slow however with 1 million cases, so just imagine you have the actual population

```
wholepop <- readRDS("complete_USA_population_data_with_matched_survey_results.RDS")
```

Want to see how you design an imputation model? Have a look at the weeks where we talked about imputation.

Ok now what?

The idea is that you now develop your own adjustment model. You may try out different models (raking, poststratification, mass imputation), and can of course also try different covariates. Please send your R-code of your final model (just 1!) in a plan-text e-mail to Peter before the next class. Please do not change the names of the variables so your code will actually run.

The person who has designed the best adjustment model (closest to the poll aggregate on July 1st), which is 45% Clinton/ 43% Trump will win a prize. Note I will run your model on my holdout sample to prevent overfitting.

Good luck, and have fun!