

R Lab 1 - Defining the Causal Parameter & Introduction to Simulations in R

Introduction to Causal Inference

Goals:

1. Review the structural causal model (SCM), counterfactuals, and causal parameters, defined with and without a working marginal structural model (MSM).
2. Given a specific data generating process, obtain the value of the target causal parameter $\Psi^{\mathcal{F}}(\mathbb{P}_{U,X})$ in closed form.
3. Introduce simulations in R. Translate a specific data generating process, which is an element of the causal model, into an R simulation.
4. Simulate the counterfactual outcomes Y_a for levels of the exposure of interest, and obtain the value of target causal parameter using simulations.

Next lab:

We will obtain the value of the statistical estimand, which under the needed identifiability assumptions equals the target causal parameter. We will also implement the simple substitution estimator, based on the G-computation identifiability result, and use simulations to evaluate the properties of estimators.

Reminder:

This is not an R class. However, software is an important bridge between the statistical concepts and implementation.

1 Background Story

Suppose we are interested in the causal effect of butterbeer consumption on happiness among wizards at Hogwarts. Specifically, we want to know if the average happiness would be higher if all wizards consumed butterbeer or if all wizards did not. Let $W1$ be a summary measure the student's pre-exposure covariates, including age, house, gender, friendship with Dumbledore and enemy status with Snape. Let $W2$ be an additional baseline covariate, indicating whether the student had permission to travel to Hogsmeade, a location where butterbeer is sold. We consider a binary exposure A , indicating consumption of butterbeer ($A = 1$) or not ($A = 0$). Denote the outcome happiness with Y . Finally, suppose having a permission $W2$ only affects the exposure A , but has no direct effect on the happiness Y .



<http://www.funcage.com/blog/10-incredibly-nerdy-dog-costumes/>

This study can be translated into the following, structural causal model (SCM) $\mathcal{M}^{\mathcal{F}}$:

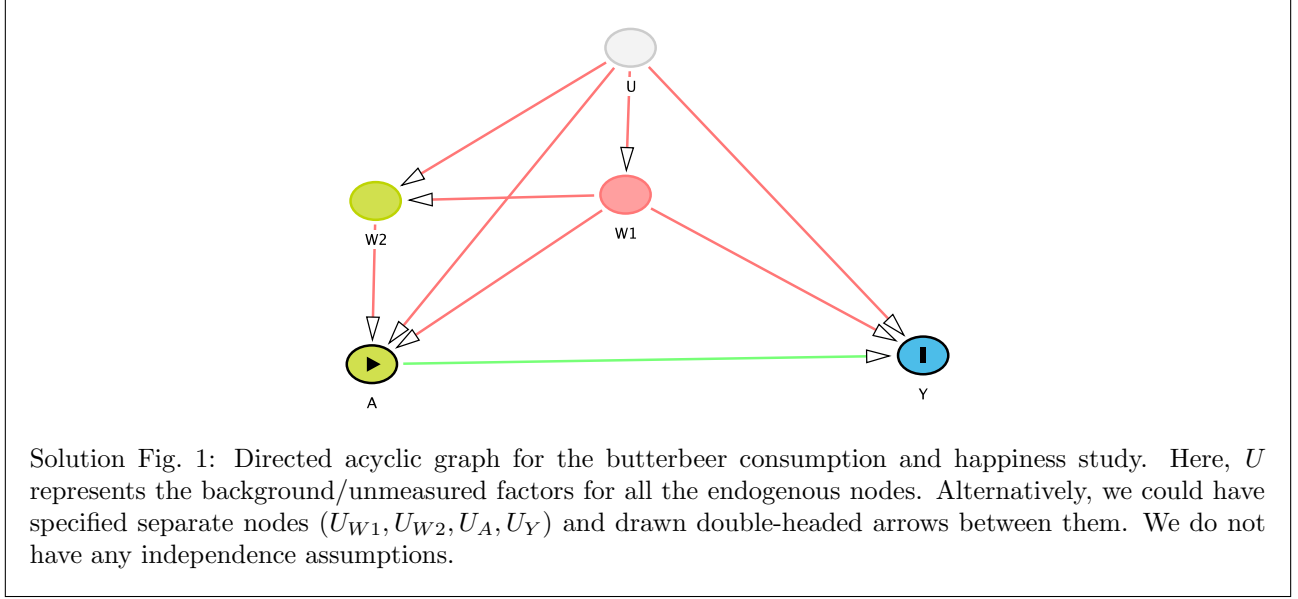
Endogenous Nodes: $X = (W1, W2, A, Y)$
 Exogenous (Unmeasured) Nodes: $U = (U_{W1}, U_{W2}, U_A, U_Y) \sim \mathbb{P}_U$
 Structural Equations F :

