4.3 — Categorical Data

ECON 480 • Econometrics • Fall 2022

Dr. Ryan Safner
Associate Professor of Economics
safner@hood.edu
ryansafner/metricsF22
metricsF22.classes.ryansafner.com

Contents

Categorical Variables

- **Categorical variables** place an individual into one of several possible *categories*
  - e.g. sex, season, political party
  - may be responses to survey questions
  - can be quantitative (e.g. age, zip code)
- In R: `character` or `factor` type data
  - `factor` $\implies$ specific possible categories

Question	Categories or Responses
Do you invest in the stock market?	__ Yes __ No
What kind of advertising do you use?	__ Newspapers __ Internet __ Direct mailings
What is your class at school?	__ Freshman __ Sophomore __ Junior __ Senior
I would recommend this course to another student.	__ Strongly Disagree __ Slightly Disagree __ Slightly Agree __ Strongly Agree
How satisfied are you with this product?	__ Very Unsatisfied __ Unsatisfied __ Satisfied __ Very Satisfied

Working with `factor` Variables in R

Factors in R I

- `factor` is a special type of `character` object class that indicates membership in a category (called a `level`)
- Suppose I have data on students:

id <dbl>	rank <chr>	grade <dbl>
1	Freshman	76
2	Junior	82
3	Sophomore	73
4	Sophomore	95
5	Senior	74

5 rows

- See that `Rank` is a `character` (`<chr>`) variable, just a string of text

Factors in R II

- We can make `rank` a `factor` variable, to indicate a student is a member of one of the possible categories: (freshman, sophomore, junior, senior)

```
students <- students %>%
  mutate(rank = as.factor(rank)) # overwrite and change class of Rank to factor

students %>% head(n = 5)
```

id <dbl>	rank <fct>	grade <dbl>
1	Freshman	76
2	Junior	82
3	Sophomore	73
4	Sophomore	95
5	Senior	74

5 rows

- See now it's a `factor` (`<fct>`)

Factors in R III

```
# what are the categories?
students %>%
  group_by(rank) %>%
  count()
```

rank <fct>	n <int>
Freshman	4
Junior	1
Senior	3
Sophomore	2

4 rows

```
# note the order is arbitrary! This is an "unordered" factor
```

Ordered Factors in R I

- If there is a rank order you wish to preserve, you can make an `ordered` (`factor`) variable
  - list the `levels` from 1st to last

```
students <- students %>%
  mutate(rank = ordered(rank, # overwrite and change class of Rank to ordered
                        # next, specify the levels, in order
                        levels = c("Freshman", "Sophomore", "Junior", "Senior")
                        )
         )

students %>% head(n = 5)
```

id <dbl>	rank <ord>	grade <dbl>
1	Freshman	76
2	Junior	82
3	Sophomore	73
4	Sophomore	95

id	rank	grade
<dbl>	<ord>	<dbl>
5	Senior	74

5 rows

Ordered Factors in R II

```r
students %>%
  group_by(rank) %>%
  count()
```

rank	n
<ord>	<int>
Freshman	4
Sophomore	2
Junior	1
Senior	3

4 rows

Example Research Question with Categorical Data

Example

How much higher wages, on average, do men earn compared to women?

A Difference in Group Means

- Basic statistics: can test for statistically significant difference in group means with a **t-test**[1], let:
- Y_M: average earnings of a sample of n_M men
- [Y_W]{.pink: average earnings of a sample of n_M women
- **Difference** in group averages: $d = \bar{Y}_M - \bar{Y}_W$
- The hypothesis test is:
  - $H_0 : d = 0$
  - $H_1 : d \neq 0$

1. See today's appendix for this example

Plotting `factors` in R

- Plotting `wage` vs. a `factor` variable, e.g. `gender` (which is either `Male` or `Female`) looks like this

Plot | Code

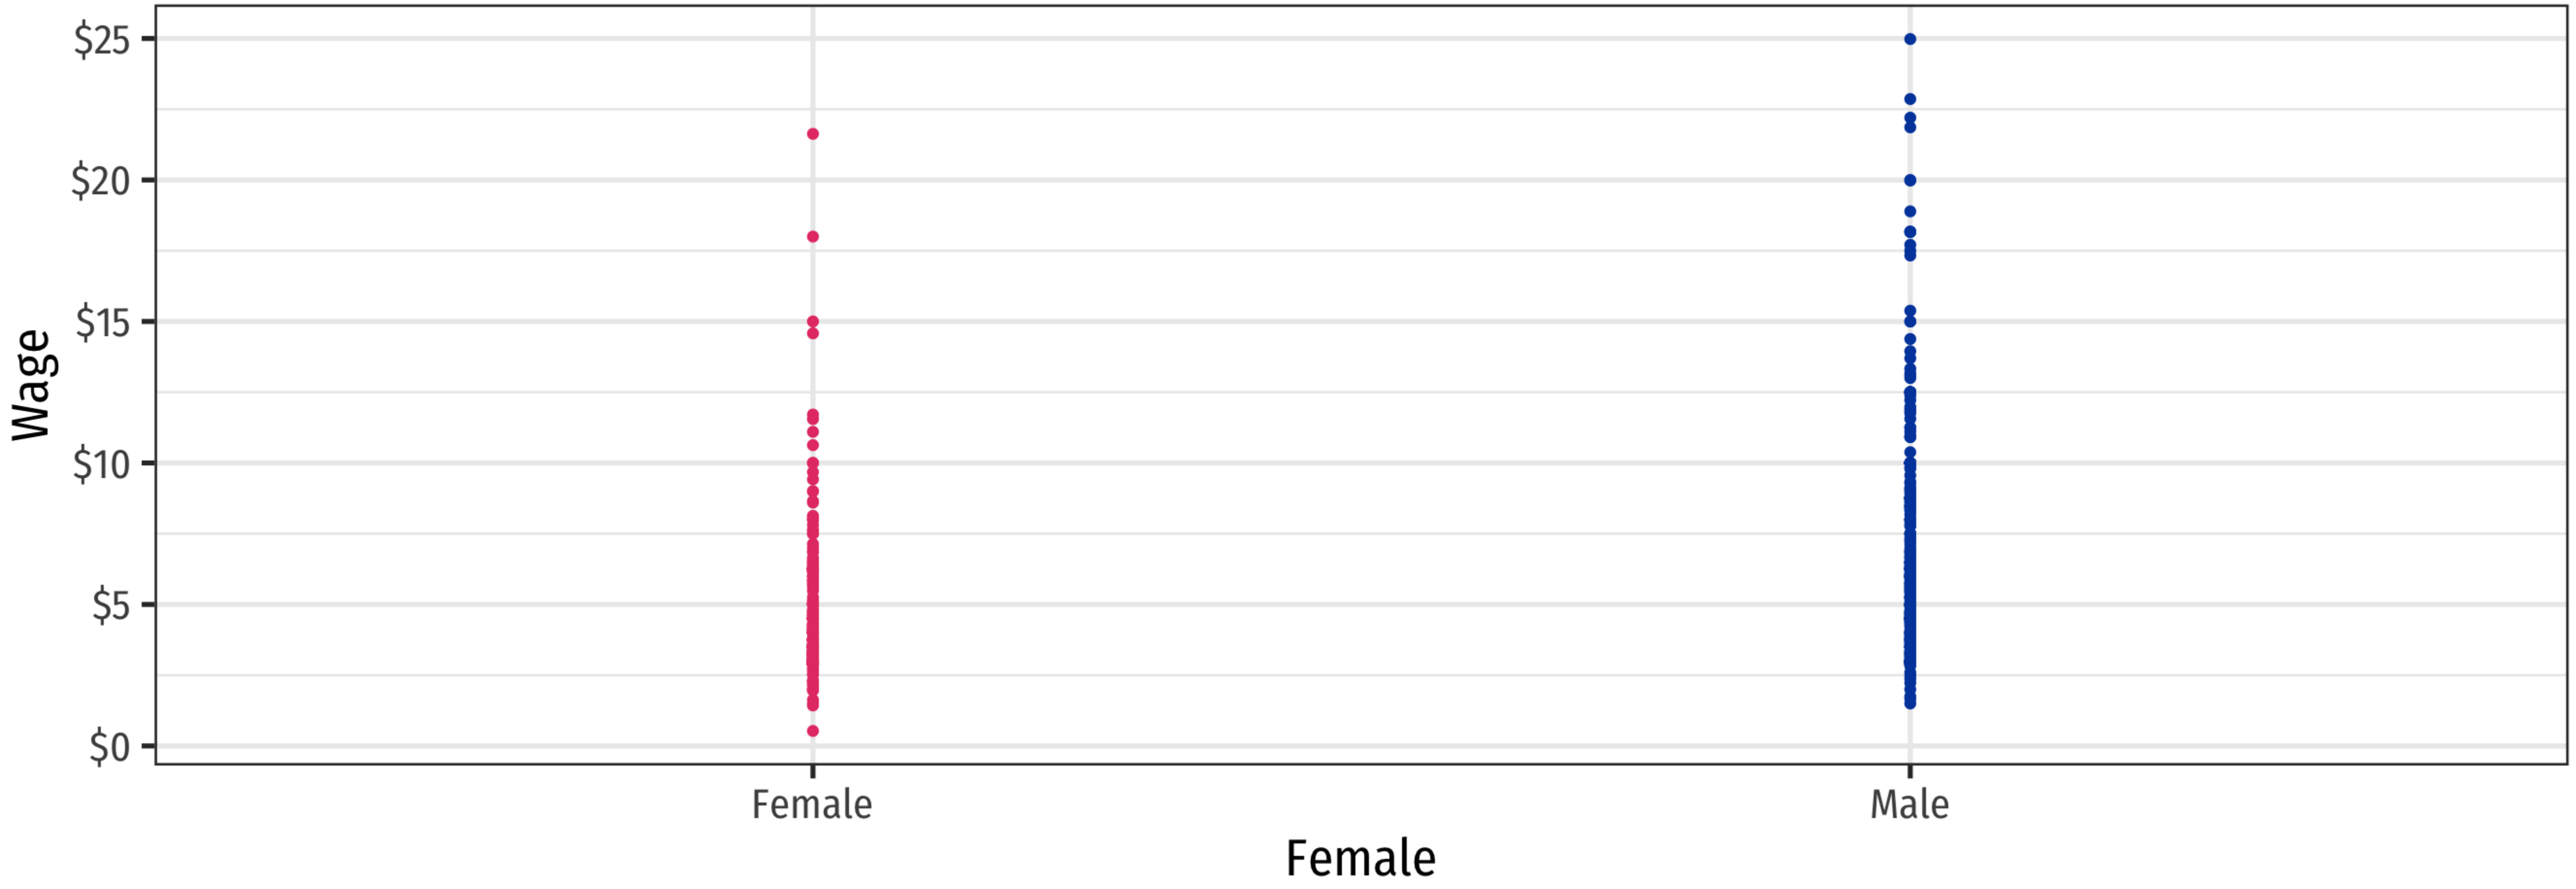

- Effectively R treats values of a factor variable as integers (e.g. `"Female"` = 0, `"Male"` = 1)
- Let's make this more explicit by making a **dummy variable** to stand in for gender

Regression with Dummy Variables

Comparing Groups with Regression

- In a regression, we can easily compare across groups via a **dummy variable**[1]
- Dummy variable *only* $= 0$ or $= 1$, if a condition is TRUE vs. FALSE
- Signifies whether an observation belongs to a category or not

Example

$$\widehat{Wage_i} = \hat{\beta_0} + \hat{\beta_1}\, Female_i \qquad \text{where } Female_i = \begin{cases} 1 & \text{if individual } i \text{ is } Female \\ 0 & \text{if individual } i \text{ is } Male \end{cases}$$

- Again, $\hat{\beta_1}$ makes less sense as the "slope" of a line in this context

1. Also called a **binary variable** or **dichotomous variable** since it only takes on 2 values.

Comparing Groups in Regression: Scatterplot

Plot | Code

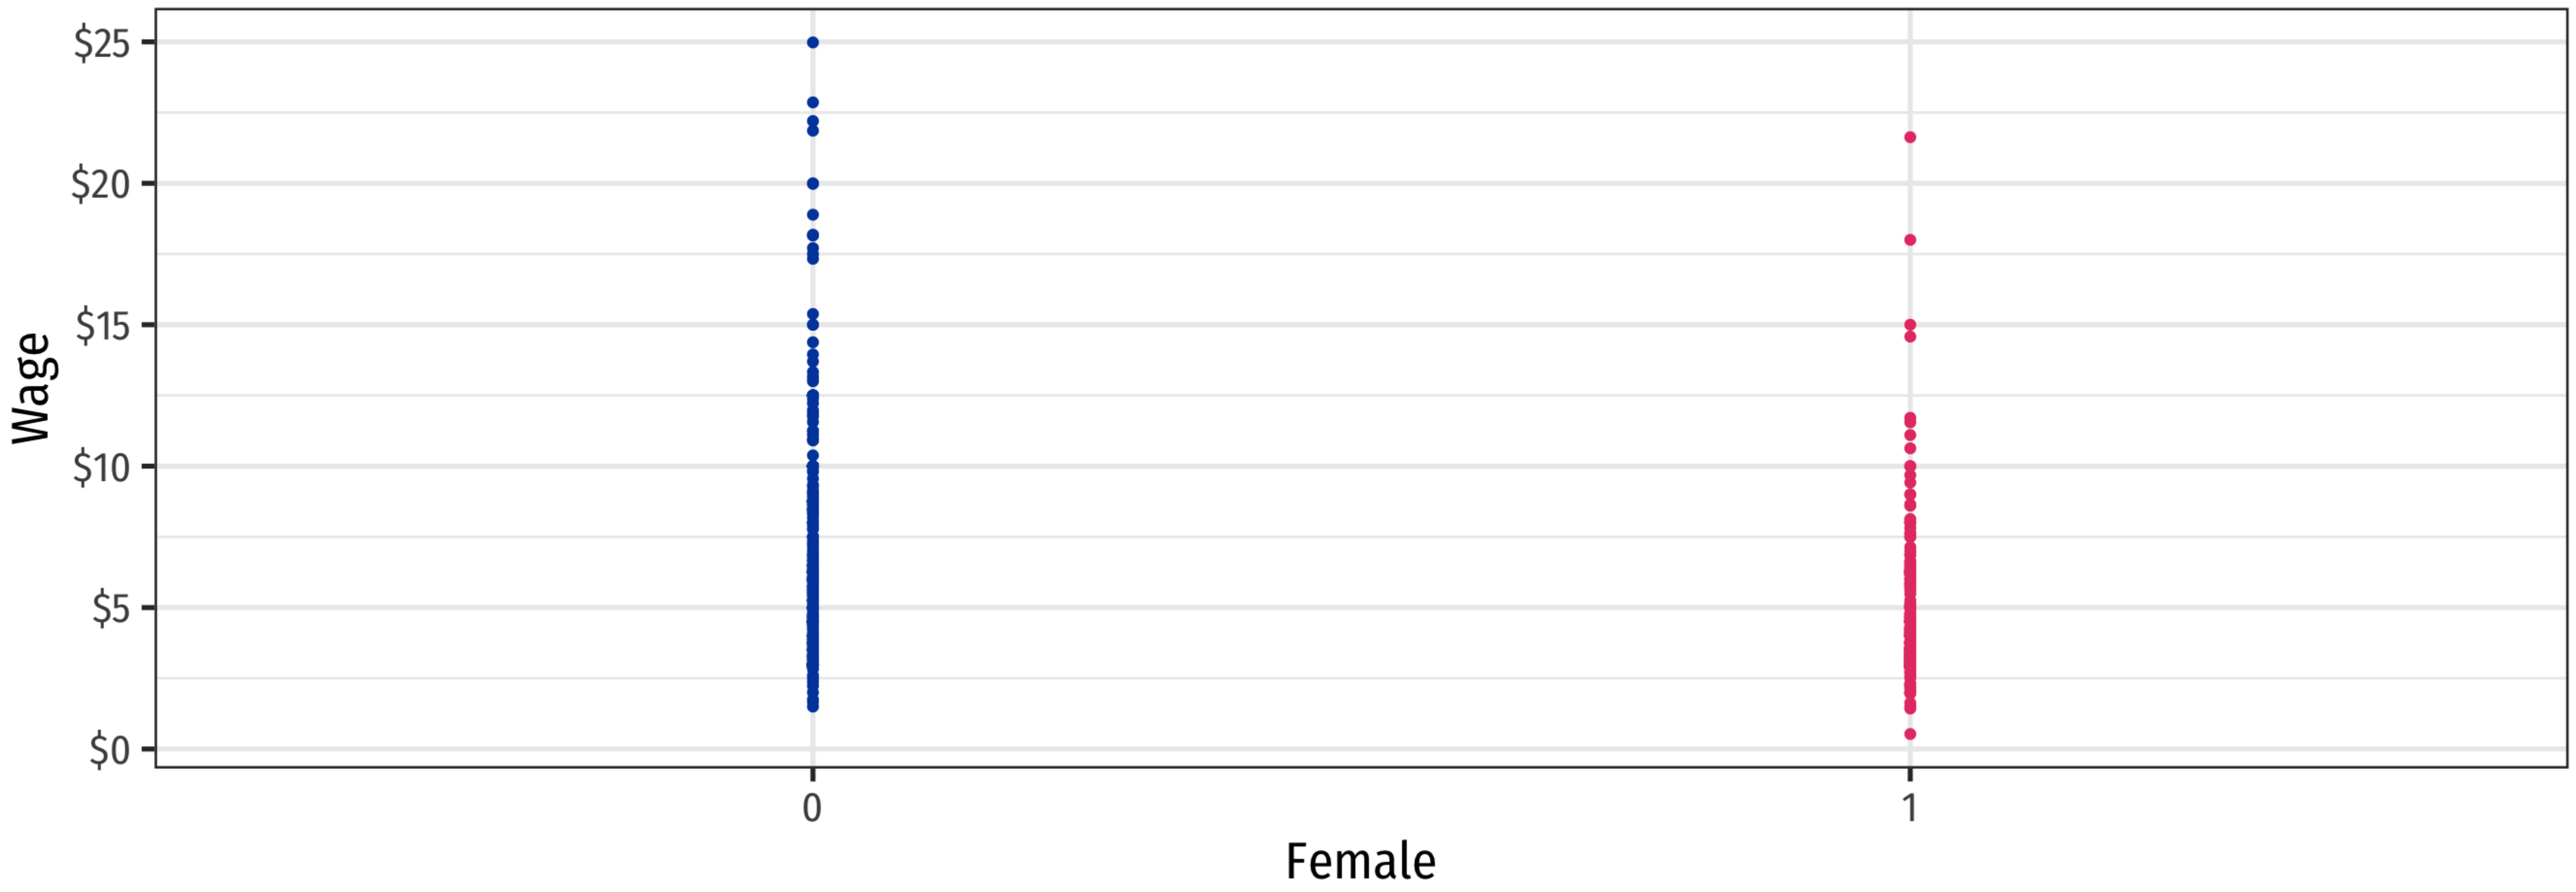

- Hard to see relationships because of **overplotting** . . .

