Workshop on Causality 2019, Uder, Germany
Mediation workshop

Introduction

This document serves to illustrate some of the available mediation approaches that you can implement in the R software language. In particular, we will focus on three packages: (1) `sem`, (2) `mediation`, and (3) `medflex`.

The dataset we will be using for our analyses is the JOBS II dataset, which is available as a download from

https://doi.org/10.3886/ICPSR02739.v1

or as part of the `mediation` package. For convenience, we will use the version of the dataset that comes with the `mediation` package. This version of the dataset contains $N = 899$ complete observations on 17 demographic, socioeconomic, and health-related variables. The JOBS II study was a randomized experiment designed to measure the efficacy of a job training intervention on unemployed workers, not only in terms of reemployment but also improved mental health. After completing a baseline questionnaire, participants were randomly assigned to a treatment ($X = 1$) or control ($X = 0$) condition. The treatment condition consisted of participating in a workshop that taught job seeking skills and strategies to deal with disappointment when searching for a job, while the control condition consisted of receiving a booklet that provided tips for finding a job. The two outcomes of interest were (1) a continuous measure of depressive symptoms and (2) a binary measure of whether the participant had found a job by the end of the study or not.

After the `mediation` package has been loaded, you can get the JOBS II dataset via

```r
data(jobs)
```

If you want to check the structure of the JOBS II data, or have questions about what variables are, you can access these through

```r
#Check the structure of the data
str(jobs)

#Bring up the help file for the data
?jobs
```

A good practice for replicable data analysis is to set a random seed.

```r
set.seed(1234)
```

For each approach, we will carry out 2 mediation analyses corresponding to the 2 outcomes that the JOBS II intervention was hypothesized to affect, namely participants' mental health and reemployment. For each, the treatment will be the JOBS II intervention, which is the variable `treat` in the JOBS II data. We will also use the same mediator, namely job search self-efficacy (i.e., how strongly a participant believed that he or she is capable of looking for a job). The variable is called `job_seek` in the dataset, and it is a continous measurement in which higher scores indicate more self-efficacy. For the first analysis, our outcome will be a continuous measure of depressive symptoms, i.e., `depress2`, in which higher scores indicate stronger depressive symptoms. For the second analysis, our outcome will be a binary indicator of whether the participant found a job at the end of the study, i.e., `work1`. To address the possibility of confounding, we will adjust for three covariates: (1) age, (2) sex, and (3) econonic hardship.

Exercise 1: Based on the information provided, please draw the DAGs corresponding to the first mediation question, namely whether the JOBS II intervention impacts depressive symptoms through job search self-efficacy.

Exercise 2: Suppose a subject matter expert tells you that the JOBS II intervention was very time intensive, and a number of participants did not adhere to their assigned treatment. How can you reflect this in the DAG you drew? (Note: We will assume the DAG in Exercise 1 for all analyses).

Mediation analysis with `sem`

The linear structural equation modeling (LSEM) approach to mediation analysis involves specifying linear structural models for both the mediator and outcome, and then estimating mediation effects as functions of the resulting regression parameters. All of this can be done using the `sem` package.

Exercise 3: If we are using LSEMs to answer a mediation question, what assumptions must we make? Do you think these assumptions are plausible?

The first mediation question is whether the JOBS II intervention affects depressive symptoms, possibly through job search self-efficacy. The LSEMs for this mediation analysis are:

$$M = \beta_0 + \beta_1 X + \beta_2 C_1 + \beta_3 C_2 + \beta_4 C_3 + \epsilon_m$$
$$Y = \theta_0 + \theta_1 X + \theta_2 M + \theta_3 C_1 + \theta_4 C_2 + \theta_5 C_3 + \epsilon_y$$

where the variables Y, M, A, C_1, C_2, and C_3 are post-baseline depressive symptoms, job search self-efficacy, treatment, age, sex, and economic hardship, respectively.

