36-617: Applied Linear Regression

Causal Inference

Brian Junker
132E Baker Hall
brian@stat.cmu.edu

Welcome back!

- I hope you had a great midsemester break!
- Here are a few notes that you may be interested in:
  - There are solutions for the take-home midterm in the **"take-home midterm"** folder on Canvas.
  - There is a new folder on Canvas called **"0 - midterm grades"** that contains a summary of grades for the whole class, for midsemester grade reports.
  - There is another new folder on Canvas called **"0 – midsemester evaluation"** with a report on how you evaluated the class so far this semester.

Announcements

- Quiz next Monday (not today!)
  - Focus on Sheather 10.1
- HW06 out soon
  - Nonparametric regression & causal reasoning
  - Is Wed better than Mon for HW due date?
- Reading
  - This week: Causal Reasoning
    - G&H Ch's 9 & 10 (see pdf in week08 area)
  - Next week: Intro to Multilevel Models
    - Sheather 10.1
    - Intercepts: G&H Ch 12 (see pdf in week08 area)
    - Slopes: G&H Ch 13 (see pdf in week08 area)

Outline

- 15 Causal Inference [G&H Ch 9]
  - The Fundamental Problem
  - Confounders, and how Controlled Randomized Trials control them
  - Adjusting an analysis for pre-treatment covariates (but not post-treatment ones!)
- 16 More sophisticated tools for causal inference [G&H Ch 10]
  - Observational Studies
  - Instrumental Variables
  - Matching and propensity scores
  - Regression discontinuity designs

Causal Inference

- Want to test a new pain reliever for headaches
- Have 200 patients i=1,…,200.
  - T_i=1 (patient gets drug) for i=1..100,
  - T_i=0 (patient gets nothing) for i=101..200.
- Suppose drug is worthless, but
  - i=1..100 are healthy and
  - i=101..200 all have flu, colds, etc.
  - *How will the drug look?*
- Suppose drug is effective, but
  - i=1..100 have colds & flu, and
  - i=101..200 are healthy.
  - *How will the drug look now?*
- What is wrong with these examples?

Causal Inference—The Fundamental Problem

- We really would like to see the difference between pain level "with the drug" vs pain level "without", *for each individual patient.*

$$\begin{aligned} y_i^0 &= \text{outcome without treatment} \\ y_i^1 &= \text{outcome with treatment} \\ y_i^1 - y_i^0 &= \text{treatment effect for unit } i \end{aligned}$$

- But we cannot try the drug, and then go back in time and try without the drug.
  - *For each patient i, can see either y_i^0 or y_i^1 but not both!*

Causal Inference—The Fundamental Problem

- If we average the individual treatment effect over all patients, get the average causal effect (ACE):

$$\begin{aligned} \text{ACE} &= \frac{1}{N}\sum_{i=1}^{N}(y_i^1 - y_i^0) = \frac{1}{N}\sum_{i=1}^{N} y_i^1 - \frac{1}{N}\sum_{i=1}^{N} y_i^0 \\ &= E[y^1] - E[y^0] \end{aligned}$$

- Most studies try to estimate ACE. A good way to do this would be:
  - Estimate $E[y^1] \approx \bar{y}^1$ from unbiased sample $y_1^1, \ldots y_{n_1}^1$
  - Estimate $E[y^0] \approx \bar{y}^0$ from unbiased sample $y_1^0, \ldots y_{n_0}^0$

Causal Inference—The Fundamental Problem

- The problem with the examples we started with was that ***the samples were not unbiased***.
- There are basically two ways to deal with bias
  - Design a study for which the samples are guaranteed to be unbiased
  - Do some statistical adjustment to account for the bias
- To understand how to design an "unbiased" study, we first consider how "bias" arises…

Causal inference - Confounders

- If some patients have $T_i = 1$ and others have $T_i = 0$, we know that $E[y^1] - E[y^0] \approx \hat{\beta}_1$ in the regression

$$y_i = \beta_0 + \beta_1 T_i + \epsilon_i$$

- However, if there is a "confounding" variable x_i , the correct $\hat{\beta}_1$ should come from

$$y_i = \beta_0 + \beta_1 T_i + \beta_2 x_i + \epsilon_i$$

- How bad can the bias be if we omit x_i?