Comparing Groups in Regression: Scatterplot

Plot | Code

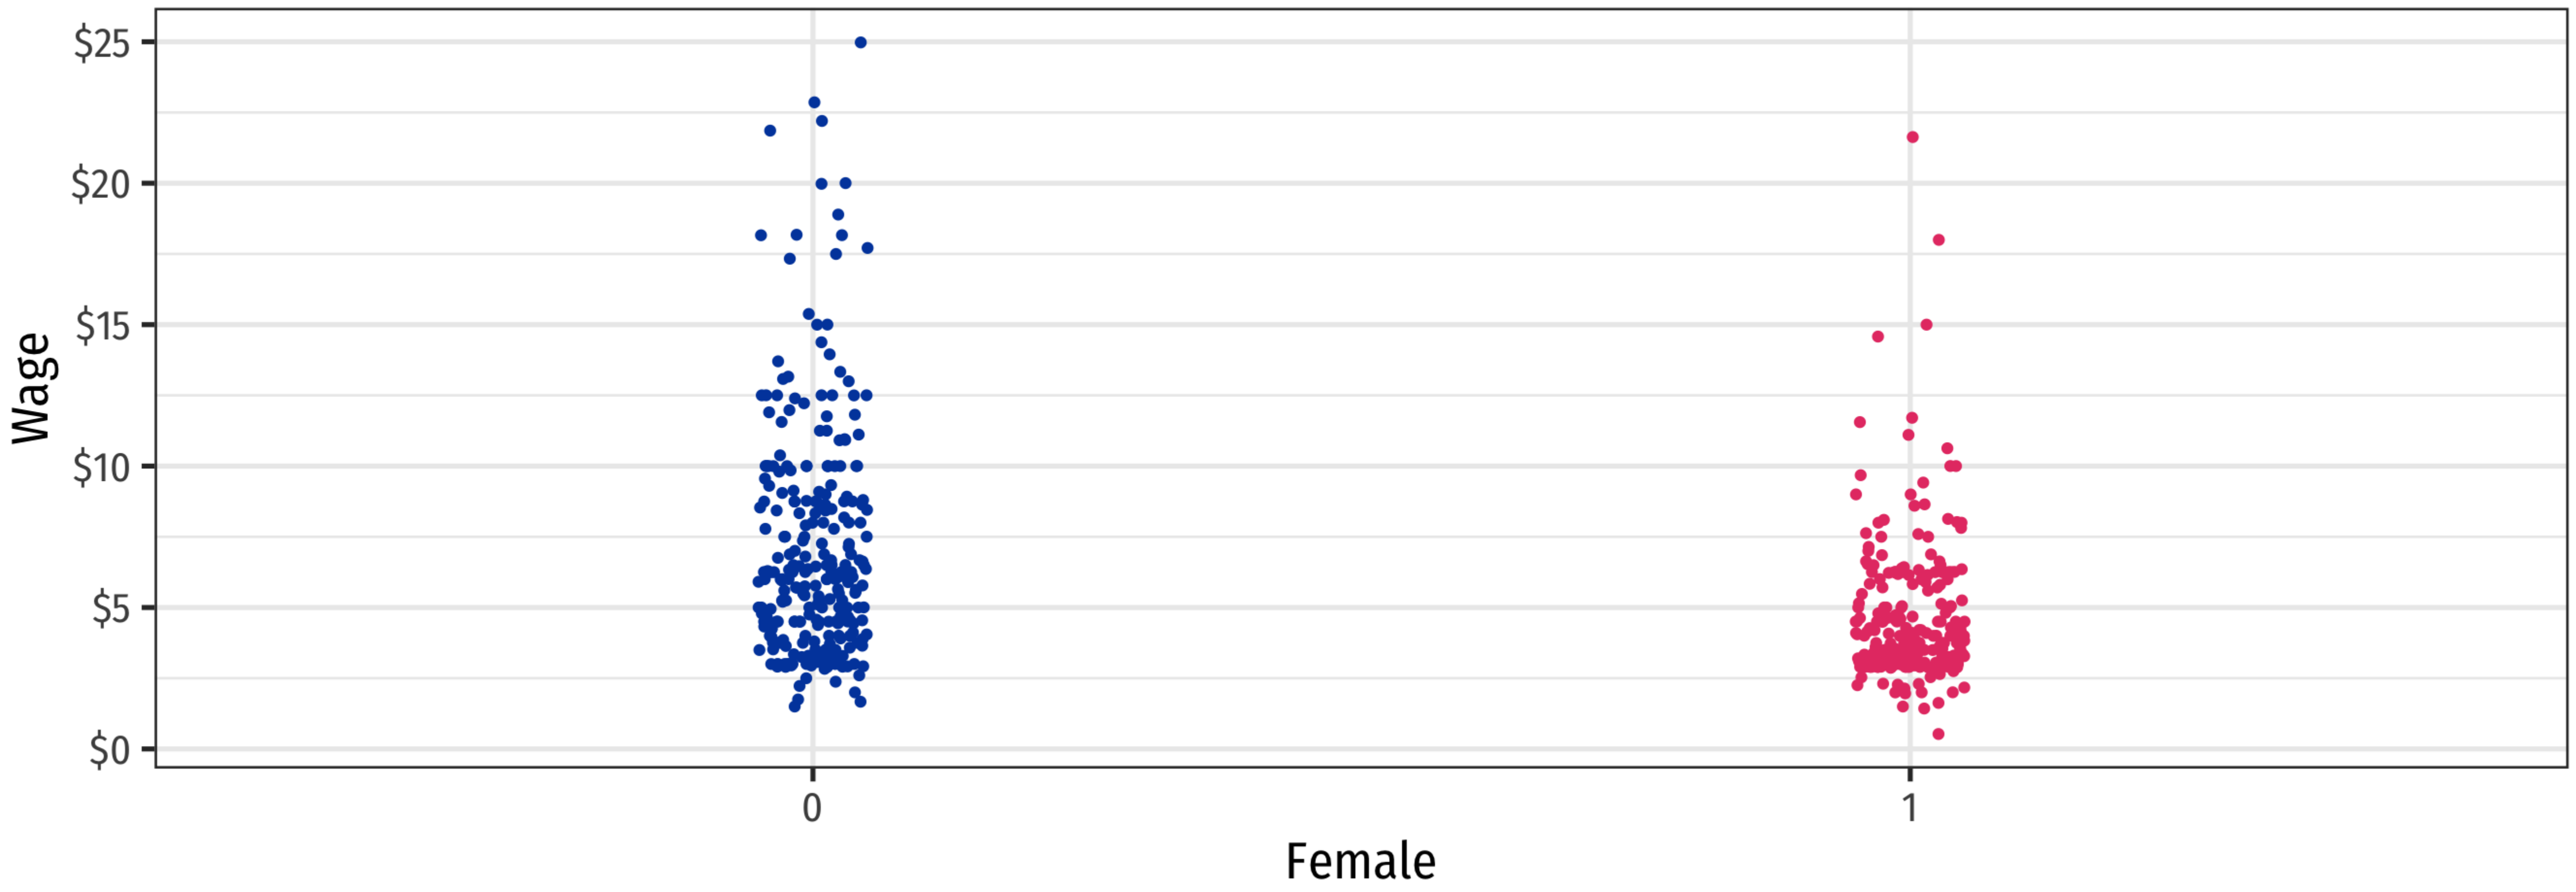

- Tip: use `geom_jitter()` instead of `geom_point()` to *randomly* nudge points!
  - Only used for *plotting*, does not affect actual data, regression, etc.

Dummy Variables as Group Means

$$\hat{Y}_i = \hat{\beta}_0 + \hat{\beta}_1 D_i \qquad \text{where } D_i = \{0, 1\}$$

- **When $D_i = 0$ ("Control group"):**
  - $\hat{Y}_i = \hat{\beta}_0$
  - $E[Y_i|D_i = 0] = \hat{\beta}_0 \iff$ the mean of Y when $D_i = 0$
- **When $D_i = 1$ ("Treatment group"):**
  - $\hat{Y}_i = \hat{\beta}_0 + \hat{\beta}_1 D_i$
  - $E[Y_i|D_i = 1] = \hat{\beta}_0 + \hat{\beta}_1 \iff$ the mean of Y when $D_i = 1$
- So the **difference** in group means:

$$\begin{aligned} &= E[Y_i|D_i = 1] - E[Y_i|D_i = 0] \\ &= (\hat{\beta}_0 + \hat{\beta}_1) - (\hat{\beta}_0) \\ &= \hat{\beta}_1 \end{aligned}$$

Dummy Variables as Group Means: Our Example

Example

$$\widehat{Wage_i} = \hat{\beta}_0 + \hat{\beta}_1\, Female_i$$

- Mean wage for men:

$$E[Wage|Female = 0] = \hat{\beta}_0$$

- Mean wage for women:

$$E[Wage|Female = 1] = \hat{\beta}_0 + \hat{\beta}_1$$

- Difference in wage between men & women:

$$\hat{\beta}_1$$

Comparing Groups in Regression: Scatterplot

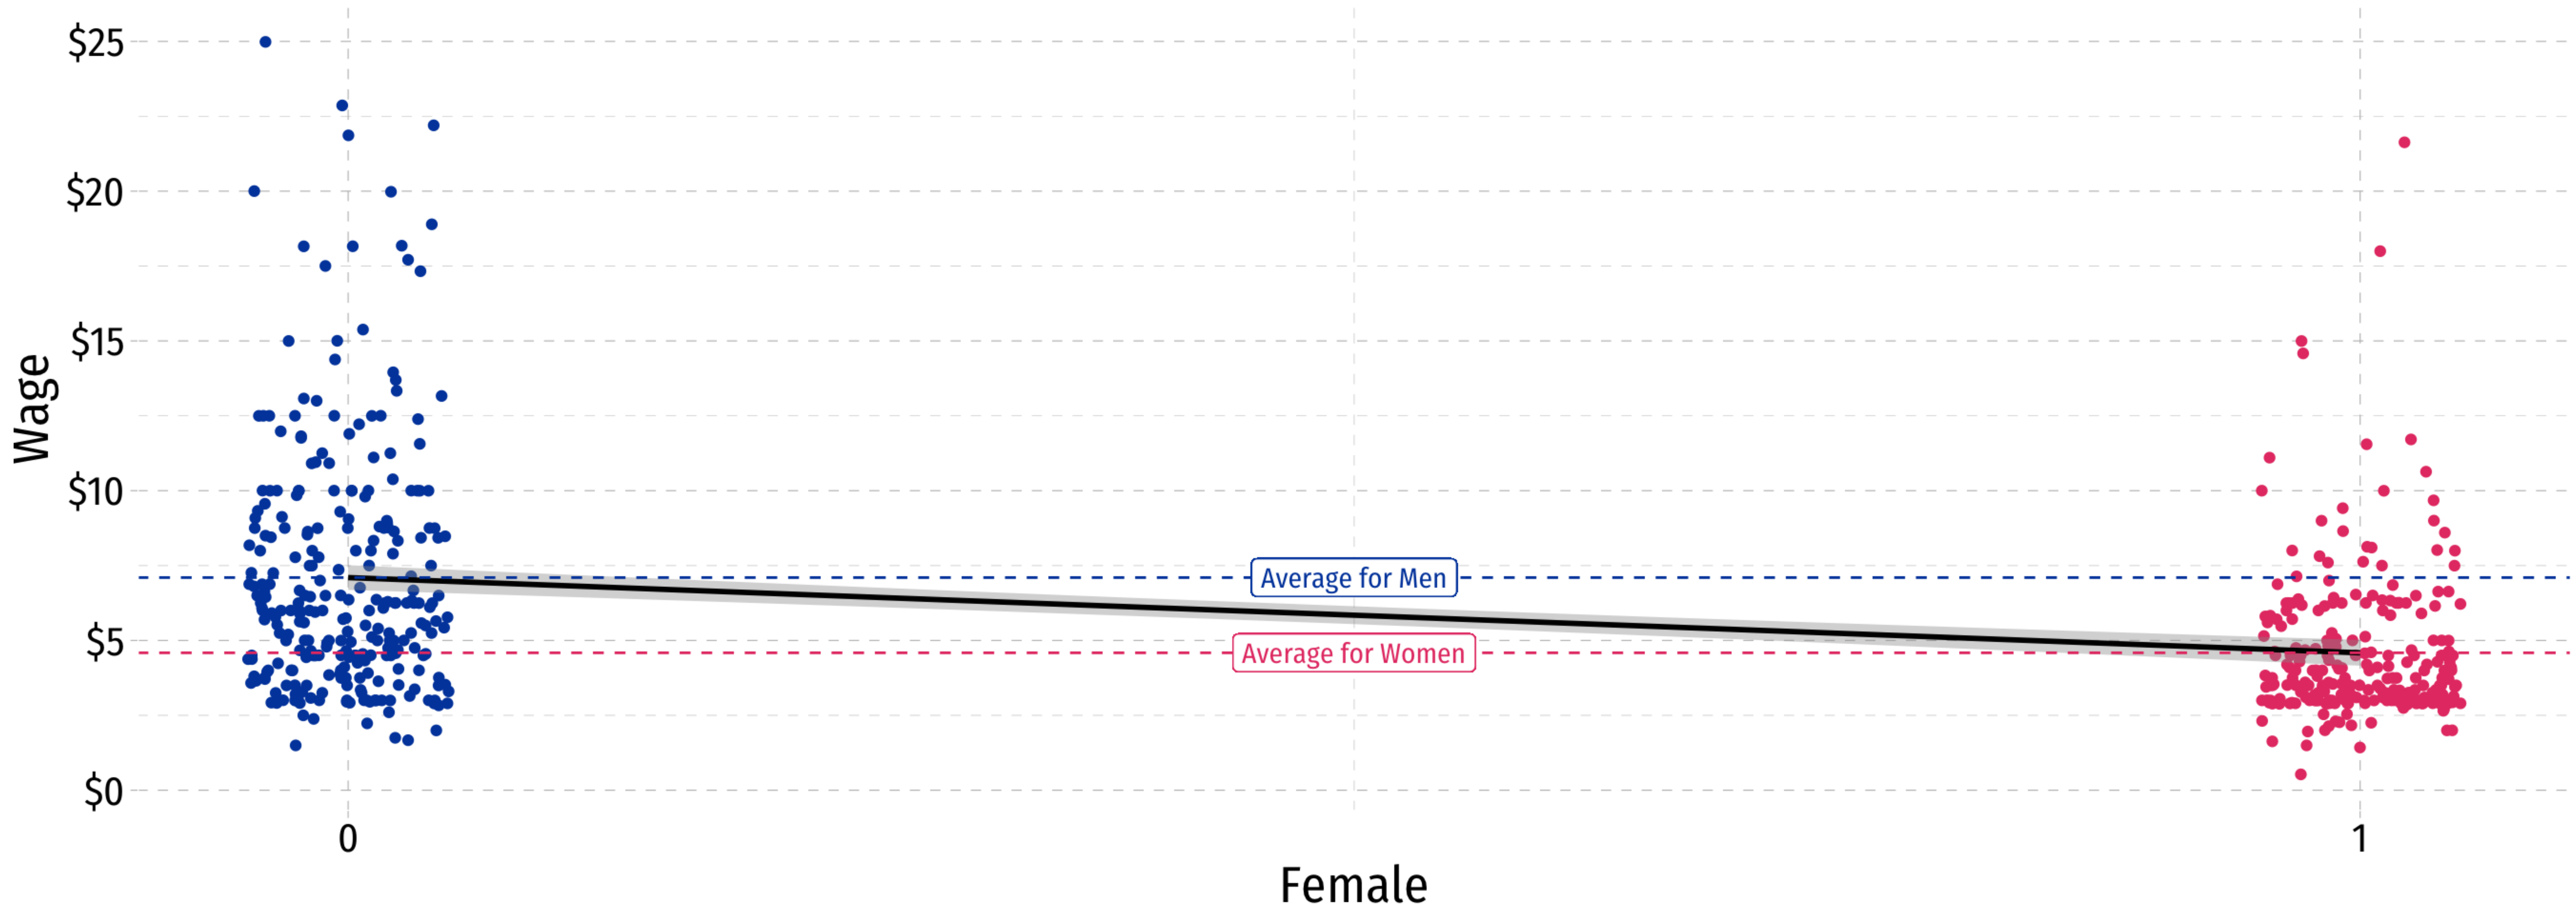

$$\widehat{Wage_i} = \hat{\beta}_0 + \hat{\beta}_1 \, Female_i$$

The Data

	wage <dbl>	gender <fct>	educ <int>	exper <int>	tenure <int>
1	3.10	Female	11	2	0
2	3.24	Female	12	22	2
3	3.00	Male	11	2	0
4	6.00	Male	8	44	28
5	5.30	Male	12	7	2
6	8.75	Male	16	9	8
7	11.25	Male	18	15	7
8	5.00	Female	12	5	3
9	3.60	Female	12	26	4
10	18.18	Male	17	22	21