To specify the LSEMS, the `specifyEquations()` function can be used. This function requires that the system of LSEMs is input as text, and it translates it into an `R` object that can be used by other functions in the package:

```
jobs2sem <- specifyEquations(text="
  job_seek = beta1*treat + beta2*age + beta3*sex + beta4*econ_hard
  depress2 = theta1*treat + theta2*job_seek + theta3*age + theta4*sex + theta5*econ_hard")
```

`R` adds two variances to the model. These are the variances of the endogenous variables (`job_seek` and `depress2`). To check the model specification, enter `jobs2sem` into the console.

The `sem()` function estimates the structural parameters. This function requires a model specification (i.e., the `jobs2sem` object we created above), plus a dataset[1] and an argument (`fixed.x`) that tells the function which variables are exogenous.

```
jobs2sem_fit = sem(jobs2sem,
                   data=jobs,
                   fixed.x = c("treat", "econ_hard", "sex", "age"))

summary(jobs2sem_fit)
```

```
##
##  Model Chisquare =  1.246225e-13   Df =  0 Pr(>Chisq) = NA
##  AIC =  22
##  BIC =  1.246225e-13
##
```

[1] You could alternatively specify an observed data covariance matrix and the number of observations in the dataset

```
##  Normalized Residuals
##       Min.    1st Qu.     Median       Mean    3rd Qu.       Max.
## -1.698e-15 -1.690e-16  0.000e+00 -1.555e-16  0.000e+00  6.377e-16
##
##  R-square for Endogenous Variables
## job_seek depress2
##   0.0115   0.1203
##
##  Parameter Estimates
##             Estimate      Std Error   z value    Pr(>|z|)
## beta1        0.0656150034 0.051356976  1.2776259 2.013814e-01
## beta2        0.0045864917 0.002313725  1.9822977 4.744593e-02
## beta3       -0.0076373362 0.048615938 -0.1570953 8.751697e-01
## beta4        0.0531624129 0.024543589  2.1660407 3.030807e-02
## theta1      -0.0402647000 0.043384588 -0.9280877 3.533621e-01
## theta2      -0.2399549527 0.028164594 -8.5197376 1.599153e-17
## theta3       0.0006488642 0.001957048  0.3315525 7.402272e-01
## theta4       0.1068048699 0.041032342  2.6029435 9.242717e-03
## theta5       0.1485433818 0.020768795  7.1522388 8.537385e-13
## V[job_seek]  0.5245259758 0.024753911 21.1896201 1.190539e-99
## V[depress2]  0.3736373989 0.017633039 21.1896201 1.190539e-99
##
## beta1       job_seek <--- treat
## beta2       job_seek <--- age
## beta3       job_seek <--- sex
## beta4       job_seek <--- econ_hard
## theta1      depress2 <--- treat
## theta2      depress2 <--- job_seek
## theta3      depress2 <--- age
## theta4      depress2 <--- sex
## theta5      depress2 <--- econ_hard
## V[job_seek] job_seek <--> job_seek
## V[depress2] depress2 <--> depress2
##
##  Iterations =  0
```

Exercise 4: Which parameter above is the direct effect? How can we interpret this effect?

Since the `sem` package is not a mediation-specific package, a small amount of work must be done to get an estimate of the indirect effect. By the product method, the indirect effect is equal to $\hat{\beta}_1\hat{\theta}_2$, which is approximately -0.016.

Calculating the standard error for the indirect effect is not as straightforward, and it requires the use of the delta method[2]. Using this method, the estimated standard error is 0.012. You could then use this standard error estimate to test the significance of an indirect effect, or you could use another method, like bootstrapping[3].

A slightly easier way of doing the same analysis using LSEMs is via the the `lavaan` package. The syntax is

[2] see Sobel, Michael E. (1986). "Some new results on indirect effects and their standard errors in covariance structure". *Sociological Methodology.* 16: 159-186.

[3] Those who are interested in learning more about non-parametric bootstrapping in R are encouraged to read the useful description of the `boot` package at https://www.statmethods.net/advstats/bootstrapping.html

```r
#Code is not run to save space.
#Make sure you have lavaan installed before running

lavaan_model <- ' # direct effect
                    depress2 ~ direct*treat + econ_hard + age + sex
                  # mediator
                    job_seek ~ beta*treat + econ_hard + age + sex
                    depress2 ~ theta*job_seek
                  # indirect effect (beta*theta)
                    indirect := beta*theta
                  # total effect
                    total := direct + (beta*theta)
                '
lavaan_fit <- lavaan::sem(lavaan_model, data = jobs)
summary(lavaan_fit)
```

Because `lavaan` allows for mediation effect to be defined as part of the model specification, the direct, indirect, and total effect estimates (and their standard errors) are given in the output.