$$\begin{aligned} W1 &= f_{W1}(U_{W1}) \\ W2 &= f_{W2}(W1, U_{W2}) \\ A &= f_A(W1, W2, U_A) \\ Y &= f_Y(W1, A, U_Y) \end{aligned}$$

1. Draw the corresponding directed acyclic graph (DAG).
2. Are there any exclusion restrictions? Are there any independence assumptions?
3. Define the counterfactual outcomes of interest with formal notation and in words.
4. In the SCM framework, how are counterfactuals derived? What does the SCM tell us about all possible distributions for these counterfactuals?

Solution:

1. This study can be translated into the directed acyclic graph, given in Fig. 1.
2. We are making an exclusion restriction; the baseline covariate $W2$ does not directly affect the outcome Y . We have not made any independence assumptions. In other words, there are no restrictions on the joint distribution of the unmeasured factors \mathbb{P}_U .
3. The counterfactuals of interest are $(Y_a : a \in \mathcal{A} = \{0, 1\})$. Here, Y_1 is the counterfactual happiness if, possibly contrary to fact, a wizard drank butterbeer and Y_0 is the counterfactual happiness if, possibly contrary to fact, a wizard did not drink butterbeer.
4. Counterfactuals are derived by intervening on SCM to set $A = a$. The distribution of counterfactuals is implied by the joint distribution of (U, X) , which is denoted $\mathbb{P}_{U,X}$. The SCM $\mathcal{M}^{\mathcal{F}}$ defines for the set of allowed counterfactual distributions. (This is why it is called a “model”.)



2 Target causal parameters defined without a MSM

Suppose our target causal parameter is the average treatment effect:

$$\begin{aligned}\Psi^{\mathcal{F}}(\mathbb{P}_{U,X}) &= \mathbb{E}_{U,X}(Y_1) - \mathbb{E}_{U,X}(Y_0) \\ &= \mathbb{E}_{U,X}[f_Y(W1, 1, U_Y)] - \mathbb{E}_{U,X}[f_Y(W1, 0, U_Y)]\end{aligned}$$

This is the difference in the expected counterfactual happiness if all wizards were to drink butterbeer and the expected counterfactual happiness if all wizards were not to drink butterbeer. In the second line, we have replaced the counterfactual outcome Y_a with the corresponding structural equation f_Y .

In the previous section, we specified a SCM $\mathcal{M}^{\mathcal{F}}$, reflecting our limited knowledge of the data generating system. We did not place any assumptions on the joint distribution of the exogenous nodes. We made only one exclusion restriction. Finally, we did not make any assumptions about the functional form of the structural equations.

Now, we consider a particular data generating process $\mathbb{P}_{U,X}$, **one** of many compatible with $\mathcal{M}^{\mathcal{F}}$.

- Each of the exogenous factors U is drawn independently from the following distributions:

$$U_{W1} \sim \text{Uniform}(\min = 0, \max = 1)$$

$$U_{W2} \sim \text{Bernoulli}(p = 0.5)$$

$$U_A \sim \text{Normal}(\mu = -3, \sigma^2 = 1^2)$$

$$U_Y \sim \text{Normal}(\mu = 0, \sigma^2 = 0.3^2)$$

- Let us also specify the structural equations F :

$$W1 = f_{W1}(U_{W1}) = \mathbb{I}[U_{W1} < 0.35]$$

$$W2 = f_{W2}(W1, U_{W2}) = W1 + 2 * U_{W2}$$

$$A = f_A(W1, W2, U_A) = \mathbb{I}[(1 + W1 + 2 * W2 + U_A) > 0]$$

$$Y = f_Y(W1, A, U_Y) = 1 + 2.5 * A + 3 * W1 - 0.25 * A * W1 + U_Y$$

where $\mathbb{I}[\cdot]$ is the indicator function and equal to 1 if the statement in the brackets is true.