1-10 of 526 rows | 1-6 of 25 columns

Previous **1** 2 3 4 5 6 ... 53 Next

Conditional Group Means

```
# Summarize for Men

wages %>%
  filter(gender=="Male") %>%
  summarize(mean = mean(wage),
            sd = sd(wage))
```

mean <dbl>
7.099489

1 row | 1-1 of 2 columns

```
# Summarize for Women

wages %>%
  filter(gender=="Female") %>%
  summarize(mean = mean(wage),
            sd = sd(wage))
```

mean <dbl>
4.587659

1 row | 1-1 of 2 columns

Visualize Differences

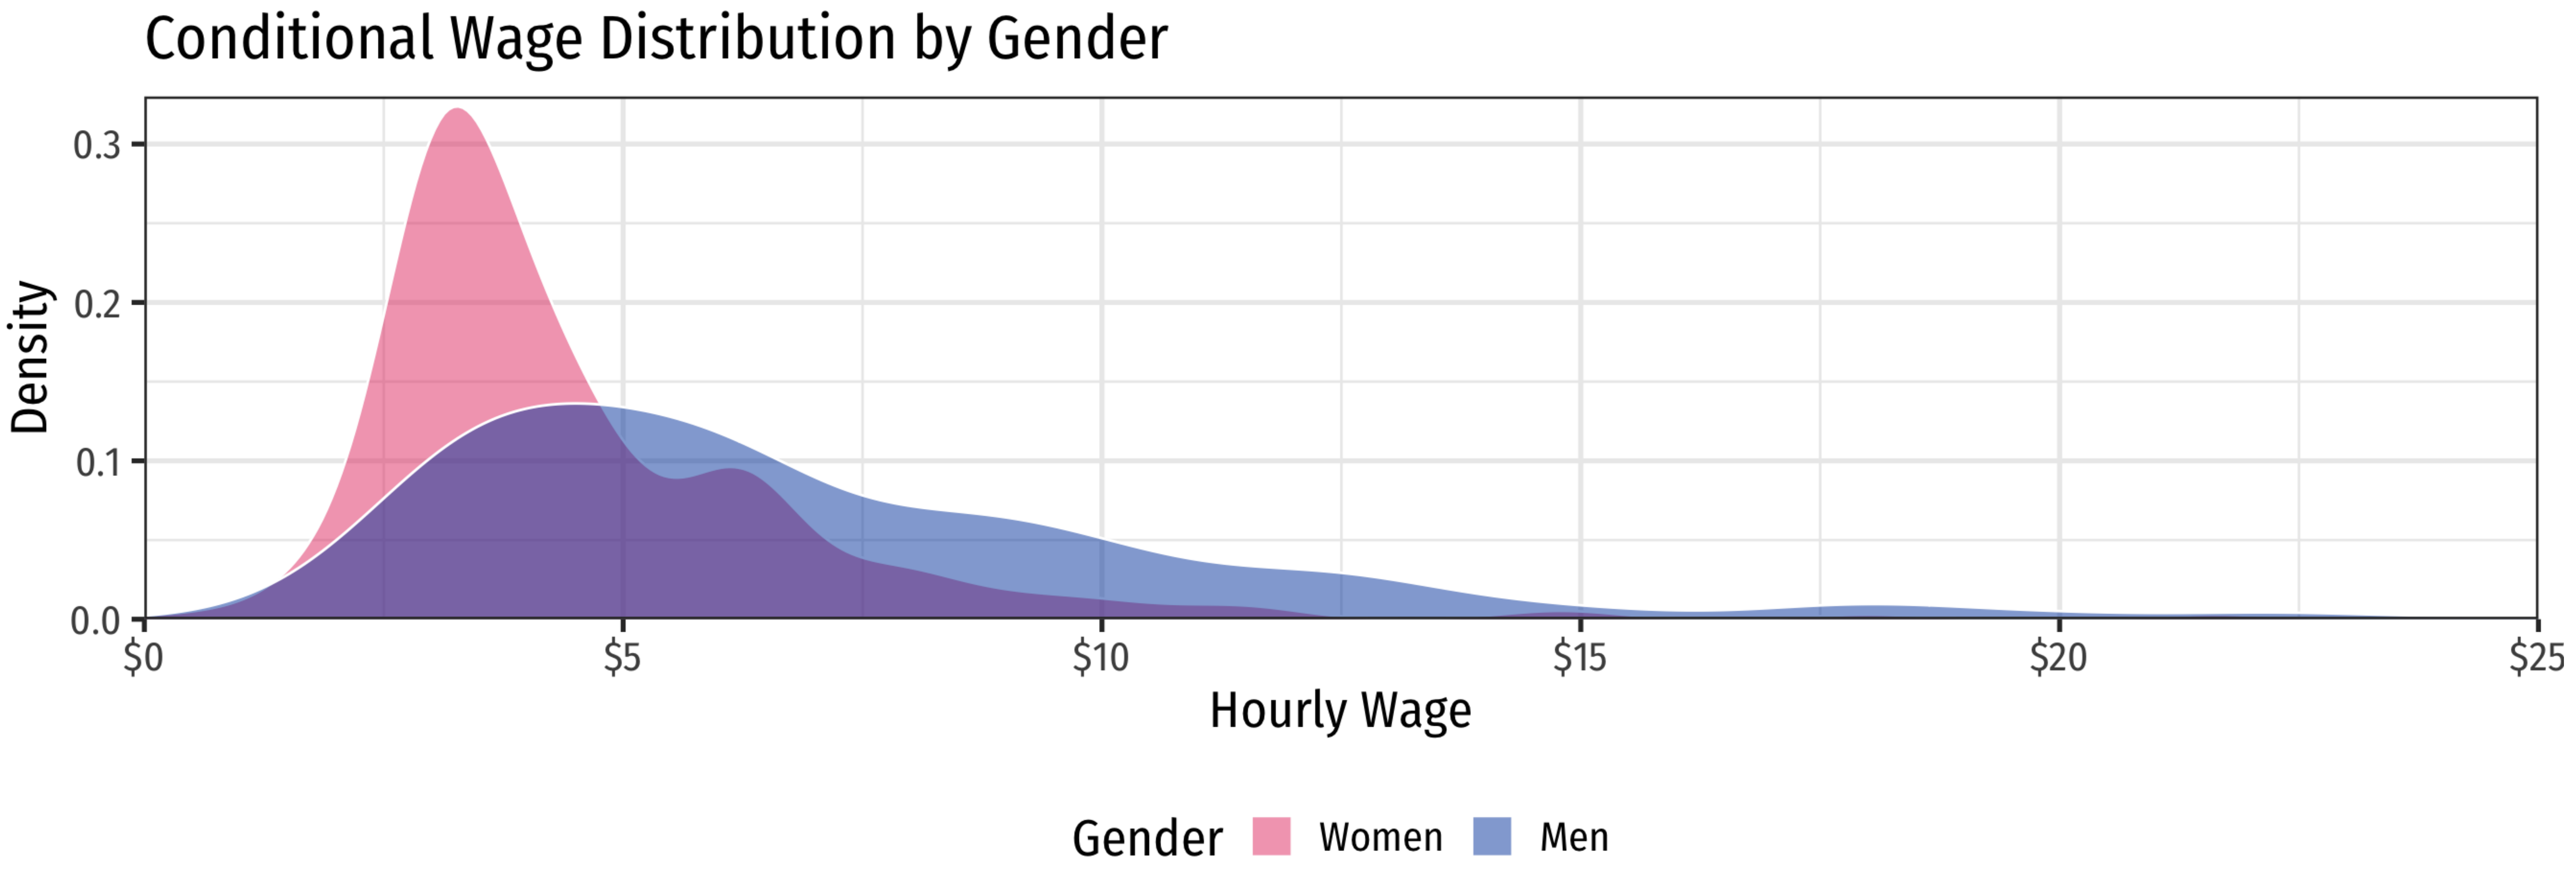

The Regression (`factor` variables)

```r
reg <- lm(wage ~ gender, data = wages)
summary(reg)
```

```
Call:
lm(formula = wage ~ gender, data = wages)

Residuals:
    Min      1Q  Median      3Q     Max 
-5.5995 -1.8495 -0.9877  1.4260 17.8805 

Coefficients:
            Estimate Std. Error t value Pr(>|t|)    
(Intercept)   4.5877     0.2190  20.950  < 2e-16 ***
genderMale    2.5118     0.3034   8.279 1.04e-15 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 3.476 on 524 degrees of freedom
Multiple R-squared:  0.1157,	Adjusted R-squared:  0.114 
F-statistic: 68.54 on 1 and 524 DF,  p-value: 1.042e-15
```

```r
library(broom)
tidy(reg)
```

term <chr>
(Intercept)
genderMale

2 rows | 1-1 of 5 columns

- Putting the `factor` variable `gender` in, `R` automatically chooses a value to set as `TRUE`, in this case `Male = TRUE`
  - `genderMALE` $= 1$ for Male, $= 0$ for Female
- According to the data, men earn, on average, $2.51 more than women

The Regression: Dummy Variables

- Let's explicitly make `gender` into a dummy variable for `female`:

```r
# add a female dummy variable
wages <- wages %>%
  mutate(female = ifelse(test = gender == "Female",
                         yes = 1,
                         no = 0))
```

```r
wages
```

	wage <dbl>	female <dbl>	educ <int>	exper <int>	tenure <int>
1	3.10	1	11	2	0
2	3.24	1	12	22	2
3	3.00	0	11	2	0
4	6.00	0	8	44	28
5	5.30	0	12	7	2
6	8.75	0	16	9	8
7	11.25	0	18	15	7
8	5.00	1	12	5	3
9	3.60	1	12	26	4
10	18.18	0	17	22	21

1-10 of 526 rows | 1-6 of 26 columns

Previous **1** 2 3 4 5 6 ... 53 Next

The Regression (Dummy variables)

```r
female_reg <- lm(wage ~ female, data = wages)
summary(female_reg)
```

```
Call:
lm(formula = wage ~ female, data = wages)

Residuals:
    Min      1Q  Median      3Q     Max 
-5.5995 -1.8495 -0.9877  1.4260 17.8805 

Coefficients:
            Estimate Std. Error t value Pr(>|t|)    
(Intercept)   7.0995     0.2100  33.806  < 2e-16 ***
female       -2.5118     0.3034  -8.279 1.04e-15 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 3.476 on 524 degrees of freedom
Multiple R-squared:  0.1157,	Adjusted R-squared:  0.114 
F-statistic: 68.54 on 1 and 524 DF,  p-value: 1.042e-15
```

```r
library(broom)
tidy(female_reg)
```

term <chr>
(Intercept)
female

2 rows | 1-1 of 5 columns

Dummy Regression vs. Group Means

From tabulation of group means

Gender	**Avg. Wage**	**Std. Dev.**	n
Female	4.59	2.33	252
Male	7.10	4.16	274
Difference	2.51	0.30	$-$

From t-test of difference in group means

term
<chr>
(Intercept)
female

2 rows | 1-1 of 5 columns

$$\widehat{\text{Wages}}_i = 7.10 - 2.51\,\text{Female}_i$$

Recoding Dummy Variables

Recoding Dummy Variables

Example

Suppose instead of `female` we had used:

$$\widehat{Wage}_i = \hat{\beta}_0 + \hat{\beta}_1\, Male_i \qquad \text{where } Male_i = \begin{cases} 1 & \text{if person } i \text{ is } Male \\ 0 & \text{if person } i \text{ is } Female \end{cases}$$

Recoding Dummies in the Data

```r
wages <- wages %>%
  mutate(male = ifelse(female == 0, # condition: is female equal to 0?
                       yes = 1, # if true: code as "1"
                       no = 0)) # if false: code as "0"

# verify it worked
wages %>%
  select(wage, female, male) %>%
  head(n = 5)
```

	wage	female	male
	<dbl>	<dbl>	<dbl>
1	3.10	1	0
2	3.24	1	0
3	3.00	0	1
4	6.00	0	1
5	5.30	0	1