Since the `sem` package is for linear structural equation modeling, we do not recommend it be used to estimate mediation effects for binary outcomes. Therefore, we do not attempt to use it for the second mediation case in which the outcome is whether the JOBS II participant was reemployed vs. not.

Mediation analysis with `mediation`

The `mediation` package implements the parametric causal mediation analysis techniques described in Imai, Keele, and Tingley (2010)[4]. In their framework, one can estimate two mediation effects, the average causal mediation effect (ACME) and average direct effect (ADE). These are defined as follows for a binary treatment:

$$\text{ACME} = \mathbb{E}\left[Y\left(x, M(1)\right)\right] - \mathbb{E}\left[Y\left(x, M(0)\right)\right]$$
$$\text{ADE} = \mathbb{E}\left[Y(1, M(x))\right] - \mathbb{E}\left[Y(0, M(x))\right]$$

for $x = 0, 1$.

Please note that these effects are the same as you saw in the lecture, but with different names. The ACME is the natural indirect effect (NIE), while the ADE is the natural direct effect (NDE).

Exercise 5: What are the assumptions that now underlie this mediation analysis? Are they plausible?

We start with the same analysis above. In the `mediation` package, the models are specified as

```r
imai_m <- lm(job_seek ~ treat + econ_hard + sex + age, data=jobs)
imai_y <- lm(depress2 ~ treat + job_seek + econ_hard + sex + age, data=jobs)
```

The `mediate()` function takes these equations and estimates mediation effects, and the variables corresponding to the treatment and the mediator must be given. There are also several arguments related to confidence interval calculations that you can change, including whether the non-parametric bootstrap or a quasi-Bayesian approximation (the default) is used.

[4]Imai, K., Keele, L. and Tingley, D., 2010. A general approach to causal mediation analysis. *Psychological methods*, 15(4), p.309-334

```r
imai_mediation <- mediate(model.m = imai_m,
                          model.y = imai_y,
                          treat="treat",
                          mediator="job_seek",
                          #use non-param bootstrap
                          boot = TRUE)
summary(imai_mediation)
```

```
##
## Causal Mediation Analysis
##
## Nonparametric Bootstrap Confidence Intervals with the Percentile Method
##
##                  Estimate 95% CI Lower 95% CI Upper p-value
## ACME              -0.0157      -0.0423         0.01    0.18
## ADE               -0.0403      -0.1230         0.05    0.40
## Total Effect      -0.0560      -0.1404         0.03    0.26
## Prop. Mediated     0.2811      -1.8890         2.57    0.36
##
## Sample Size Used: 899
##
##
## Simulations: 1000
```

Exercise 6: What are the estimates of the NDE and NIE? How would you interpret them?

Exercise 7: Add a treatment-mediator interaction to the outcome model, and then reestimate the NDE and NIE. What do you notice about the output of the package now? How can you interpret these effects?

The `mediation` package also works with nonlinear models. The effect of the JOBS II intervention on reemployment, possibly through job search self-efficacy, can also be studied. The syntax for this analysis will be very similar to what was done above, with a logistic regression replacing a linear regression for the outcome.

```r
imai_y2 <- glm(work1 ~ treat + job_seek + econ_hard + sex + age,
               data=jobs,
               family=binomial(link="logit"))

imai_mediation2 <- mediate(model.m = imai_m,
                           model.y = imai_y2,
                           treat="treat",
                           mediator="job_seek")

summary(imai_mediation2)
```

```
##
## Causal Mediation Analysis
##
## Quasi-Bayesian Confidence Intervals
##
##                           Estimate 95% CI Lower 95% CI Upper p-value
## ACME (control)             0.00261     -0.00188         0.01   0.258
## ACME (treated)             0.00283     -0.00202         0.01   0.258
```

```
## ADE (control)              0.05256     -0.00999         0.11   0.092 .
## ADE (treated)              0.05278     -0.00999         0.11   0.092 .
## Total Effect               0.05539     -0.00821         0.11   0.078 .
## Prop. Mediated (control)  0.03401     -0.15591         0.42   0.328
## Prop. Mediated (treated)  0.03807     -0.15267         0.42   0.328
## ACME (average)             0.00272     -0.00197         0.01   0.258
## ADE (average)              0.05267     -0.00999         0.11   0.092 .
## Prop. Mediated (average)  0.03604     -0.15435         0.42   0.328
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Sample Size Used: 899
##
##
## Simulations: 1000
```