- The distribution of exogenous terms U and the set of structural equations F give us the joint distribution of (U, X) , denoted $\mathbb{P}_{U,X}$. This distribution is an element of (i.e. is one possible distribution compatible with) the SCM: $\mathbb{P}_{U,X} \in \mathcal{M}^{\mathcal{F}}$.
1. **Evaluate $\Psi^{\mathcal{F}}(\mathbb{P}_{U,X})$ for this data generating process.**
 2. **Interpret $\Psi^{\mathcal{F}}(\mathbb{P}_{U,X})$.**

Solution:

1. In this particular data generating system (one of many compatible with the SCM), the expectation of the counterfactual outcome is a linear function of the treatment level a , covariate $W1$ and random error U_Y :

$$\begin{aligned}\mathbb{E}_{U,X}[Y_a] &= \mathbb{E}_{U,X}[1 + 2.5*a + 3*W1 - 0.25*a*W1 + U_Y] \\ &= 1 + 2.5*a + 3*\mathbb{E}_{U,X}[W1] - 0.25*a*\mathbb{E}_{U,X}[W1] + \mathbb{E}_{U,X}[U_Y]\end{aligned}$$

We also know that $W1$ is a Bernoulli random variable with probability 0.35:

$$\mathbb{P}_{U,X}(W1 = 1) = \mathbb{E}_{U,X}[W1] = \mathbb{E}_{U,X}(\mathbb{I}[U_{W1} < 0.35]) = 0.35$$

Finally, the mean value of U_Y is zero:

$$\mathbb{E}_{U,X}[U_Y] = 0$$

Therefore, the true value of the target causal parameter is

$$\begin{aligned}\Psi^{\mathcal{F}}(\mathbb{P}_{U,X}) &= \mathbb{E}_{U,X}(Y_1 - Y_0) \\ &= 1 + 2.5*1 + 3*0.35 - 0.25*1*0.35 + 0 - (1 + 2.5*0 + 3*0.35 - 0.25*0*0.35 + 0) = 2.4125\end{aligned}$$

2. The counterfactual expected happiness would be 2.4125 units higher if all wizards consumed butterbeer than if none of the wizards consumed butterbeer.

2.1 Translating this data generating process for $\mathbb{P}_{U,X}$ into simulations

1. **First, set the seed to 252. Type `set.seed(252)`.** The `set.seed` function ensures the same values are generated each time each time we run all of our code. More information about this function can be accessed with `?set.seed`. Briefly, R generates the exogenous input U by calling a pseudorandom number generator to simulate a *Uniform*(0,1) variable and then transforming it to correspond with a draw from the specified distribution. For example, a Bernoulli random variable with probability p is generated as an indicator that this *Uniform*(0,1) is $< p$.
2. **Set `n=5000` as the the number of draws from $\mathbb{P}_{U,X}$.**
3. **Simulate the background factors U .** To simulate from the uniform distribution, use `runif()` function, setting the minimum and maximum values. To simulate from a Bernoulli with probability p , use the `rbinom()` function, setting the `size=1`. To simulate from a normal distribution, use the `rnorm()` function, setting the mean and *standard deviation*. More information on these functions can be found by typing `?runif`, `?rnorm` or `?rbinom`.
4. **Evaluate the structural equations F to deterministically generate the endogenous nodes X .** There are several ways to code indicator functions. One way is with logical variables and the `as.numeric` function. Consider the following example.

```

> # assign x=20 and x2=30
> x <- 20
> x2 <- 30
> # create a logical variable x3 from the result of the test x > x2
> x3 <- (x > x2)
> x3

[1] FALSE

> # we can convert the logical variable into binary with the as.numeric function
> # R converts TRUE to 1 and FALSE to 0
> x3 <- as.numeric(x > x2)
> x3

[1] 0

```

5. **Create a data frame X to hold these values.** The rows are the n repetitions of the experiment and the columns are the random variables. Use the `data.frame` function. Then use the `head` and `summary` functions to get a better understanding of the data. (More information on these functions can be obtained with `?data.frame`, `?head` and `?summary`.)