5 rows

Scatterplot with `Male`

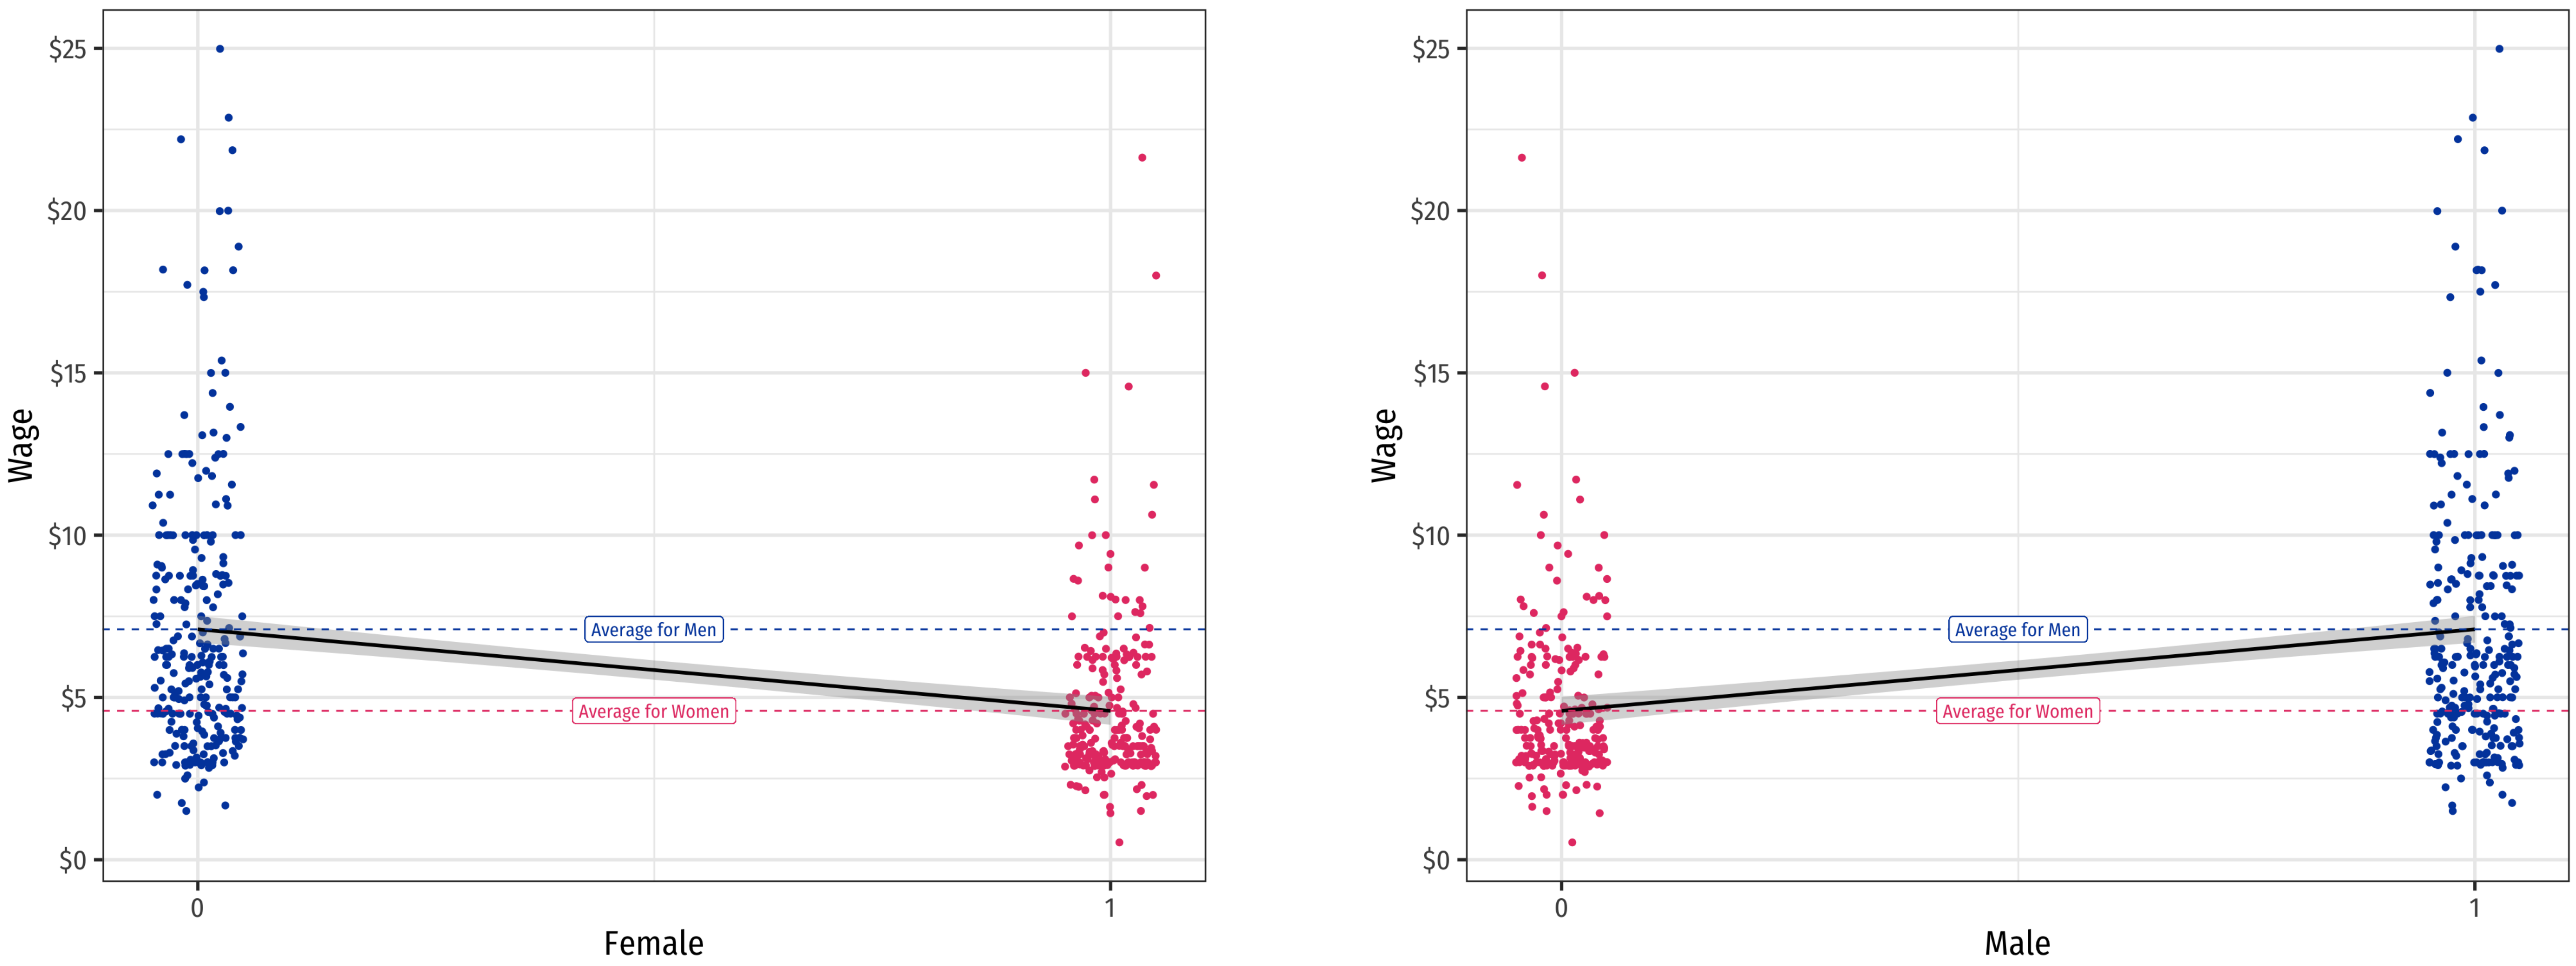

Dummy Variables as Group Means: With Male

Example

$$\widehat{Wage}_i = \hat{\beta}_0 + \hat{\beta}_1 \, Male_i$$

- Mean wage for men:

$$E[Wage|Male = 1] = \hat{\beta}_0 + \hat{\beta}_1$$

- Mean wage for women:

$$E[Wage|Male = 0] = \hat{\beta}_0$$

- Difference in wage between men & women:

$$\hat{\beta}_1$$

Scatterplot & Regression Line with `Male`

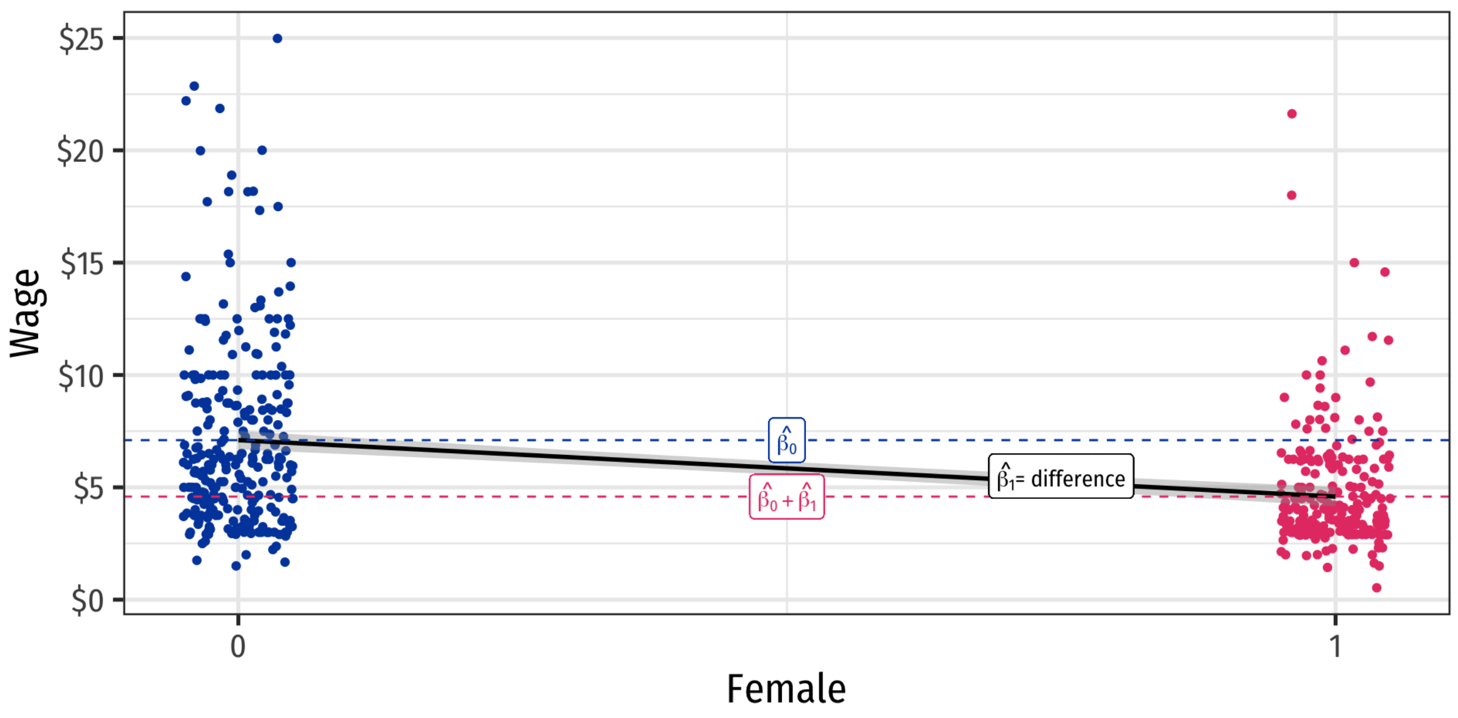

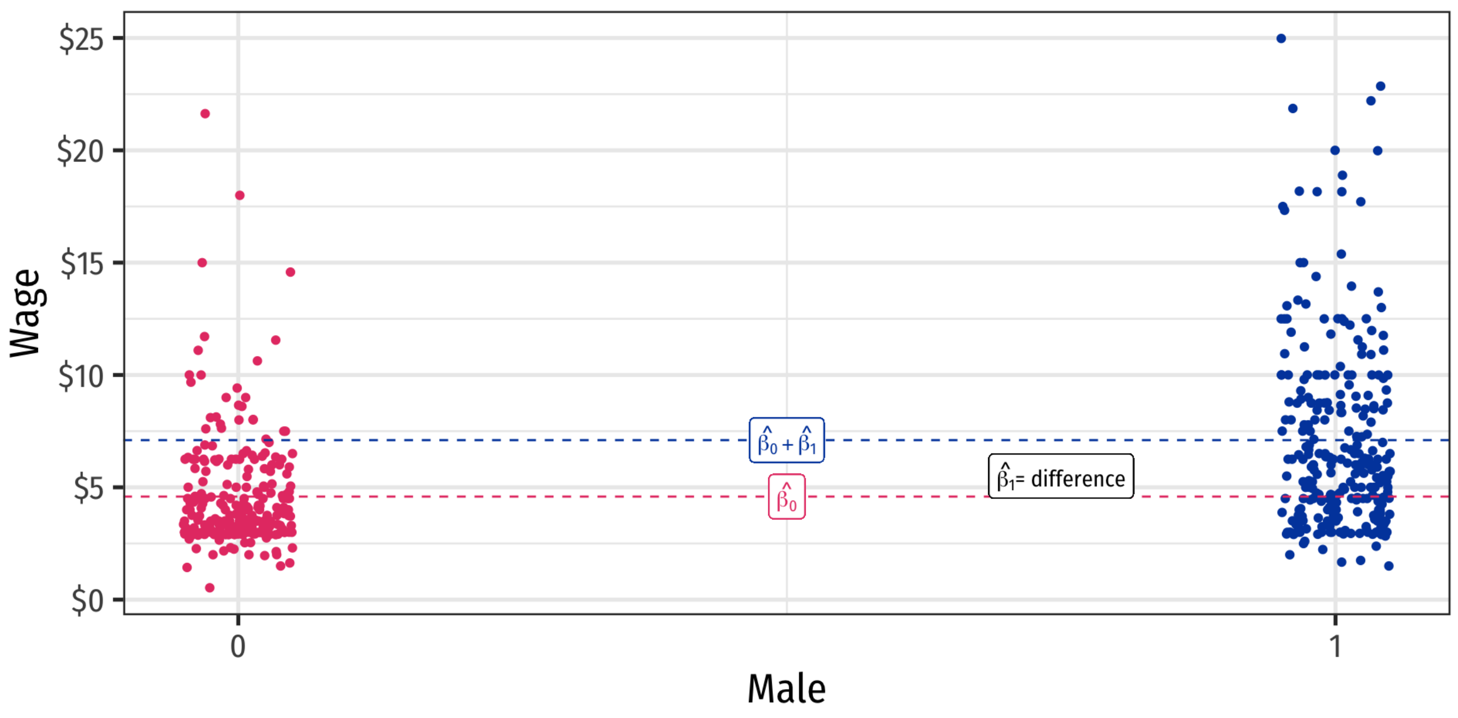

The Regression with `Male`

```r
male_reg <- lm(wage ~ male, data = wages)
summary(male_reg)
```

```

Call:
lm(formula = wage ~ male, data = wages)

Residuals:
    Min      1Q  Median      3Q     Max 
-5.5995 -1.8495 -0.9877  1.4260 17.8805 

Coefficients:
            Estimate Std. Error t value Pr(>|t|)    
(Intercept)   4.5877     0.2190  20.950  < 2e-16 ***
male          2.5118     0.3034   8.279 1.04e-15 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.'
```

```r
library(broom)
tidy(male_reg)
```

term <chr>
(Intercept)
male

2 rows | 1-1 of 5 columns

The Dummy Regression: Male or Female

	Wage	Wage
Constant	7.10***	4.59***
	(0.21)	(0.22)
female	−2.51***	
	(0.30)	
male		2.51***
		(0.30)
n	526	526
Adj. R^2	0.11	0.11
SER	3.47	3.47

* p < 0.1, ** p < 0.05, *** p < 0.01

- Note it doesn't matter if we use `male` or `female`, difference is always $2.51
- Compare the constant (average for the $D=0$ group)
- Should you use `male` AND `female` in a regression? We'll come to that...

Categorical Variables (More than 2 Categories)

Categorical Variables with More than 2 Categories

- A **categorical variable** expresses membership in a category, where there is no ranking or hierarchy of the categories
  - We've looked at categorical variables with 2 categories only
  - e.g. Male/Female, Spring/Summer/Fall/Winter, Democratic/Republican/Independent
- Might be an **ordinal variable** expresses rank or an ordering of data, but not necessarily their relative magnitude
  - e.g. Order of finalists in a competition (1st, 2nd, 3rd)
  - e.g. Highest education attained (1=elementary school, 2=high school, 3=bachelor's degree, 4=graduate degree)
  - in R, an `ordered factor`

Using Categorical Variables in Regression I

Example

How do wages vary by region of the country? Let $Region_i = \{Northeast, Midwest, South, West\}$

- Can we run the following regression?

$$\widehat{Wages_i} = \hat{\beta}_0 + \hat{\beta}_1 \, Region_i$$

Using Categorical Variables in Regression II

Example

How do wages vary by region of the country? Let $Region_i = \{Northeast, Midwest, South, West\}$

- Code region numerically:

$$Region_i = \begin{cases} 1 & \text{if } i \text{ is in } Northeast \\ 2 & \text{if } i \text{ is in } Midwest \\ 3 & \text{if } i \text{ is in } South \\ 4 & \text{if } i \text{ is in } West \end{cases}$$

- Can we run the following regression?

$$\widehat{Wages_i} = \hat{\beta_0} + \hat{\beta_1}\, Region_i$$

Using Categorical Variables in Regression III

Example

How do wages vary by region of the country? Let $Region_i = \{Northeast, \; Midwest, \; South, \; West\}$

- Create a dummy variable for *each* region:
  - $Northeast_i = 1$ if i is in Northeast, otherwise $= 0$
  - $Midwest_i = 1$ if i is in Midwest, otherwise $= 0$
  - $South_i = 1$ if i is in South, otherwise $= 0$
  - $West_i = 1$ if i is in West, otherwise $= 0$
- Can we run the following regression?

$$\widehat{Wages_i} = \hat{\beta}_0 + \hat{\beta}_1 Northeast_i + \hat{\beta}_2 Midwest_i + \hat{\beta}_3 South_i + \hat{\beta}_4 West_i$$

The Dummy Variable Trap

Example

$$\widehat{Wages}_i = \hat{\beta}_0 + \hat{\beta}_1 Northeast_i + \hat{\beta}_2 Midwest_i + \hat{\beta}_3 South_i + \hat{\beta}_4 West_i$$

- If we include *all* possible categories, they are **perfectly multicollinear**, an exact linear function of one another:

$$Northeast_i + Midwest_i + South_i + West_i = 1 \quad \forall i$$

- This is known as the **dummy variable trap**, a common source of perfect multicollinearity

The Reference Category

- To avoid the dummy variable trap, always omit one category from the regression, known as the **"reference category"**
- It does not matter which category we omit!
- **Coefficients on each dummy variable measure the *difference* between the *reference* category and each category dummy**

The Reference Category: Example

Example

$$\widehat{Wages_i} = \hat{\beta}_0 + \hat{\beta}_1 Northeast_i + \hat{\beta}_2 Midwest_i + \hat{\beta}_3 South_i$$ :::

- $West_i$ is omitted (arbitrarily chosen)
- $\hat{\beta}_0$: average wage for i in the West
- $\hat{\beta}_1$: difference between West and Northeast
- $\hat{\beta}_2$: difference between West and Midwest
- $\hat{\beta}_3$: difference between West and South