Here, ACME and ADE effects given for the treatment and control groups separately, even though there is no interaction terms in the model. This is because we used a non-linear model for the outcome, so by default the `mediate()` function gives separate effects for treated and control units (but they are quite similar, given the absence of any specified interactions). From this output, the ACME is approximately 0.003 and the ADE is approximately 0.053. Both of these effects can be interpreted as absolute risk differences between the treatment and control conditions, since the `mediation` package only gives natural effect estimates on an additive scale.

Mediation analysis with `medflex`

The `medflex` package implements the natural effects models of Lange, Vansteelandt, and Bekaert (2012)[5] and Vansteelandt, Bekaert, and Lange (2012)[6]. In essence, the natural effects model approach to mediation analysis consists of two steps: (1) expanding the data and (2) estimating mediation effects. To complete the first step, an analyst could either choose a weighting approach in which s/he specifies a model for the mediator given the treatment and covariates, or s/he can choose an imputation approach in which s/he specifies instead a model for the outcome given the treatment, mediator, and covariates.

This tutorial will focus on the imputation-based approach. Consider the second mediation question of whether the JOBS II intervention affects reemployment, possibly through job search self-efficacy. The imputation model for this question will be a logistic regression

$$\text{logit}\Pr(Y=1|X,M,C)=\beta_0+\beta_1 x+\beta_2 m+\beta_3 c_1+\beta_4 c_2+\beta_5 c_3+\beta_6 xm+\beta_7 xc_2$$

This model includes two interactions, one between the treatment and the mediator and another between the treatment and sex. The parameter β_1 will encode the pure natural direct effect, while the parameter β_2 will encode the pure natural indirect effect.

```
#Note: 1st variable MUST be trt and 2nd MUST be mediator
#as.factor() used to ensure trt is understood to be binary

medflex_impmod <- glm(work1 ~ as.factor(treat) + job_seek + econ_hard + sex + age
                      + as.factor(treat)*job_seek + as.factor(treat)*sex,
                      family=binomial(link='logit'),
                      data=jobs)
```

[5] Lange T, Vansteelandt S, Bekaert M (2012). A simple unified approach for estimating natural direct and indirect Effects. *American Journal of Epidemiology*, 176(3), 190-195.

[6] Vansteelandt, S., Bekaert, M. and Lange, T., 2012. Imputation strategies for the estimation of natural direct and indirect effects. *Epidemiologic Methods*, 1(1), pp.131-158.

After specifying this imputation model, it can be inputted it into the `neImpute()` function of `medflex`. This function takes as inpute a generalized linear model, and it outputs an expanded dataset in which different counterfactual outcomes are imputed according to the outcome model specified.

```
medflex_impute <- neImpute(medflex_impmod)
```

Exercise 8: Print the first 6 rows of the imputation dataset you just created, and compare it to the first 6 rows of the JOBS II dataset. Comment on their differences and similiarities. Are they what you would have expected?