Causal inference - Confounders

We suppose the correct model is

$$y_i = \beta_0 + \beta_1 T_i + \beta_2 x_i + \epsilon_i \tag{1}$$

but we fit instead

$$y_i = \beta_0^* + \beta_1^* T_i + \epsilon_i^* \tag{2}$$

Note that x_i also has some relationship with T_i that can be expressed as a linear regression:

$$x_i = \gamma_0 + \gamma_1 T_i + \nu_i \tag{3}$$

If we substitute (3) into (1) and do a little rearranging, we get

$$y_i = (\beta_0 + \beta_2\gamma_0) + (\beta_1 + \beta_2\gamma_1)T_i + (\epsilon_i + \beta_2\nu_i) \tag{4}$$

Equating coefficients in (2) and (4), we see

$$\beta_1^* = \beta_1 + \beta_2\gamma_1 \tag{5}$$

Thus, estimating $E[y^1] - E[y^0] \approx \hat{\beta}_1^*$ will be biased, *unless*

- $\gamma_1 = 0$, i.e. x_i is independent of treatment assignment T_i
- $\beta_2 = 0$, i.e. x_i has no influence on y_i after considering T_i (x_i not really a confounder!)

Causal inference - Confounders

- If X is a confounder, the total effect of T on Y is $\beta_1 + \beta_2\gamma_1$:

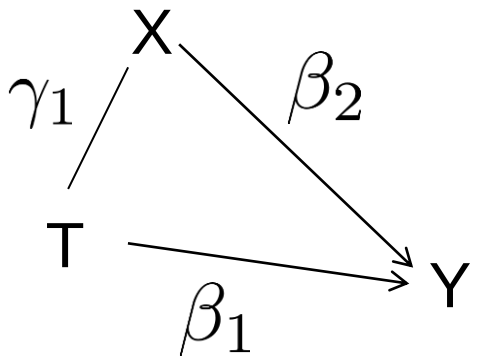

- $\beta_2 = 0$: X not really a confounder!
- $\gamma_1 = 0$: No selection effect!

- If we omit X (or it is hidden!) then we only get the right answer from y = β_o + β_1 T + ϵ, if β_2 or γ_1 is zero.

Causal inference – Estimating ACE

- We can get an unbiased estimate of ACE in any of the following ways
  - *If there are no confounders*, estimate β_1 in

$$y_i = \beta_0 + \beta_1 T_i + \epsilon_i$$

  - *If there are confounders*, ***find them all***, include them as x's, and then estimate β_1 in

$$y_i = \beta_0 + \beta_1 T_i + \beta_2 x_{2i} + \beta_3 x_{3i} + \cdots + \beta_K x_{Ki} + \epsilon_i$$

  - *Design the experiment* so that all ***confounders*** x_i are ***independent of treatment*** assignment T_i and then estimate β_1 from

$$y_i = \beta_0 + \beta_1 T_i + \epsilon_i$$

Causal inference – randomized trials

- In a ***randomized experiment***, each unit i is assigned $T_i = 1$ (treatment) or $T_i = 0$ (no tx) randomly (e.g. by random coin toss!).
  - This forces every potential confounder x_i to be independent of T_i, whether we "discover" x_i or not! $(\gamma_1 = 0)$
  - From a randomized experiment we can always estimate ACE by estimating β_1 in

$$y_i = \beta_0 + \beta_1 T_i + \epsilon_i$$

Causal inference – randomized trials

- In many settings you can't completely randomize
  - A study of effectiveness of a new math curriculum might involve several schools.
    - Can't put all math classes in all schools together in one "pot" and randomly assign some to new math curriculum
    - Instead assign ½ the classes to the new math program and ½ to the old math program *within each school*
    - Since schools contain other factors that affect math performance, school becomes an x_i and we can estimate the ACE for the new math program from

$$y_i = \beta_0 + \beta_1 T_i + \beta_2 x_i + \epsilon_i$$

- A lot of experimental design is like this…

Causal inference – pre-treatment covariates in randomized trials

- Even in a randomized experiment, if we can identify a confounder x_i, it is good to include it in the model.
- Estimating ACE = $\hat{\beta}_1$ from

$$y_i = \beta_0 + \beta_1 T_i + \epsilon_i$$

  is unbiased, but not efficient (more uncertainty)
- Estimating ACE = $\hat{\beta}_1$ from

$$y_i = \beta_0 + \beta_1 T_i + \beta_2 x_i + \epsilon_i$$

  will be more efficient (less uncertainty).

Causal inference – randomized trials

- If R is a random treatment assignment (coin flip!), then γ_1 must equal zero!