Regression in R with Categorical Variable

```r
lm(wage ~ region, data = wages) %>% summary()
```

```
Call:
lm(formula = wage ~ region, data = wages)

Residuals:
   Min     1Q Median     3Q    Max 
-6.083 -2.387 -1.097  1.157 18.610 

Coefficients:
                Estimate Std. Error t value Pr(>|t|)    
(Intercept)       5.7105     0.3195  17.871   <2e-16 ***
regionNortheast   0.6593     0.4651   1.418   0.1569    
regionSouth      -0.3236     0.4173  -0.775   0.4385    
regionWest        0.9029     0.5035   1.793   0.0735 .  
```

Regression in R with Dummies (& Dummy Variable Trap)

```r
lm(wage ~ northeast + midwest + south + west, data = wages) %>% summary()
```

```
Call:
lm(formula = wage ~ northeast + midwest + south + west, data = wages)

Residuals:
   Min     1Q Median     3Q    Max 
-6.083 -2.387 -1.097  1.157 18.610 

Coefficients: (1 not defined because of singularities)
            Estimate Std. Error t value Pr(>|t|)    
(Intercept)   6.6134     0.3891  16.995  < 2e-16 ***
northeast    -0.2436     0.5154  -0.473  0.63664    
midwest      -0.9029     0.5035  -1.793  0.07352 .  
south        -1.2265     0.4728  -2.594  0.00974 ** 
```

- R automatically drops one category to avoid perfect multicollinearity

Using Different Reference Categories in R

	Wage	Wage	Wage	Wage
Constant	6.37***	5.71***	5.39***	6.61***
	(0.34)	(0.32)	(0.27)	(0.39)
northcen	−0.66		0.32	−0.90*
	(0.47)		(0.42)	(0.50)
south	−0.98**	−0.32		−1.23***
	(0.43)	(0.42)		(0.47)
west	0.24	0.90*	1.23***	
	(0.52)	(0.50)	(0.47)	
northeast		0.66	0.98**	−0.24
		(0.47)	(0.43)	(0.52)
n	526	526	526	526
R^2	0.02	0.02	0.02	0.02
Adj. R^2	0.01	0.01	0.01	0.01
SER	3.66	3.66	3.66	3.66

* p < 0.1, ** p < 0.05, *** p < 0.01

- Constant is alsways average wage for reference (omitted) region
- Compare coefficients between Midwest in (1) and Northeast in (2)...
- Compare coefficients between West in (3) and South in (4)...
- Does not matter which region we omit!
  - Same R^2, SER, coefficients give same results

Dummy *Dependent* (Y) Variables

- In many contexts, we will want to have our *dependent* (Y) variable be a dummy variable

Example

$$\widehat{\text{Admitted}}_i = \hat{\beta}_0 + \hat{\beta}_1\, GPA_i \quad \text{where } \text{Admitted}_i = \begin{cases} 1 & \text{if } i \text{ is Admitted} \\ 0 & \text{if } i \text{ is Not Admitted} \end{cases}$$

- A model where Y is a dummy is called a **linear probability model**, as it measures the **probability of Y occurring given the X's, i.e. $P(Y_i = 1|X_1, \cdots, X_k)$**
  - e.g. the probability person i is Admitted to a program with a given GPA
- Special models to properly interpret and extend this (**logistic "logit"**, **probit**, etc)
- Feel free to write papers with dummy Y variables!

Interaction Effects

Sliders and Switches

- Marginal effect of dummy variable: effect on Y of going from 0 to 1
- Marginal effect of continuous variable: effect on Y of a 1 unit change in X

Interaction Effects

- Sometimes one X variable might *interact* with another in determining Y

Example

Consider the gender pay gap again.

- *Gender* affects wages
- *Experience* affects wages
- **Does experience affect wages *differently* by gender?**
  - i.e. is there an interaction effect between gender and experience?
- **Note this is *NOT the same* as just asking: "do men earn more than women *with the same amount of experience?*"**

Three Types of Interactions

- Depending on the types of variables, there are 3 possible types of interaction effects
- We will look at each in turn

1. Interaction between a **dummy** and a **continuous** variable:

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 D_i + \beta_3 (X_i \times D_i)$$

2. Interaction between a **two dummy** variables:

$$Y_i = \beta_0 + \beta_1 D_{1i} + \beta_2 D_{2i} + \beta_3 (D_{1i} \times D_{2i})$$

3. Interaction between a **two continuous** variables:

$$Y_i = \beta_0 + \beta_1 X_{1i} + \beta_2 X_{2i} + \beta_3 (X_{1i} \times X_{2i})$$

Interactions Between a Dummy and Continuous Variable

Interactions: A Dummy & Continuous Variable

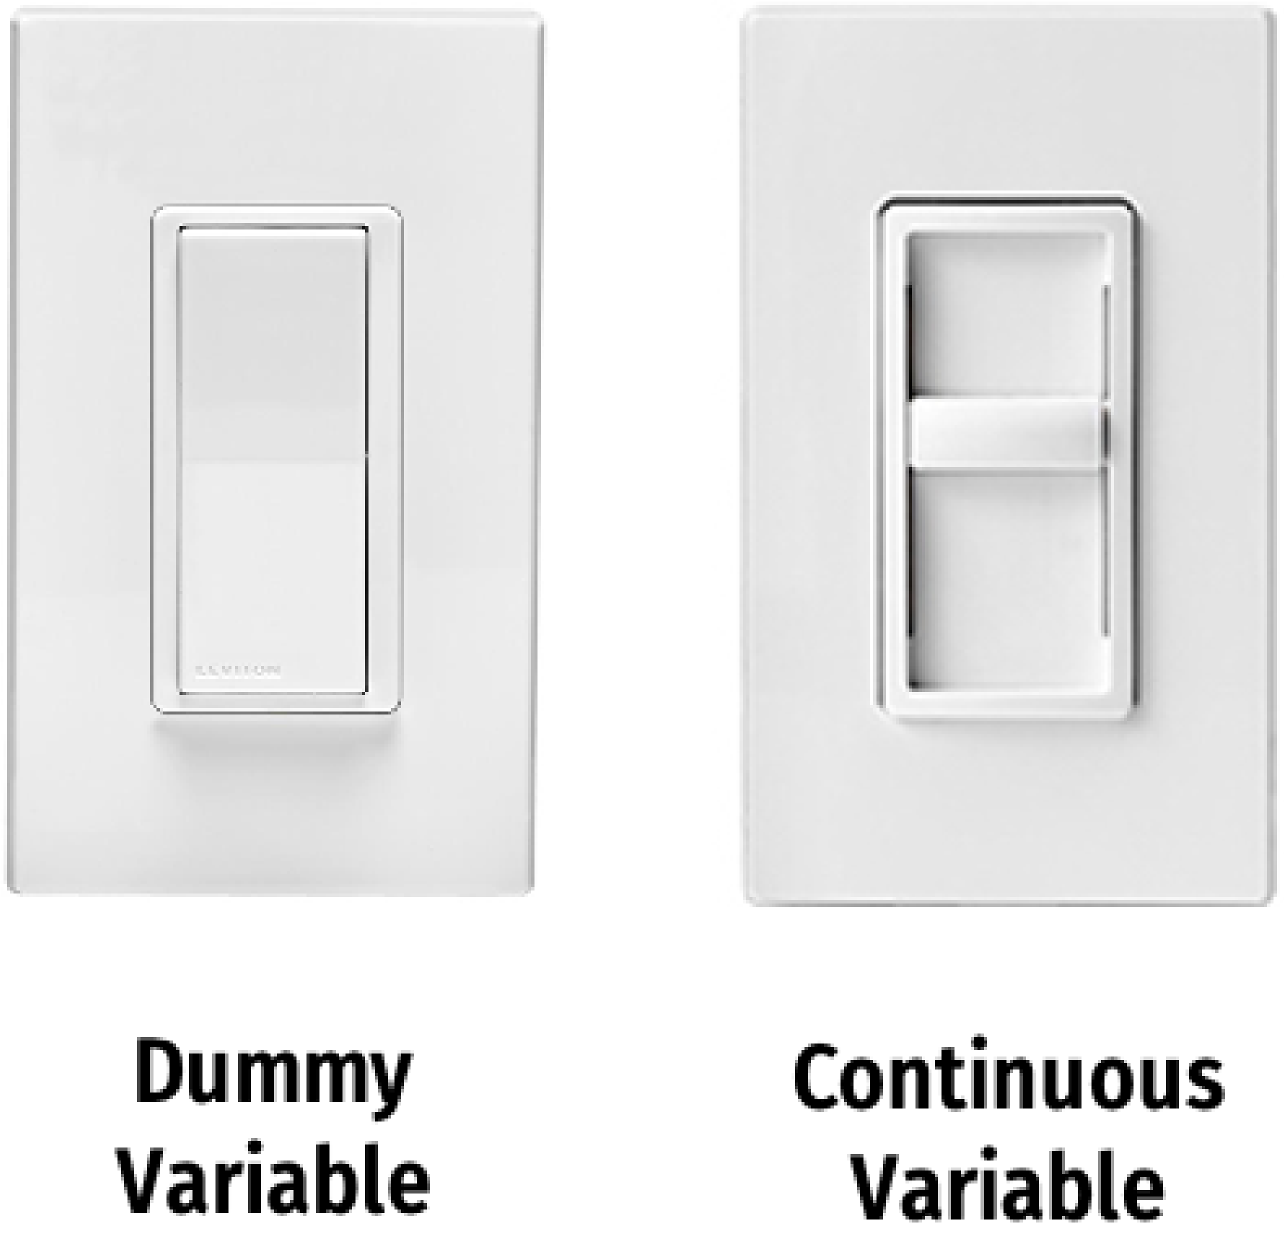

- Does the marginal effect of the continuous variable on Y change depending on whether the dummy is "on" or "off"?

Interactions: A Dummy & Continuous Variable I

- We can model an interaction by introducing a variable that is an .hi[interaction term] capturing the interaction between two variables:

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 D_i + \beta_3 (X_i \times D_i) \quad \text{where } D_i = \{0, 1\}$$

- β_3 estimates the **interaction effect** between X_i and D_i on Y_i
- What do the different coefficients (β)'s tell us?
  - Again, think logically by examining each group $(D_i = 0$ or $D_i = 1)$

Dummy-Continuous Interaction Effects as Two Regressions I

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 D_i + \beta_3 X_i \times D_i$$

- When $D_i = 0$ ("Control group"):

$$\begin{aligned} \hat{Y_i} &= \hat{\beta_0} + \hat{\beta_1} X_i + \hat{\beta_2}(0) + \hat{\beta_3} X_i \times (0) \\ \hat{Y_i} &= \hat{\beta_0} + \hat{\beta_1} X_i \end{aligned}$$

- When $D_i = 1$ ("Treatment group"):

$$\begin{aligned} \hat{Y_i} &= \hat{\beta_0} + \hat{\beta_1} X_i + \hat{\beta_2}(1) + \hat{\beta_3} X_i \times (1) \\ \hat{Y_i} &= (\hat{\beta_0} + \hat{\beta_2}) + (\hat{\beta_1} + \hat{\beta_3}) X_i \end{aligned}$$

- So what we really have is *two* regression lines!

Dummy-Continuous Interaction Effects as Two Regressions II

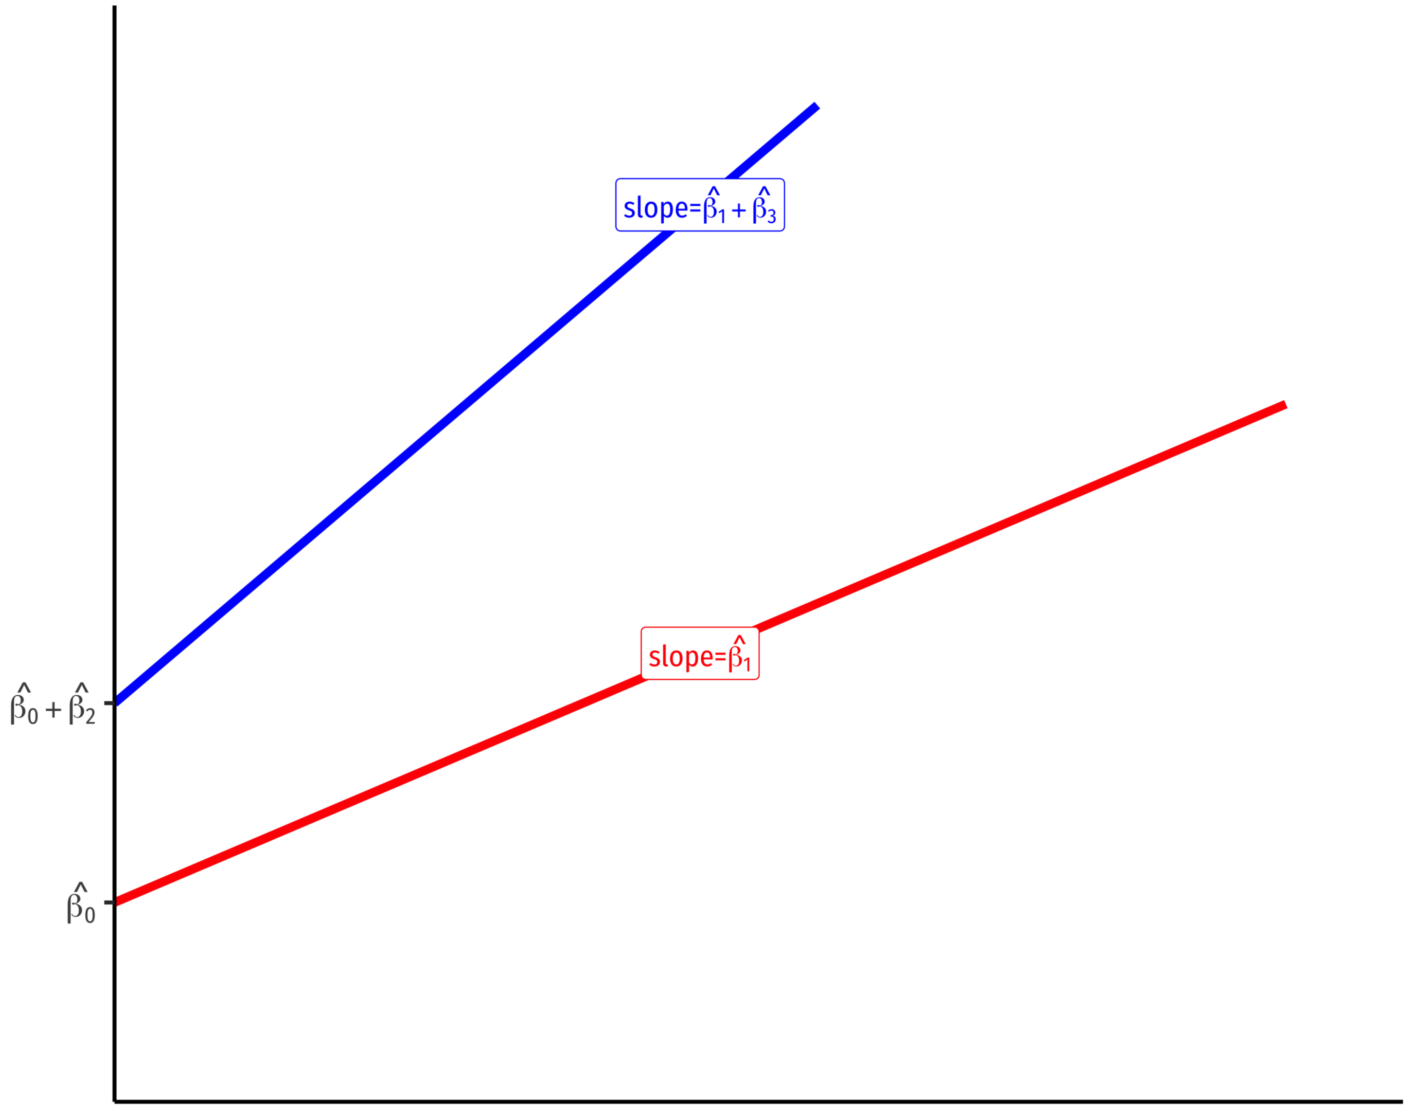

- $D_i = 0$ group:

$$Y_i = \hat{\beta}_0 + \hat{\beta}_1 X_i$$

- $D_i = 1$ group:

$$Y_i = (\hat{\beta}_0 + \hat{\beta}_2) + (\hat{\beta}_1 + \hat{\beta}_3) X_i$$

Interpretting Coefficients I

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 D_i + \beta_3 (X_i \times D_i)$$

- To interpret the coefficients, compare cases after changing X by ΔX:

$$Y_i + \Delta Y_i = \beta_0 + \beta_1 (X_i + \Delta X_i) \beta_2 D_i + \beta_3 \big((X_i + \Delta X_i) D_i \big)$$