The estimation of the natural effects model is done through the `neModel()` function. This function requires the specification of the natural effects model, including which GLM should be used to estimate it. It also requires an expanded dataset (e.g., `medflex_impute` above). If desired, default arguments related to boostrapping and confidence intervals can be changed. The syntax is as follows:

```
medflex_neModel <- neModel(work1 ~ treat0*treat1 + treat1*sex + age + sex + econ_hard,
                           family=binomial(link='logit'),
                           expData = medflex_impute,
                           #Change the default number of bootstrap resamples to 100
                           nBoot = 100,
                           #Change the default print behavior of the function to FALSE
                           progress = FALSE)
summary(medflex_neModel)
```

```
## Natural effect model
## with standard errors based on the non-parametric bootstrap
## ---
## Exposure: treat
## Mediator(s): job_seek
## ---
## Parameter estimates:
##                 Estimate Std. Error z value Pr(>|z|)
## (Intercept)      0.702448   0.356395   1.971  0.04873 *
## treat01          0.268623   0.151422   1.774  0.07606 .
## treat11          0.042991   0.034526   1.245  0.21307
## sex             -0.479240   0.163450  -2.932  0.00337 **
## age             -0.033820   0.007798  -4.337 1.45e-05 ***
## econ_hard       -0.037191   0.073646  -0.505  0.61356
## treat01:treat11 -0.053901   0.036988  -1.457  0.14505
## treat11:sex      0.009691   0.026816   0.361  0.71781
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Thus, the pure natural direct log odds ratio is approximately 0.269, while the pure natural indirect effect log odds ratio is approximately 0.043. Since there is a treatment-mediator interaction and modification by sex, these effect estimates do not give the complete story. A more informative summary is given by the `neEffdecomp()` function

```
#sex = 0 for males and = 1 for females
decomp_males <- neEffdecomp(medflex_neModel, covLev = c(sex=0))
summary(decomp_males)
```

```
## Effect decomposition on the scale of the linear predictor
```

```
## with standard errors based on the non-parametric bootstrap
## ---
## conditional on: sex = 0, age, econ_hard
## with x* = 0, x = 1
## ---
##                        Estimate Std. Error z value Pr(>|z|)
## pure direct effect      0.26862    0.15142   1.774   0.0761 .
## total direct effect     0.21472    0.15353   1.399   0.1619
## pure indirect effect    0.04299    0.03453   1.245   0.2131
## total indirect effect -0.01091    0.02503  -0.436   0.6629
## total effect            0.25771    0.15365   1.677   0.0935 .
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
## (Univariate p-values reported)
```

```r
decomp_females <- neEffdecomp(medflex_neModel, covLev = c(sex=1))
summary(decomp_females)
```

```
## Effect decomposition on the scale of the linear predictor
## with standard errors based on the non-parametric bootstrap
## ---
## conditional on: sex = 1, age, econ_hard
## with x* = 0, x = 1
## ---
##                         Estimate Std. Error z value Pr(>|z|)
## pure direct effect      0.268623   0.151422   1.774   0.0761 .
## total direct effect     0.214722   0.153528   1.399   0.1619
## pure indirect effect    0.052682   0.033193   1.587   0.1125
## total indirect effect -0.001219   0.026306  -0.046   0.9630
## total effect            0.267404   0.150574   1.776   0.0758 .
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
## (Univariate p-values reported)
```

Exercise 9: Use the medflex package to investigate the first mediation question of whether the JOBS II intervention affects depressive symptoms, possibly through job search self-efficacy. How do the effect estimates you obtain compare to those obtained by the mediation package?

Glossary of causal mediation

Pure direct effect: The contrast $\mathbb{E}\{Y(1, M(0)) - Y(0, M(0))\}$

Total direct effect: The contrast $\mathbb{E}\{Y(1, M(1)) - Y(0, M(1))\}$

Pure indirect effect: The contrast $\mathbb{E}\{Y(0, M(1)) - Y(0, M(0))\}$

Total indirect effect: The contrast $\mathbb{E}\{Y(1, M(1)) - Y(1, M(0))\}$

Average causal mediation effect (ACME): Term used by Imai et al. to describe the natural indirect effect. For a binary treatment, it is defined as $\mathbb{E}\{Y(x, M(1)) - Y(x, M(0))\}$ for $x = \{0, 1\}$. When $x = 0$, the ACME is the pure indirect effect. When $x = 1$, the ACME is the total indirect effect.

Average direct effect (ADE): Term used by Imai et al. to describe the natural direct effect. For a binary treatment, it is defined as $\mathbb{E}\{Y(1, M(a)) - Y(0, M(a))\}$ for $x = \{0, 1\}$. When $x = 0$, the ADE is the pure direct effect. When $x = 1$, the ADE is the total direct effect.