Solution:

```

> # 1. setting the seed
> set.seed(252)

> #2. setting the number of observations
> n <- 5000

> #3. simulating the Us
> U.W1 <- runif(n, min=0, max=1)
> U.W2 <- rbinom(n, size=1, prob=0.5)
> U.A <- rnorm(n, mean=-3, sd=1)
> U.Y <- rnorm(n, mean=0, sd=0.3)

> #4. Given the random input, deterministically evaluate the structural equations F
> W1 = as.numeric( U.W1 < 0.35)
> W2 = W1 + 2*U.W2
> A = as.numeric( 1+W1+2*W2 + U.A > 0)
> Y = 1 +2.5*A + 3*W1 - 0.25*A*W1 + U.Y

> X <- data.frame(W1, W2, A, Y)
> head(X)

```

	W1	W2	A	Y
1	0	2	1	3.3561085
2	0	0	0	0.8793806
3	1	1	0	3.3546226
4	0	2	1	3.5474294
5	0	0	0	0.4332530
6	0	2	1	3.6398616

```
> summary(X)
```

	W1	W2	A	Y
Min.	:0.0000	Min. :0.000	Min. :0.0000	Min. :0.07914
1st Qu.	:0.0000	1st Qu.:0.000	1st Qu.:0.0000	1st Qu.:1.22898
Median	:0.0000	Median :2.000	Median :1.0000	Median :3.50644
Mean	:0.3372	Mean :1.338	Mean :0.6456	Mean :3.55352
3rd Qu.	:1.0000	3rd Qu.:2.000	3rd Qu.:1.0000	3rd Qu.:6.00153
Max.	:1.0000	Max. :3.000	Max. :1.0000	Max. :7.16974

2.2 Generate the counterfactual outcomes & obtain the value of the target causal parameter

1. Intervene on the above data generating system to set butterbeer consumption $a = \{0, 1\}$ and generate counterfactual outcomes Y_a .
2. Add the resulting columns to the X matrix using either the `data.frame` or `cbind` function.
3. Does the counterfactual value Y_a equal the observed Y when $A = a$?
4. Above, we evaluated the average treatment effect in closed form. **Estimate $\Psi^{\mathcal{F}}(\mathbb{P}_{U,X})$ with the simulated counterfactuals.** Given a large sample of i.i.d. draws of (Y_0, Y_1) , we can closely approximate $\mathbb{E}_{U,X}[Y_1 - Y_0]$ with the sample average of the observed values of $Y_1 - Y_0$.

Solution:

```
> # 1. intervene to set A=a and generate the counterfactual outcomes Y.a
> Y.1<- 1 +2.5*1 + 3*W1 - 0.25*1*W1 + U.Y
> Y.0<- 1 +2.5*0 + 3*W1 - 0.25*0*W1 + U.Y

> # 2. add these columns to the dataframe
> X<- data.frame(X, Y.1, Y.0)
> head(X)

  W1 W2 A      Y      Y.1      Y.0
1  0  2  1 3.3561085 3.356108 0.8561085
2  0  0  0 0.8793806 3.379381 0.8793806
3  1  1  0 3.3546226 5.604623 3.3546226
4  0  2  1 3.5474294 3.547429 1.0474294
5  0  0  0 0.4332530 2.933253 0.4332530
6  0  2  1 3.6398616 3.639862 1.1398616

>
> #3. YES! We see the outcome Y when A=0 is equal to the counterfactual outcome Y.0.
> # Likewise, the outcome Y when A=1 is equal to the counterfactual outcome Y.1.
> # Given input of the background factors U, the structural equations are deterministic.

> #4. Evaluation
> Psi.F<- mean(Y.1 - Y.0)
> Psi.F
```

```
[1] 2.4157
```

```
> # Our answers match! :)
```

3 Target causal parameters defined with a MSM

Suppose we are now interested in summarizing how the expectation of counterfactual happiness Y_a varies as a function of cups of butterbeer consumed a . We will assume that the latest wizarding technology allows precise measurement of the amount of butterbeer consumed so that a may take on decimal values. For simplicity, ignore baseline covariates.