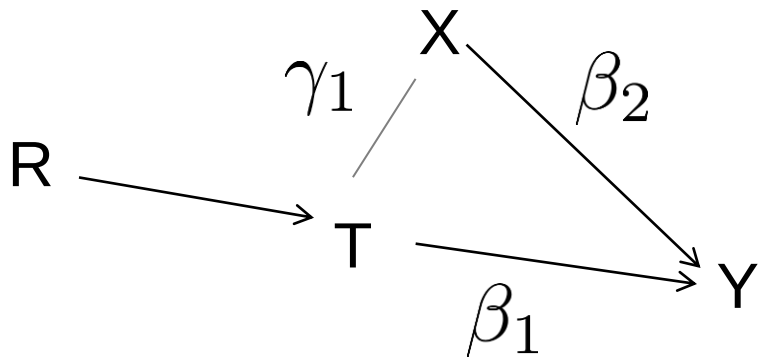

- $\gamma_1 = 0$: No selection effect!

- We can now get the right treatment effect from

$$y = \beta_0 + \beta_1 T + \epsilon.$$

- It is still worth including X in the model if possible,

$$y = \beta_0 + \beta_1 T + \beta_2 X + \epsilon$$

because including X will reduce SE(β_1) !

Randomized trials – pre-treatment covariates – uniform tx effect

```
> n <- 100; x <- rnorm(n)
> y <- 60 + 10*x + 5*rnorm(n)
  # x is a confounder
> T <- rbinom(n,1,.5)
  # treatment by random experiment
> y <- ifelse(T==1,y+20,y)
  # add treatment effect for treated
> plot(x,y,col=T+2)
> legend(-3,100,pch=c(1,1),col=2:3,
      legend=c("Non-treated","Treated"))
> (ACE <- mean(y[T==1]) - mean(y[T==0]))
[1] 20.26647
>
> summary(lm(y ~ T))$coef[,1:2]
             Estimate Std. Error
(Intercept) 60.63675   1.854682
T           20.26647   2.523902
>
> summary(lm(y ~ T + x))$coef[,1:2]
             Estimate Std. Error
(Intercept) 60.13741  0.6815005
T           19.49961  0.9275130
x           10.49448  0.4182943
```

ACE is estimated better when covariate in the model

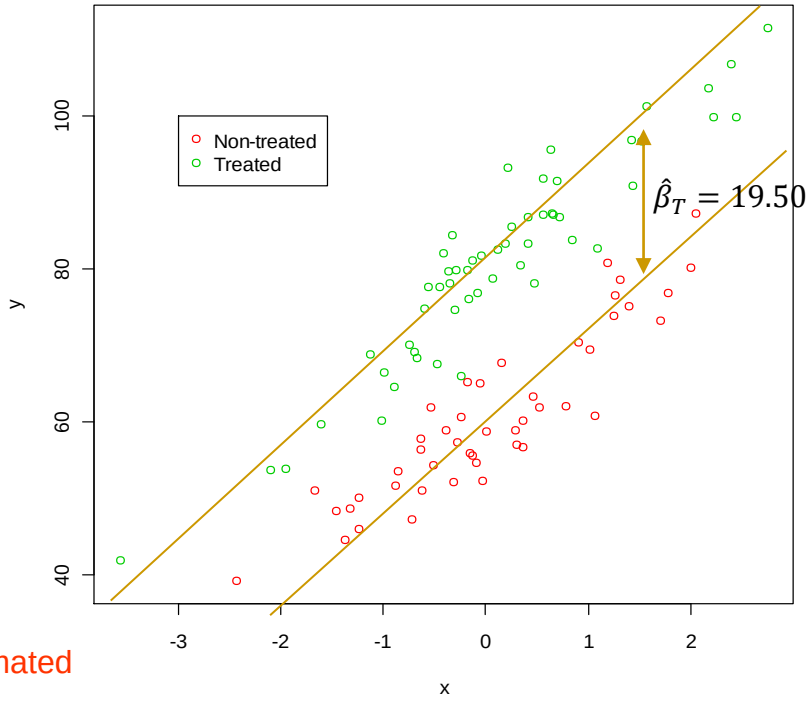

- x is a pretest score
- y is a post-test score, of course affected by x
- T is treatment (new curriculum)