- Subtracting these two equations, the difference is:

$$\Delta Y_i = \beta_1 \Delta X_i + \beta_3 D_i \Delta X_i$$
$$\frac{\Delta Y_i}{\Delta X_i} = \beta_1 + \beta_3 D_i$$

- **The effect of $X \rightarrow Y$ depends on the value of D_i!**
- **β_3: *increment* to the effect of $X \rightarrow Y$ when $D_i = 1$ (vs. $D_i = 0$)**

Interpretting Coefficients II

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 D_i + \beta_3 (X_i \times D_i)$$

- $\hat{\beta}_0$: $E[Y_i]$ for $X_i = 0$ and $D_i = 0$
- β_1: Marginal effect of $X_i \rightarrow Y_i$ for $D_i = 0$
- β_2: Marginal effect on Y_i of difference between $D_i = 0$ and $D_i = 1$
- β_3: The **difference** of the marginal effect of $X_i \rightarrow Y_i$ between $D_i = 0$ and $D_i = 1$
- This is a bit awkward, easier to think about the two regression lines:

Interpretting Coefficients III

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 D_i + \beta_3 (X_i \times D_i)$$

- For $D_i = 0$ Group: $\hat{Y}_i = \hat{\beta}_0 + \hat{\beta}_1 X_i$
  - Intercept: $\hat{\beta}_0$
  - Slope: $\hat{\beta}_1$
- For $D_i = 1$ Group: $\hat{Y}_i = (\hat{\beta}_0 + \hat{\beta}_2) + (\hat{\beta}_1 + \hat{\beta}_3) X_i$
  - Intercept: $\hat{\beta}_0 + \hat{\beta}_2$
  - Slope: $\hat{\beta}_1 + \hat{\beta}_3$

- $\hat{\beta}_2$: difference in intercept between groups
- $\hat{\beta}_3$: difference in slope between groups

- How can we determine if the two lines have the same slope and/or intercept?
  - Same intercept? t-test $H_0: \beta_2 = 0$
  - Same slope? t-test $H_0: \beta_3 = 0$

Interactions in Our Example

Example

$$\widehat{wage}_i = \hat{\beta}_0 + \hat{\beta}_1 \text{ experience}_i + \hat{\beta}_2 \text{ female}_i + \hat{\beta}_3 \text{ (experience}_i \times \text{female}_i)$$

- For men $female = 0$:

$$\widehat{wage_i} = \hat{\beta}_0 + \hat{\beta}_1 \text{ experience}_i$$

- For women $female = 1$:

$$\widehat{wage}_i = \underbrace{(\hat{\beta}_0 + \hat{\beta}_2)}_{\text{intercept}} + \underbrace{(\hat{\beta}_1 + \hat{\beta}_3)}_{\text{slope}} \text{experience}_i$$