We could use a marginal structural model (MSM) to summarize how the expected counterfactual outcome changes as function of the exposure. Suppose, however, that we do not know the exact shape of the butterbeer-happiness (dose-response) curve. Therefore, we use the following *working* MSM to define the target parameter β as the projection of the true causal curve $\mathbb{E}_{U,X}(Y_a)$ onto a summary model $m(a|\beta)$:

$$\beta(\mathbb{P}_{U,X}|m) = \operatorname{argmin}_{\beta} \mathbb{E}_{U,X} \left[\sum_{a \in \mathcal{A}} (Y_a - m(a|\beta))^2 \right]$$

$$m(a|\beta) = \beta_0 + \beta_1 a$$

The causal parameter is then the value of the β coefficients that minimize the sum of squared residuals between the counterfactual outcomes Y_a and the predicted $m(a|\beta)$ for all possible exposure levels $a \in \mathcal{A}$. In this case, we are using a linear summary.

3.1 A specific data generating process

Consider the following data generating process for $\mathbb{P}_{U,X}$:

$$\begin{aligned} U_A &\sim \text{Uniform}(\min = 0, \max = 2) \\ U_Y &\sim \text{Normal}(\mu = 0, \sigma^2 = 0.3^2) \\ A &= f_A(U_A) = 2 * U_A \\ Y &= f_Y(A, U_Y) = 4 + 9 * A - 2.25 * A^2 + U_Y \end{aligned}$$

where $U_A \perp\!\!\!\perp U_Y$. Suppose we are interested in the counterfactual happiness Y_a under butterbeer consumption $a \in \mathcal{A} = \{0, 1, 2, 3, 4\}$ cups.

1. For $n = 5$ wizards, generate the exogenous factors $U = (U_A, U_Y)$. Then set $A = 0$ to simulate the counterfactual happiness under 0 cups of butterbeer $Y.0$. Repeat for each of the 4 other doses of butterbeer to generate counterfactual outcomes $Y.1$, $Y.2$, $Y.3$ and $Y.4$.
2. Use the `c()` function to combine (stack) all the counterfactual outcomes ($Y.0$, $Y.1$, $Y.2$, $Y.3$, $Y.4$) into a single vector $Y.a$.

```
> n <- 5
```

```
> Y.a <- c(Y.0, Y.1, Y.2, Y.3, Y.4)
```

3. Use the `rep()` function to create vector a with the corresponding exposures.

```
> a <- c(rep(0, n), rep(1, n), rep(2, n), rep(3, n), rep(4, n))
```

```
> a
```

```
[1] 0 0 0 0 0 1 1 1 1 1 2 2 2 2 2 3 3 3 3 3 4 4 4 4 4
```

4. Create data frame $X = (a, Y_a)$, consisting of the set treatment levels and corresponding counterfactual outcomes.
5. Plot the counterfactual outcomes Y_a as a function of a . More information can be accessed with `?plot` and more plotting options with `?par`.

 > # Hint: this code would plot x2 as a function of x and add a title
 > plot(x, x2, main='Plot of x2 by x')
6. Based on your knowledge of the data generating system as well as the graph, do you think a working MSM with a linear form is a good summary?

Solution:

```
> # 1. making the data frame
> n=5
> # generating the exogenous Us
> U.A <- runif(n, min=0, max=2)
> U.Y <- rnorm(n, mean=0, sd=0.3)
> # Evaluating the counterfactuals
> Y.0<- 4 +9*0-2.25*0*0+ U.Y
> Y.1<- 4 +9*1-2.25*1*1+ U.Y
> Y.2<- 4 +9*2-2.25*2*2+ U.Y
> Y.3<- 4 +9*3-2.25*3*3+ U.Y
> Y.4<- 4 +9*4-2.25*4*4+ U.Y

> # 2. making the dataframe
> # c() combines values into a vector.
> Y.a<- c(Y.0, Y.1, Y.2, Y.3, Y.4)
> #
> # To get a vector of the corresponding treatments use the rep() function
> a<- c( rep(0,n), rep(1, n), rep(2, n), rep(3, n), rep(4, n) )
> #
>
> X<- data.frame(a,Y.a)
> head(X)

  a      Y.a
1 0  4.244533
2 0  3.930170
3 0  4.502727
4 0  3.775410
5 0  4.077615
6 1 10.994533


> summary(X)

      a      Y.a
Min.   :0    Min.   : 3.775
```



```
1st Qu.:1    1st Qu.: 4.245
Median :2    Median :10.680
Mean   :2    Mean   : 8.606
3rd Qu.:3    3rd Qu.:11.253
Max.   :4    Max.   :13.503
```

```
> # 3. Plotting the counterfactuals
>
> # plot the counterfactual happiness as a function of cups of butterbeer consumed
> # adding a title and label axes
> plot(a, Y.a, main='Counterfactual happiness as a function of cups of butterbeer',
+       xlab='Cups of butterbeer a', ylab='Counterfactual happiness Y.a')
```