Removing variability due to X from Y reduces uncertainty about $\hat{\beta}_T$

```
> y <- 60 + 10*x + 5*rnorm(n)
> T <- rbinom(n,1,.5)
> y <- ifelse(T==1,y+20,y)
>
> par(mfrow=c(2,1))
> plot(density(simple<-y[T==1]-y[T==0]))
> abline(v=mean(y[T==1])-mean(y[T==0]))
>
> kn <- c(-100, seq(min(x), max(x),
+    length=20))
> deltas <- NULL
> for (k in 2:length(kn)) {
+    ok <- ((x > kn[k-1])&(x<=kn[k]))
+    y.tx <- y[(T==1)&ok]
+    y.ct <- y[(T==0)&ok]
+    if(length(y.tx)&length(y.ct)) {
+        deltas <- c(deltas,mean(y.tx)-
+                    mean(y.ct))
+     }}
> plot(density(deltas))
> abline(v=mean(deltas))
```

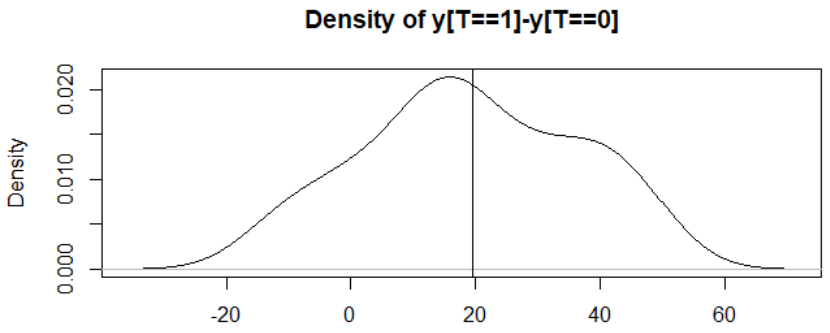

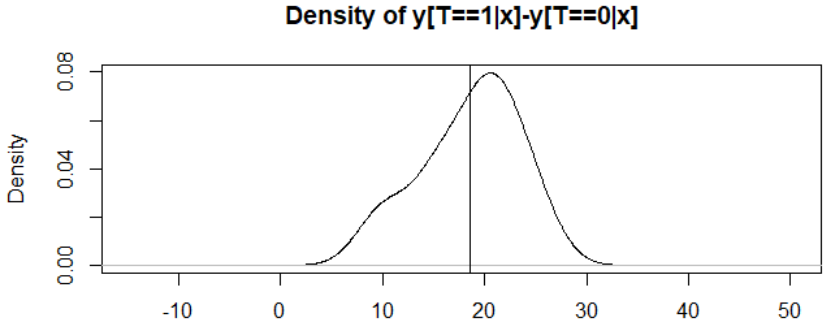

Randomized trials – pre-treatment covariates – nonuniform tx effect

```
> n <- 100
> x <- rnorm(n)
> y <- 60 + 10*x + 5*rnorm(n)
> T <- rbinom(100,1,.5)
> y <- ifelse(T==1,y+5+15*x,y)
> plot(x,y,col=T+2)
> legend(-2,100,pch=c(1,1),col=2:3,
     legend=c("Non-treated","Treated"))
> (ACE <- mean(y[T==1]) - mean(y[T==0]))
[1] 5.684276
> summary(lm(y ~ T))$coef[,1:2]
             Estimate Std. Error
(Intercept) 62.599809   3.164975
T            5.684276   4.229376
> (coef <- summary(lm(y ~ T + x +
    T:x))$coef[,1:2])
             Estimate Std. Error
(Intercept) 59.205524  0.8095489
T            6.149310  1.0646086
x            9.499872  0.6574682
T:x         15.653435  0.9527179
> mean(coef[2,1] + coef[4,1]*x)
[1] 9.631048
```

Tx affects not only the intercept but also the slope!