Interactions in Our Example: Scatterplot

- Code

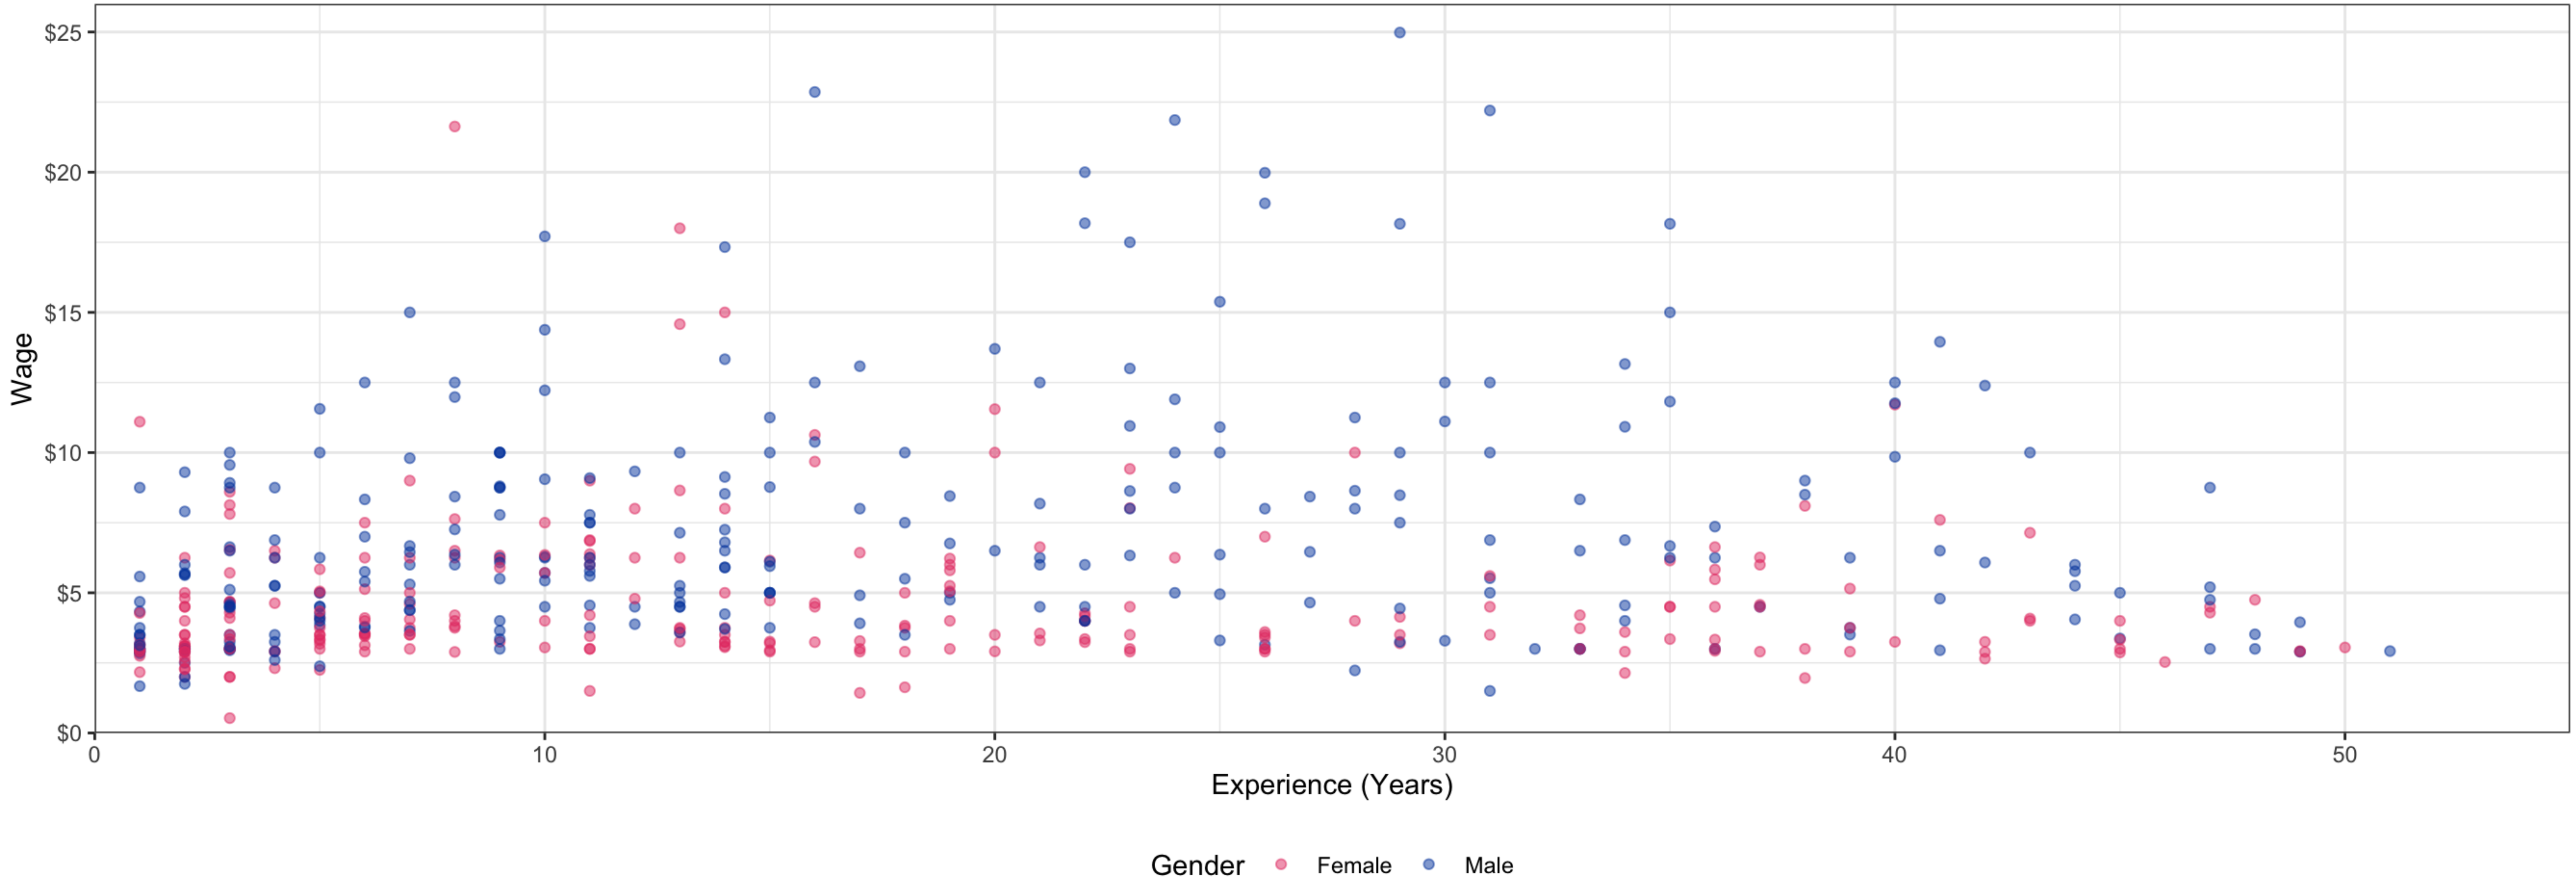

Interactions in Our Example: Scatterplot

- Code

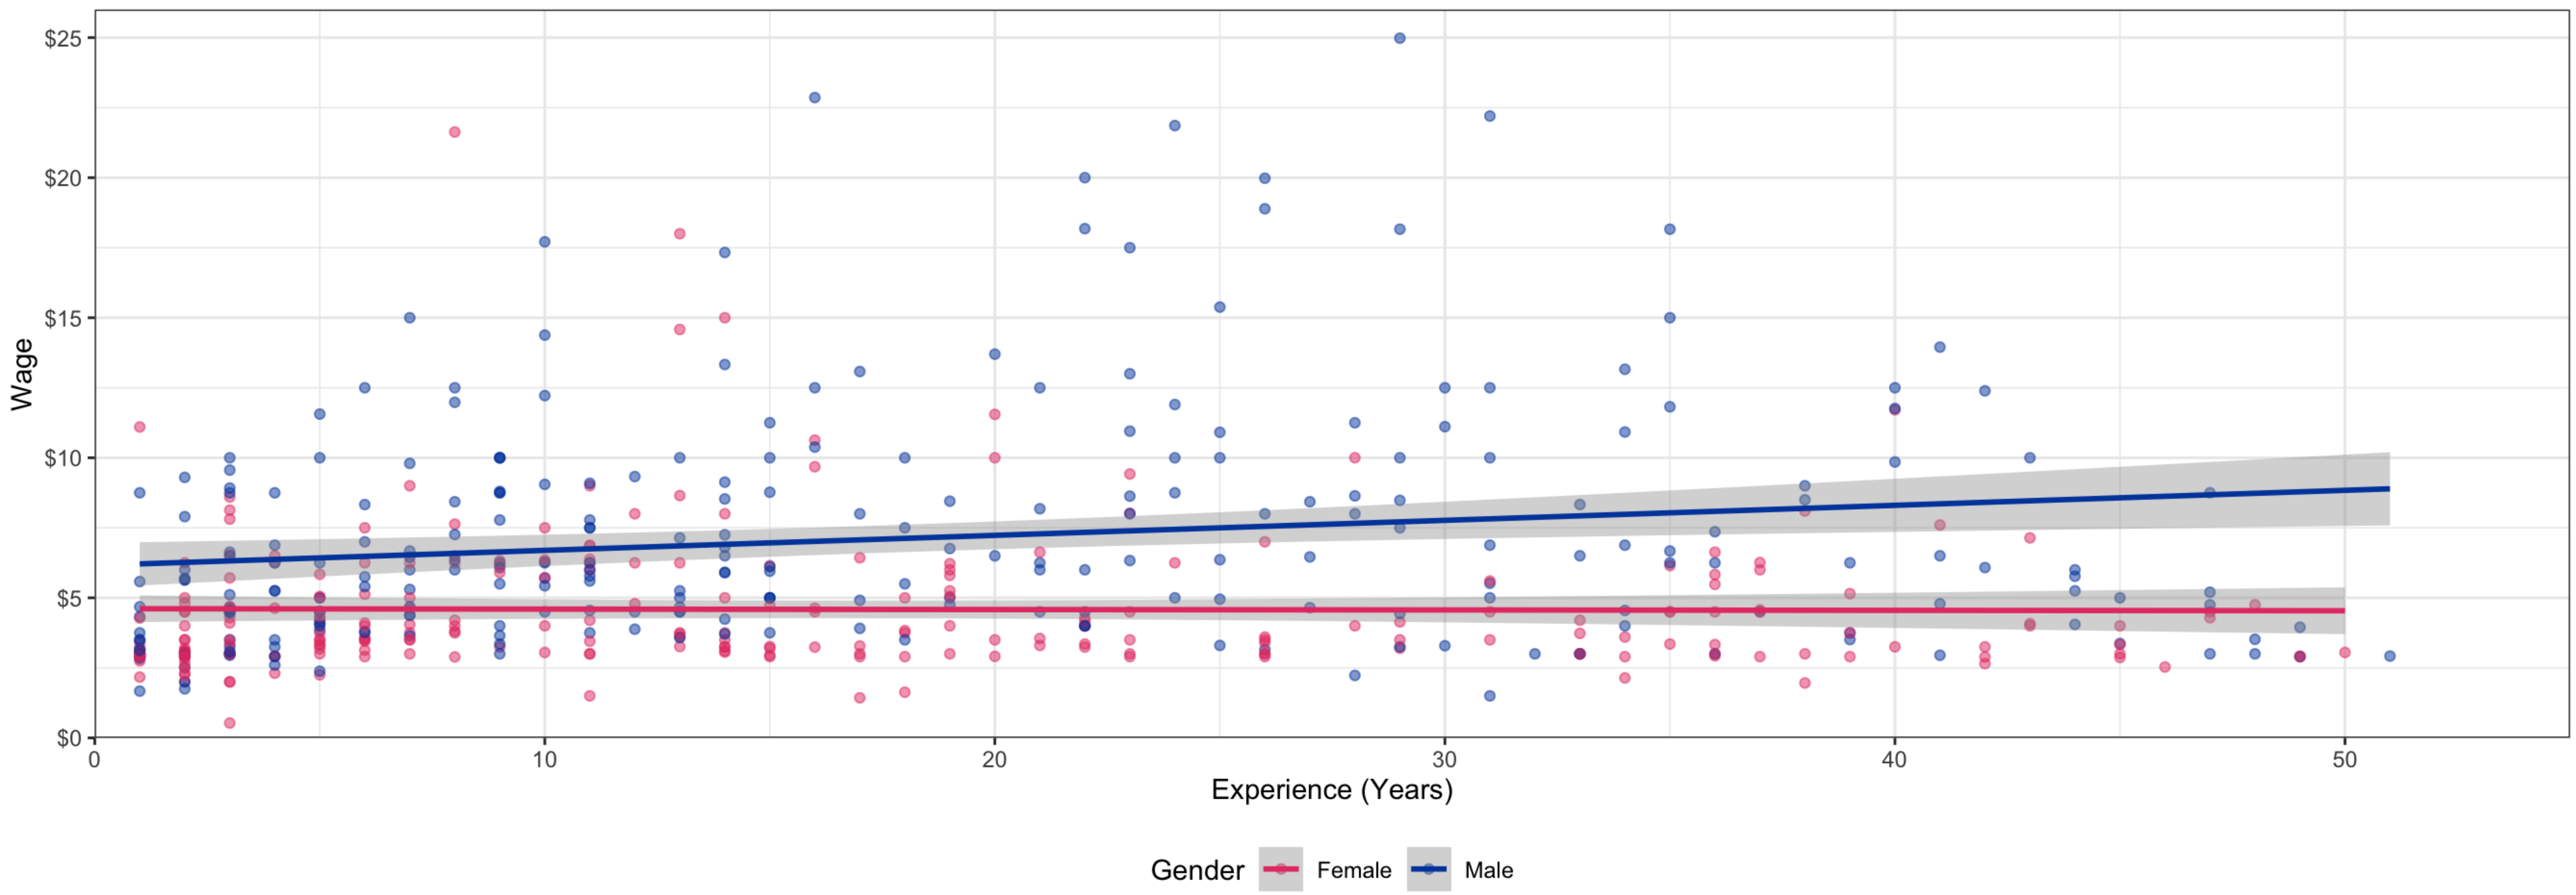

Interactions in Our Example: Scatterplot

- Code

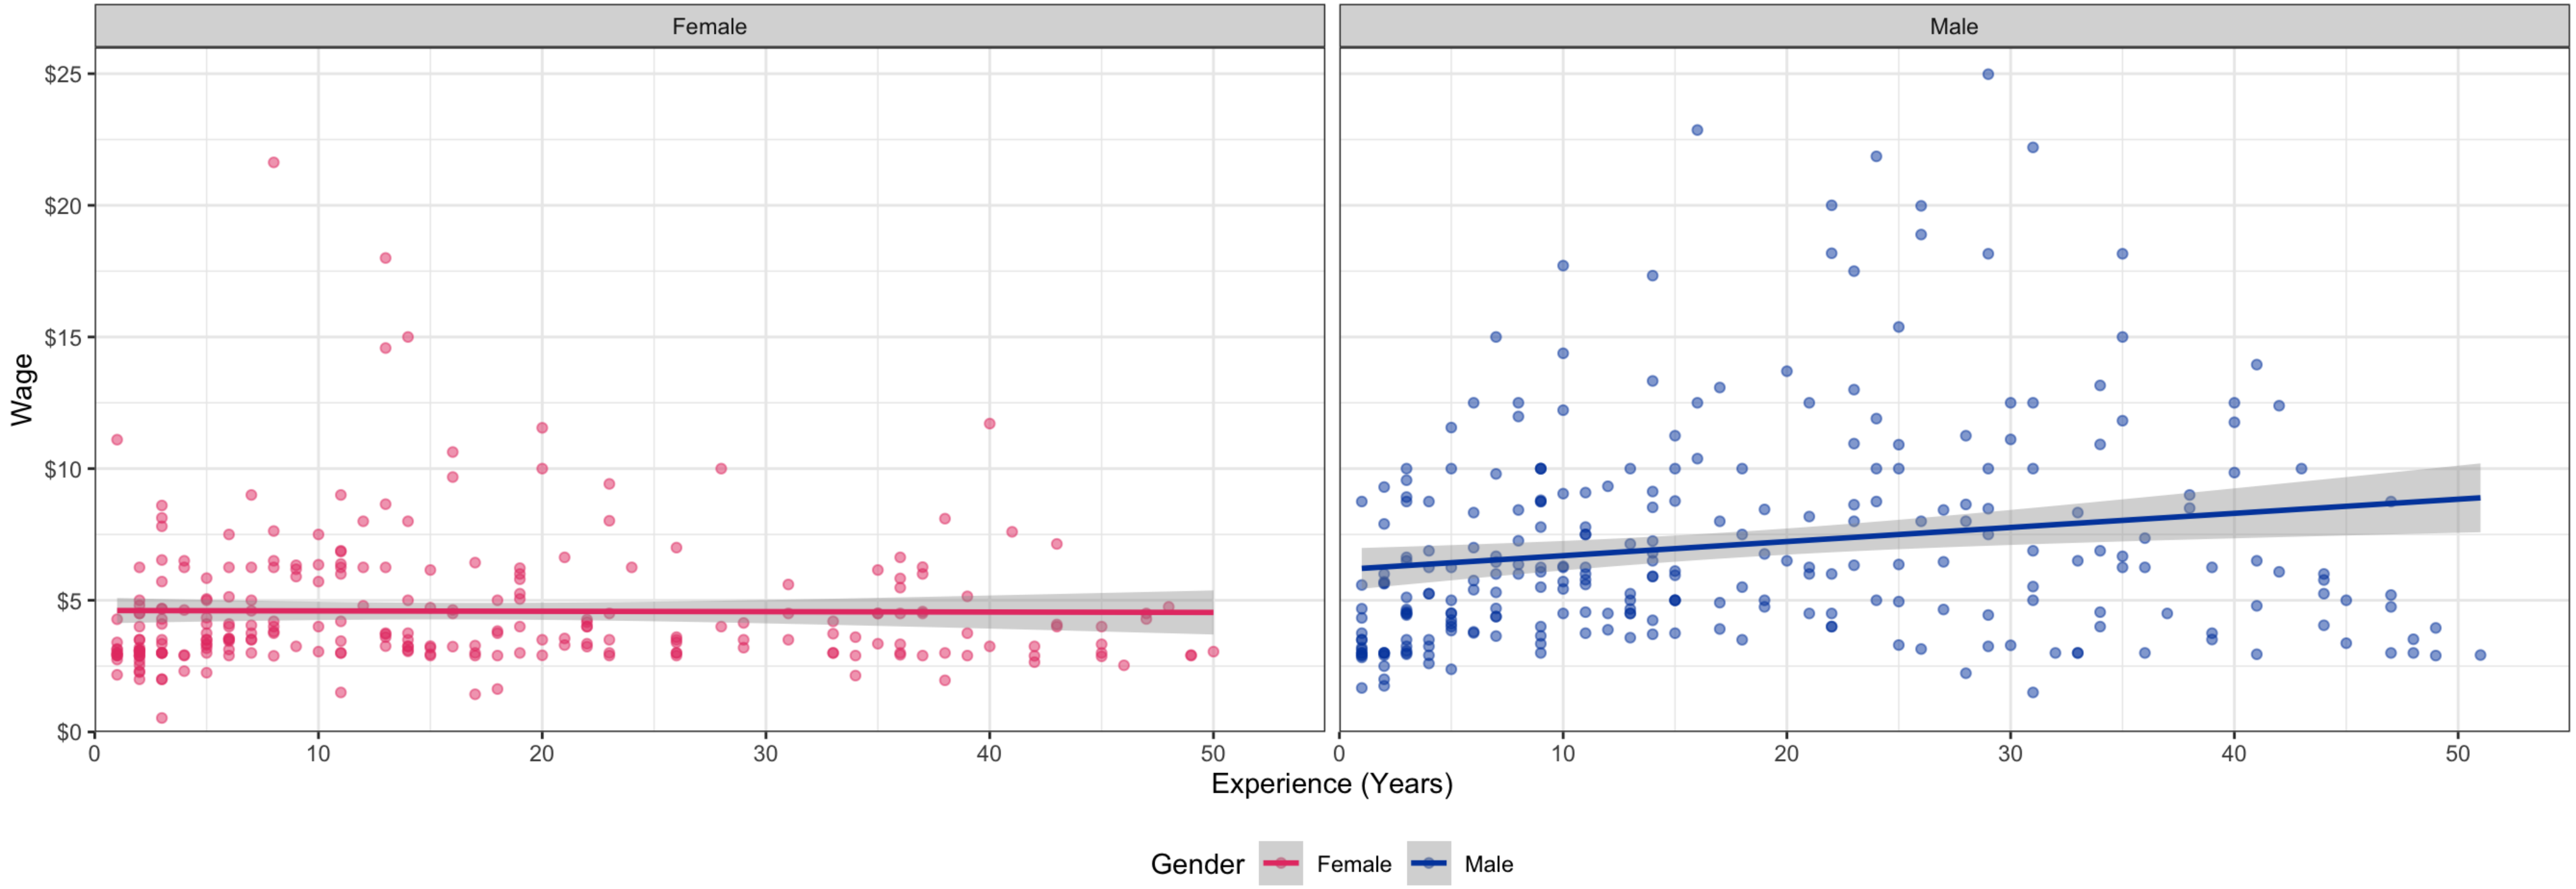

Interactions in Our Example: Regression in R

- Syntax for adding an interaction term is easy[1] in R: `x1 * x2`
  - Or could just do `x1 * x2` (multiply)

```
# both are identical in R
interaction_reg <- lm(wage ~ exper * female, data = wages)
interaction_reg <- lm(wage ~ exper + female + exper * female, data = wages)
```

term <chr>	estimate <dbl>	std.error <dbl>	statistic <dbl>
(Intercept)	6.15827549	0.34167408	18.023830
exper	0.05360476	0.01543716	3.472450
female	-1.54654677	0.48186030	-3.209534
exper:female	-0.05506989	0.02217496	-2.483427

4 rows | 1-4 of 5 columns

1. There are several options here. (1) Using `:`, running `y ~ x1:x2` will run $Y = \beta_0 + \beta_3(X_1 \times X_2)$ only (i.e. not including `x1` and `x2` terms). You of course can add them in yourself by running `y ~`

Interactions in Our Example: Regression

- Code

	Wage
Constant	6.16***
	(0.34)
exper	0.05***
	(0.02)
female	−1.55***
	(0.48)
exper:female	−0.06**
	(0.02)
n	526
Adj. R^2	0.13
SER	3.43

* p < 0.1, ** p < 0.05, *** p < 0.01

Interactions in Our Example: Interpretting Coefficients

$$\widehat{\text{wage}}_i = 6.16 + 0.05\,\text{experience}_i - 1.55\,\text{female}_i - 0.06\,(\text{experience}_i \times \text{female}_i)$$

- $\hat{\beta}_0$: **Men** with 0 years of experience earn 6.16
- $\hat{\beta}_1$: For every additional year of experience, **men** earn \$0.05
- $\hat{\beta}_2$: **Women** with 0 years of experience earn \$1.55 **less than men**
- $\hat{\beta}_3$: **Women** earn \$0.06 **less than men** for every additional year of experience

Interactions in Our Example: As Two Regressions I

$$\widehat{\text{wage}}_i = 6.16 + 0.05\,\text{experience}_i - 1.55\,\text{female}_i - 0.06\,(\text{experience}_i \times \text{female}_i)$$

Regression for men $female = 0$

$$\widehat{\text{wage}}_i = 6.16 + 0.05\,\text{experience}_i$$

- Men with 0 years of experience earn \$6.16 on average
- For every additional year of experience, men earn \$0.05 more on average

Interactions in Our Example: As Two Regressions I

$$\widehat{\text{wage}}_i = 6.16 + 0.05\,\text{experience}_i - 1.55\,\text{female}_i - 0.06\,(\text{experience}_i \times \text{female}_i)$$

Regression for women $female = 1$

$$\begin{aligned}\widehat{\text{wage}}_i &= 6.16 + 0.05\,\text{experience}_i - 1.55(1) - 0.06\,\text{experience}_i \times (1)\\ &= (6.16 - 1.55) + (0.05 - 0.06)\,\text{experience}_i\\ &= 4.61 - 0.01\,\text{experience}_i\end{aligned}$$

- Women with 0 years of experience earn \$4.61 on average
- For every additional year of experience, women earn \$0.01 *less* on average

Interactions in Our Example: Hypothesis Testing

$$\widehat{\text{wage}}_i = 6.16 + 0.05\,\text{experience}_i - 1.55\,\text{female}_i - 0.06\,(\text{experience}_i \times \text{female}_i)$$

term <chr>	estimate <dbl>	std.error <dbl>	statistic <dbl>
(Intercept)	6.15827549	0.34167408	18.023830
exper	0.05360476	0.01543716	3.472450
female	-1.54654677	0.48186030	-3.209534
exper:female	-0.05506989	0.02217496	-2.483427

4 rows | 1-4 of 5 columns

- **Are intercepts of the 2 regressions different?** $H_0 : \beta_2 = 0$
  - Difference between men vs. women for no experience?
  - Is $\hat{\beta_2}$ significant?
  - Yes (reject) H_0: p-value = 0.00
- **Are slopes of the 2 regressions different?** $H_0 : \beta_3 = 0$
  - Difference between men vs. women for marginal effect of experience?
  - Is $\hat{\beta_3}$ significant?
  - Yes (reject) H_0: p-value = 0.01

Interactions Between Two Dummy Variables

Interactions Between Two Dummy Variables

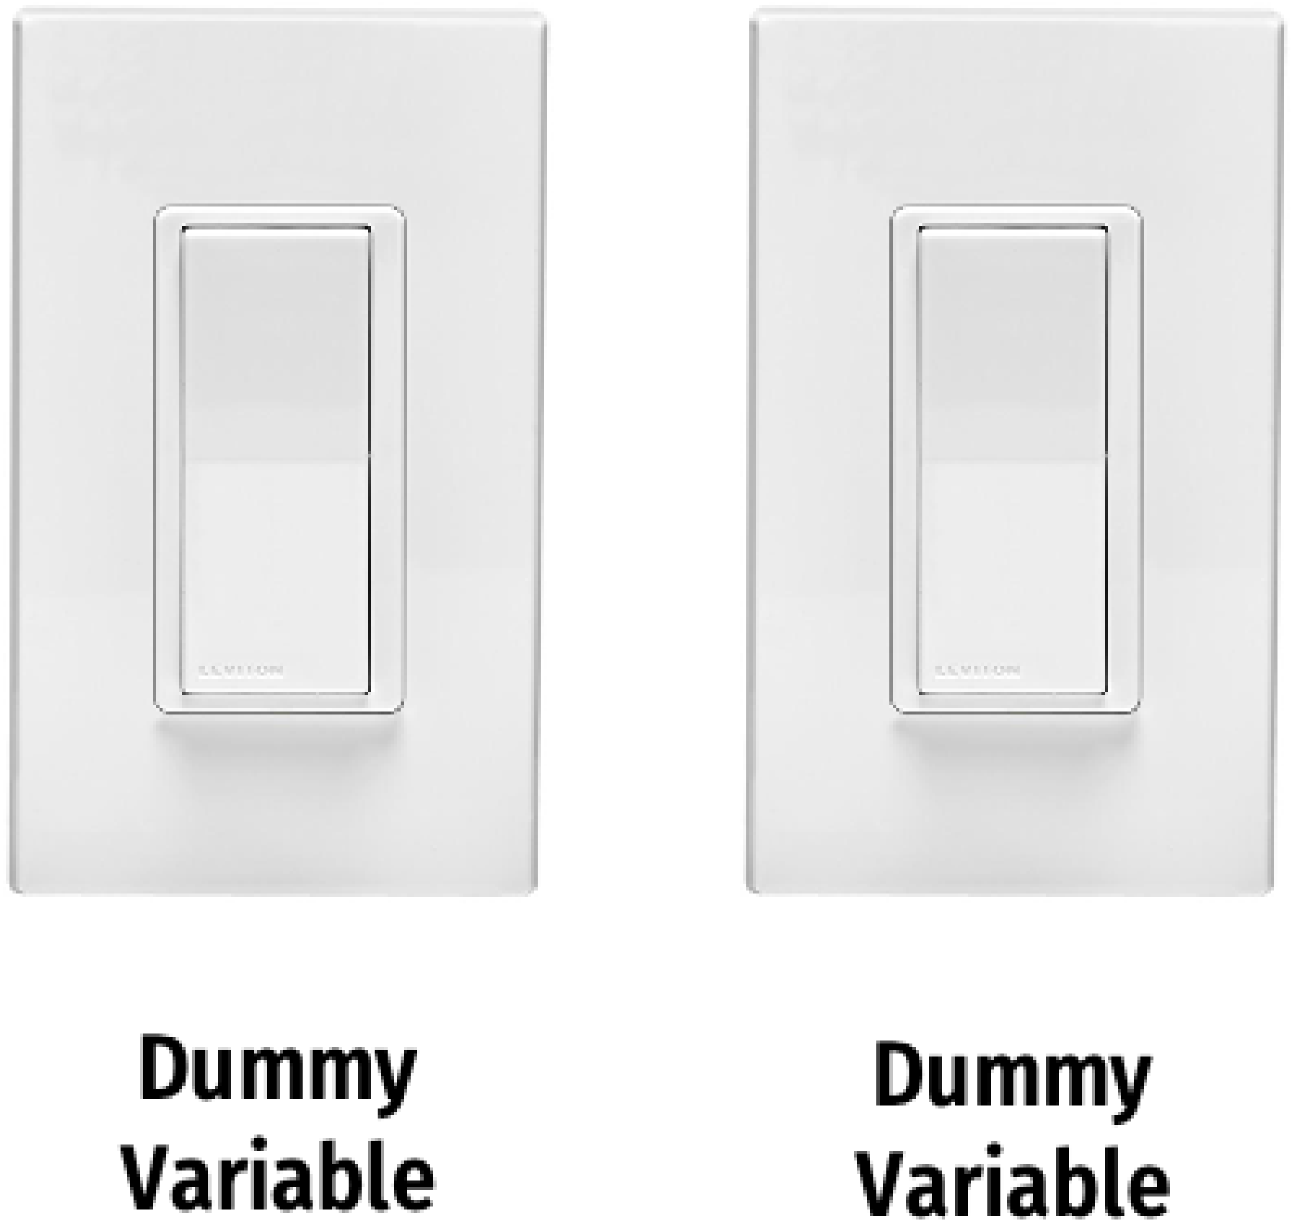

- Does the marginal effect on Y of one dummy going from "off" to "on" change depending on whether the *other* dummy is "off" or "on"?