Solution Fig. 2: Plot the counterfactual outcomes Y_a as a function of a .

- 4 Based on our knowledge of the data generating system as well as the previous graph, the working MSM is a poor summary of how counterfactual happiness changes with cups of butterbeer consumed. The true functional relationship is quadratic.

3.2 Obtain the value of the target causal parameter

1. We have the counterfactuals Y_a for $a \in \mathcal{A} = \{0, 1, 2, 3, 4\}$ and have defined the target parameter using the least squares projection (i.e. with the L2 loss function). Use `glm()` function to fit the coefficients of the $m(a|\beta)$. Interpret the results.
2. Add the projected causal curve to the graph. See `abline()`.
3. *Bonus:* Add the true causal curve to the graph. Use the `seq()` function to generate lots of doses between $A = 0$ and $A = 4$. Then evaluate the mean counterfactual $\mathbb{E}_{U,X}(Y_a) = 4 + 9a - 2.25a^2$ at these points. Finally, use the `lines()` function to add the line.

Solution:

```
> #####
> # 1. Evaluate the value of the target causal parameter.
> # We can use ordinary least squares regression to obtain betas.
> workMSM<- glm(Y.a ~ a)
> #
> # the target parameter are the coefficient value that minimize the L2 risk function
> workMSM
```

```
Call: glm(formula = Y.a ~ a)
```

```
Coefficients:
```

```
(Intercept)          a
 8.606e+00   -3.140e-16
```

```
Degrees of Freedom: 24 Total (i.e. Null);  23 Residual
```


```
Null Deviance:          356
```

```
Residual Deviance: 356      AIC: 143.3
```

Projecting the causal dose-response curve onto the working MSM yields coefficients of $\beta_0 = 8.508$ and $\beta_1 = 0$. In words, summarizing the true causal dose-response curve with a linear working MSM suggests that butterbeer consumption has no linear effect on mean counterfactual happiness.

Note: If instead we believed the MSM was true (i.e. the dose-response curve actually has a linear form), then a coefficient of $\beta_1 = 0$ would incorrectly suggest that butterbeer consumption has no effect on mean counterfactual happiness. The MSM would suggest that the expected counterfactual happiness is 8.508 units, regardless of number of cups of butterbeer consumed.

```
> #####
> #2. Add the projected causal curve to the graph.
> plot(a, Y.a, main='Counterfactual happiness as a function of cups of butterbeer',
+      xlab='Cups of butterbeer A', ylab='Counterfactual happiness Y.a')
> abline(workMSM, col='red', lwd=2)
> # we are specifying red for the color and increasing the line width (lwd)
>
> #####
> #3. Add the true causal curve to the graph.
> # using the seq() function, we can get a lot of points between A=0 and A=4
> A.cont<- seq(from=0, to=4, by=0.1)
> # then we can evaluate the mean counterfactual at all these points
> Y.a.cont<- 4 +9*A.cont-2.25*A.cont*A.cont
> lines(A.cont, Y.a.cont, col='green', lwd=2)
> # we are specifying green for the color and increasing the line width (lwd)
>
> # we can add a legend
> legend('topright', c('Working MSM', 'True'), col=c('red', 'green'), lwd=2)
```



Solution Fig. 3: Plot the counterfactual outcomes Y_a as a function of a . The true causal curve is shown in green, while its projection onto the working MSM is shown in red.