ACE not all that meaningful

Here's a kind of ACE that "might" be useful (???)…

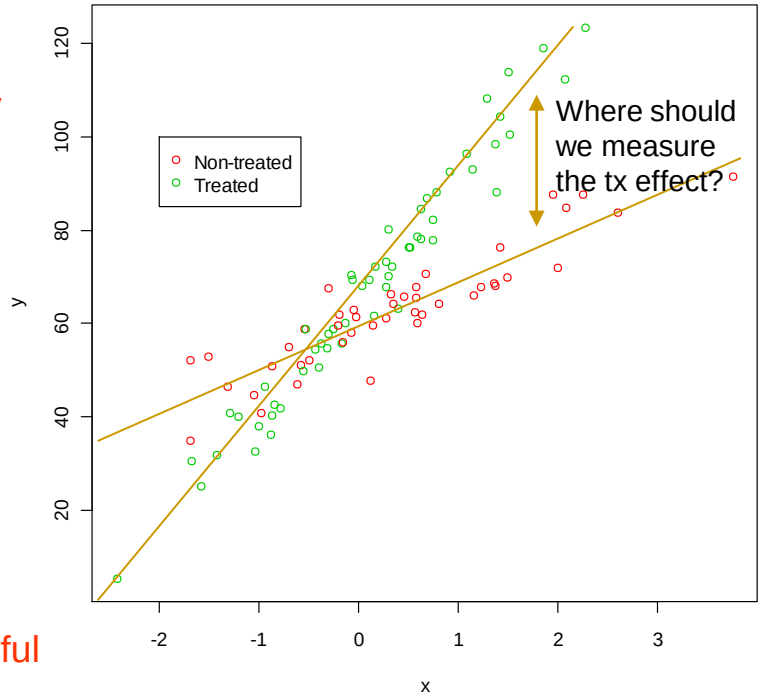

- x is a pretest score
- y is a post-test score, of course affected by x
- T is treatment (new curriculum)

Causal inference – Post-tx covariates

- If R is a random treatment assignment (coin flip!), then γ_1 must equal zero!

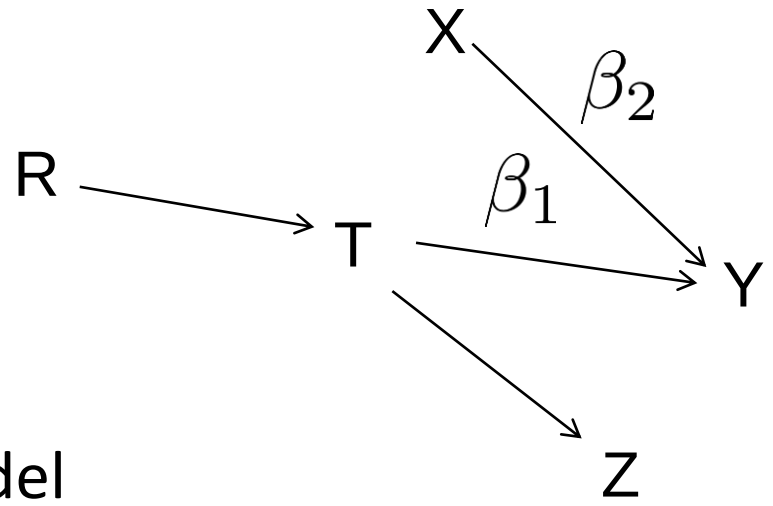

- In the model

$$y = \beta_0 + \beta_1 T + \beta_2 X + \beta_3 Z + \epsilon$$

the estimate of β_1 will only include the influence of the part of T not explained by Z… That might not be much!

Randomized trials – ***do not include*** post-treatment covariates!

```
> n <- 100
> x <- rnorm(n)
> y <- 60 + 10*x + 5*rnorm(n)
> T <- rbinom(100,1,.5)
> y <- ifelse(T==1,y+20,y)
> z <- ifelse(T==1,rnorm(100,3),
    rnorm(100,-3))
> plot(x,y,col=T+2)
> legend(-2,100,pch=c(1,1),col=2:3,
    legend=c("Non-treated","Treated"))
> (ACE <- mean(y[T==1]) -
    mean(y[T==0]))
[1] 22.43931
> summary(lm(y ~ T))$coef[,1:2]
             Estimate Std. Error
(Intercept) 58.11903   1.660045
T           22.43931   2.347659
```

```
> summary(lm(y ~ T +
    x))$coef[,1:2]
             Estimate Std. Error
(Intercept) 59.85651  0.7068169
T           20.78911  0.9959064
x           10.58185  0.4983279
> summary(lm(y ~ T + x +
    z))$coef[,1:2]
             Estimate Std. Error
(Intercept) 64.884033  1.9499540
T           10.505663  3.8573971
x           10.416234  0.4859765
z            1.608895  0.5843686
```

Including z in the model completely dilutes the effect of T that we are trying to estimate!

- x is a pretest score
- y is a post-test score, of course affected by x
- T is treatment (new curriculum)
- z is a secondary effect of T

Summary

- 15 Causal Inference [G&H Ch 9]
  - The Fundamental Problem
  - Confounders, and how Controlled Randomized Trials control them
  - Adjusting an analysis for pre-treatment covariates (but not post-treatment ones!)
- 16 More sophisticated tools for causal inference [G&H Ch 10]
  - Observational Studies
  - Instrumental Variables
  - Matching and propensity scores
  - Regression discontinuity designs