Interactions Between Two Dummy Variables

$$Y_i = \beta_0 + \beta_1 D_{1i} + \beta_2 D_{2i} + \beta_3 (D_{1i} \times D_{2i})$$

- D_{1i} and D_{2i} are dummy variables
- $\hat{\beta}_1$: effect on Y of going from $D_{1i} = 0$ to $D_{1i} = 1$ when $D_{2i} = 0$
- $\hat{\beta}_2$: effect on Y of going from $D_{2i} = 0$ to $D_{2i} = 1$ when $D_{1i} = 0$
- $\hat{\beta}_3$: effect on Y of going from $D_{1i} = 0$ to $D_{1i} = 1$ when $D_{2i} = 1$
  - *increment* to the effect of D_{1i} going from 0 to 1 when $D_{2i} = 1$ (vs. 0)
- As always, best to think logically about possibilities (when each dummy $= 0$ or $= 1$)

2 Dummy Interaction: Interpretting Coefficients

$$Y_i = \beta_0 + \beta_1 D_{1i} + \beta_2 D_{2i} + \beta_3 (D_{1i} \times D_{2i})$$

- To interpret coefficients, compare cases:
  - Hold D_2 constant (set to some value $D_2 = \mathbf{d_2}$)
  - Plug in 0s or 1s for D_1

$$E(Y|D_1 = 0, D_2 = \mathbf{d_2}) = \beta_0 + \beta_2 \mathbf{d_2}$$
$$E(Y|D_1 = 1, D_2 = \mathbf{d_2}) = \beta_0 + \beta_1(1) + \beta_2 \mathbf{d_2} + \beta_3(1)\mathbf{d_2}$$

- Subtracting the two, the difference is:

$$\beta_1 + \beta_3 \mathbf{d_2}$$

- **The marginal effect of $D_1 \rightarrow Y$ depends on the value of D_2**
  - $\hat{\beta}_3$ is the *increment* to the effect of D_1 on Y when D_2 goes from 0 to 1

Interactions Between 2 Dummy Variables: Example

Example

Does the gender pay gap change if a person is married vs. single?

$$\widehat{\text{wage}}_i = \hat{\beta}_0 + \hat{\beta}_1 \text{ female}_i + \hat{\beta}_2 \text{ married}_i + \hat{\beta}_3 \text{ (female}_i \times \text{married}_i)$$

- Logically, there are 4 possible combinations of $female_i = \{0, 1\}$ and $married_i = \{0, 1\}$

1. **Unmarried men** $(female_i = 0, married_i = 0)$

$$\widehat{wage_i} = \hat{\beta}_0$$

2. **Married men** $(female_i = 0, married_i = 1)$

$$\widehat{wage_i} = \hat{\beta}_0 + \hat{\beta}_2$$

3. **Unmarried women** $(female_i = 1, married_i = 0)$

$$\widehat{wage_i} = \hat{\beta}_0 + \hat{\beta}_1$$

4. **Married women** $(female_i = 1, married_i = 1)$

$$\widehat{wage_i} = \hat{\beta}_0 + \hat{\beta}_1 + \hat{\beta}_2 + \hat{\beta}_3$$

Conditional Group Means in the Data

```r
# get average wage for unmarried men
wages %>%
  filter(female == 0,
         married == 0) %>%
  summarize(mean = mean(wage))
```

mean
<dbl>
5.168023

1 row

```r
# get average wage for unmarried women
wages %>%
  filter(female == 1,
         married == 0) %>%
  summarize(mean = mean(wage))
```

mean
<dbl>
4.611583

1 row

```r
# get average wage for married men
wages %>%
  filter(female == 0,
         married == 1) %>%
  summarize(mean = mean(wage))
```

mean
<dbl>
7.983032

1 row

```r
# get average wage for married women
wages %>%
  filter(female == 1,
         married == 1) %>%
  summarize(mean = mean(wage))
```

mean
<dbl>
4.565909

1 row

Two Dummies Interaction: Group Means

$$\widehat{\text{wage}}_i = \hat{\beta}_0 + \hat{\beta}_1 \text{ female}_i + \hat{\beta}_2 \text{ married}_i + \hat{\beta}_3 \text{ (female}_i \times \text{married}_i)$$

	Men	Women
Unmarried	\$5.17	\$4.61
Married	\$7.98	\$4.57

Two Dummies Interaction: Regression in R I

```
reg_dummies <- lm(wage ~ female + married + female:married, data = wages)
reg_dummies %>% tidy()
```

term <chr>	estimate <dbl>	std.error <dbl>	statistic <dbl>
(Intercept)	5.1680233	0.3614348	14.298631
female	-0.5564399	0.4735578	-1.175020
married	2.8150086	0.4363413	6.451391
female:married	-2.8606829	0.6075577	-4.708496

4 rows | 1-4 of 5 columns

Two Dummies Interaction: Regression in R II

▶ Code

	Wage
Constant	5.17***
	(0.36)
female	−0.56
	(0.47)
married	2.82***
	(0.44)
female:married	-2.86***
	(0.61)
n	526
Adj. R^2	0.18
SER	3.34

* p < 0.1, ** p < 0.05, *** p < 0.01

Two Dummies Interaction: Interpretting Coefficients I

$$\widehat{\text{wage}}_i = 5.17 - 0.56\,\text{female}_i + 2.82\,\text{married}_i - 2.86\,(\text{female}_i \times \text{married}_i)$$

	Men	Women
Unmarried	\$5.17	\$4.61
Married	\$7.98	\$4.57

- Wage for **unmarried** **men**: $\hat{\beta}_0 = 5.17$
- Wage for **married** **men**: $\hat{\beta}_0 + \hat{\beta}_2 = 5.17 + 2.82 = 7.98$
- Wage for **unmarried** **women**: $\hat{\beta}_0 + \hat{\beta}_1 = 5.17 - 0.56 = 4.61$
- Wage for **married** **women**: $\hat{\beta}_0 + \hat{\beta}_1 + \hat{\beta}_2 + \hat{\beta}_3 = 5.17 - 0.56 + 2.82 - 2.86 = 4.57$

Two Dummies Interaction: Interpretting Coefficients II

$$\widehat{\text{wage}}_i = 5.17 - 0.56\,\text{female}_i + 2.82\,\text{married}_i - 2.86\,(\text{female}_i \times \text{married}_i)$$

	Men	Women	Diff
Unmarried	$5.17	$4.61	**$0.56**
Married	$7.98	$4.57	$3.41
Diff	**$2.81**	$0.04	**$2.85**

- $\hat{\beta}_0$: Wage for **unmarried** **men**
- $\hat{\beta}_1$: **Difference** in wages between **men** and **women** who are **unmarried**
- $\hat{\beta}_2$: **Difference** in wages between **married** and **unmarried** **men**
- $\hat{\beta}_3$: **Difference** in:
  - effect of **Marriage** on wages between **men** and **women**
  - effect of **Gender** on wages between **unmarried** and **married** individuals
  - **"difference in differences"**

Interactions Between Two Continuous Variables

Interactions Between Two Continuous Variables

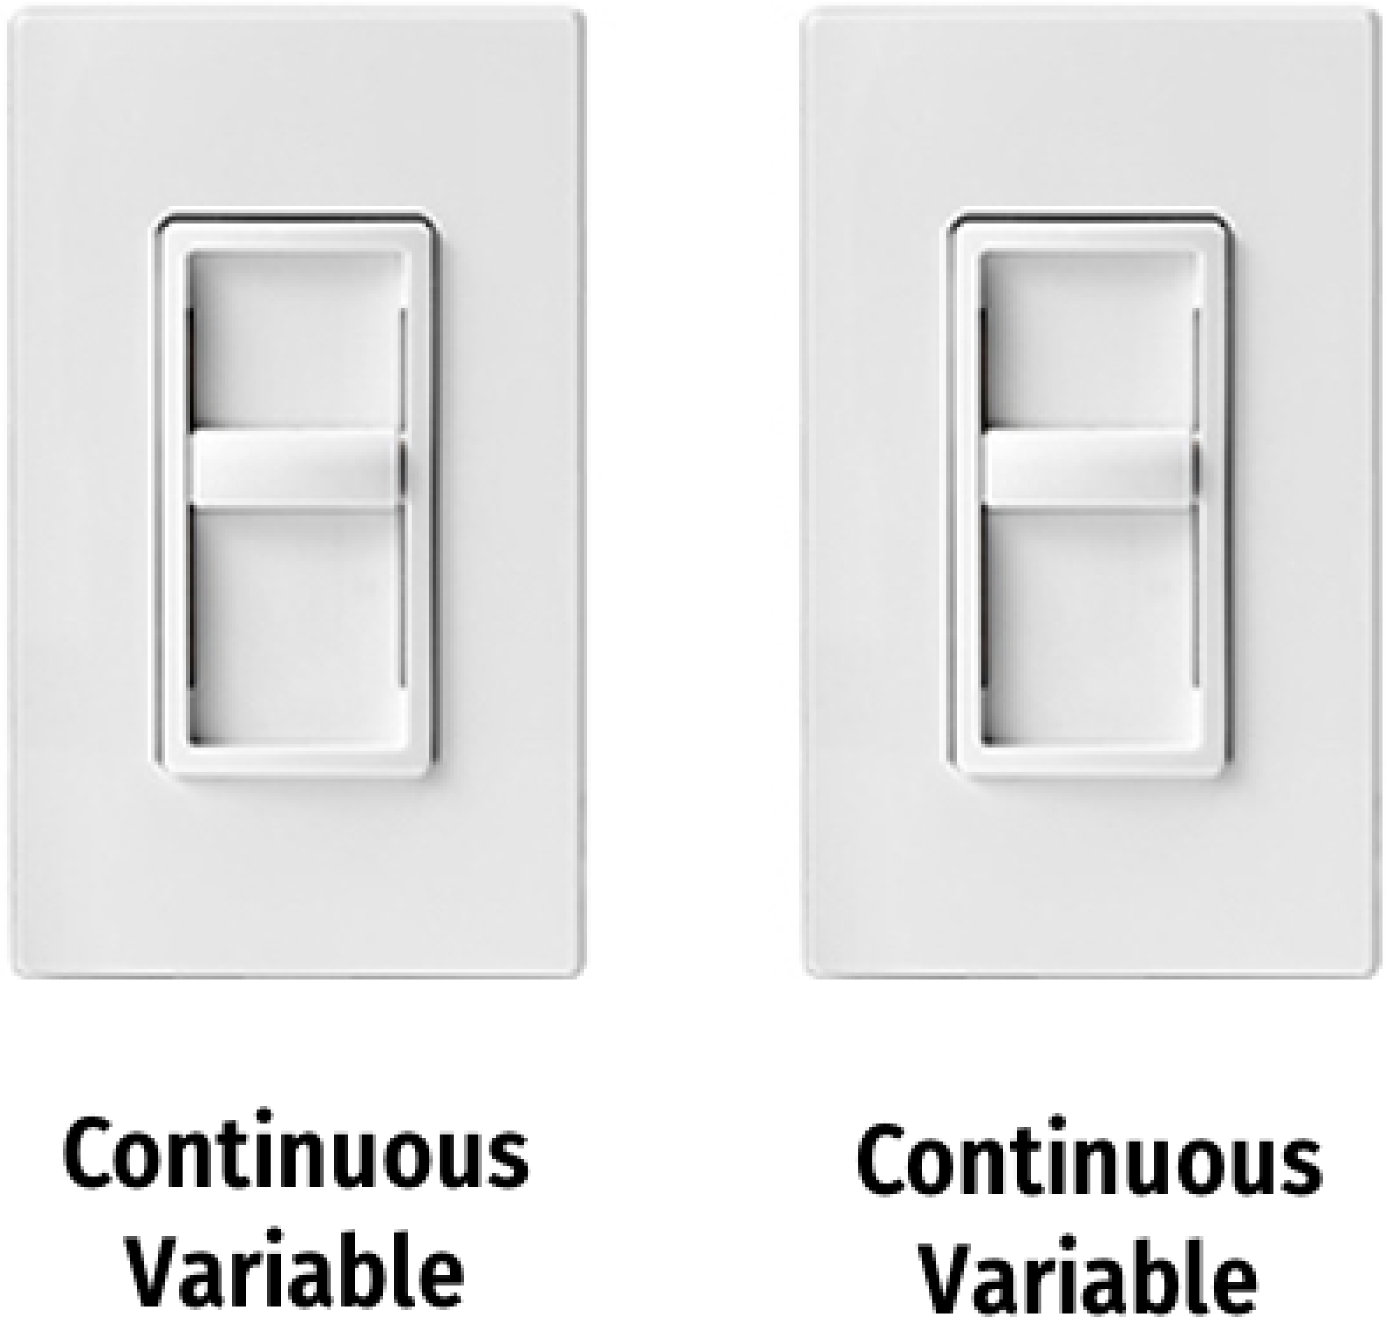

- Does the marginal effect of X_1 on Y depend on what X_2 is set to?

Interactions Between Two Continuous Variables

$$Y_i = \beta_0 + \beta_1 X_{1i} + \beta_2 X_{2i} + \beta_3 (X_{1i} \times X_{2i})$$

- To interpret coefficients, compare changes after changing ΔX_{1i} (holding X_2 constant):

$$Y_i + \Delta Y_i = \beta_0 + \beta_1 (X_1 + \Delta X_{1i}) \beta_2 X_{2i} + \beta_3 ((X_{1i} + \Delta X_{1i}) \times X_{2i})$$

- Take the difference to get:

$$\begin{aligned} \Delta Y_i &= \beta_1 \Delta X_{1i} + \beta_3 X_{2i} \Delta X_{1i} \\ \frac{\Delta Y_i}{\Delta X_{1i}} &= \beta_1 + \beta_3 X_{2i} \end{aligned}$$

- **The effect of $X_1 \rightarrow Y$ depends on the value of X_2**
  - β_3: *increment* to the effect of $X_1 \rightarrow Y$ for every 1 unit change in X_2

Continuous Variables Interaction: Example

Example

Do education and experience interact in their determination of wages?

$$\widehat{\text{wage}}_i = \hat{\beta}_0 + \hat{\beta}_1 \text{ education}_i + \hat{\beta}_2 \text{ experience}_i + \hat{\beta}_3 \text{ (education}_i \times \text{experience}_i)$$

- Estimated effect of education on wages depends on the amount of experience (and vice versa)!

$$\frac{\Delta \text{wage}}{\Delta \text{education}} = \hat{\beta}_1 + \beta_3 \text{ experience}_i$$

$$\frac{\Delta \text{wage}}{\Delta \text{experience}} = \hat{\beta}_2 + \beta_3 \text{ education}_i$$

- This is a type of nonlinearity (we will examine nonlinearities next lesson)

Continuous Variables Interaction: In R I

```
reg_cont <- lm(wage ~ educ + exper + educ:exper, data = wages)
reg_cont %>% tidy()
```

term <chr>	estimate <dbl>	std.error <dbl>
(Intercept)	-2.859915627	1.181079647
educ	0.601735470	0.089899977
exper	0.045768911	0.042613758
educ:exper	0.002062345	0.003490614

4 rows | 1-3 of 5 columns

Continuous Variables Interaction: In R II

▶ Code

	Wage
Constant	-2.86**
	(1.18)
educ	0.60***
	(0.09)
exper	0.05
	(0.04)
educ:exper	0.00
	(0.00)
n	526
Adj. R^2	0.22
SER	3.25

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Continuous Variables Interaction: Marginal Effects

$$\widehat{\text{wage}}_i = -2.860 + 0.602\,\text{education}_i + 0.047\,\text{experience}_i + 0.002\,(\text{education}_i \times \text{experi}$$

Marginal Effect of *Education* on Wages by Years of *Experience*:

Experience	$\frac{\Delta\text{wage}}{\Delta\text{education}} = \hat{\beta}_1 + \hat{\beta}_3\,\text{experience}$
5 years	$0.602 + 0.002(5) = 0.612$
10 years	$0.602 + 0.002(10) = 0.622$
15 years	$0.602 + 0.002(15) = 0.632$

- Marginal effect of education $\rightarrow$ wages **increases** with more experience

Continuous Variables Interaction: Marginal Effects

$$\widehat{\text{wage}}_i = -2.860 + 0.602\,\text{education}_i + 0.047\,\text{experience}_i + 0.002\,(\text{education}_i \times \text{experi}$$

Marginal Effect of *Experience* on Wages by Years of *Education*:

Education	$\frac{\Delta \text{wage}}{\Delta \text{experience}} = \hat{\beta}_2 + \hat{\beta}_3 \, \text{education}$
5 years	$0.047 + 0.002(5) = 0.057$
10 years	$0.047 + 0.002(10) = 0.067$
15 years	$0.047 + 0.002(15) = 0.077$

- Marginal effect of experience $\rightarrow$ wages **increases** with more education
- If you want to estimate the marginal effects more precisely, and graph them, see the appendix in today's appendix