Diff-in-diff I & II

Sessions 8–9

PMAP 8521: Program evaluation
Andrew Young School of Policy Studies

Plan for today

Quasi-experiments

Interactions & regression

Two wrongs make a right

Diff-in-diff assumptions

Quasi-experiments

RCTs are great!

Super impractical to do all the time though!

Quasi-experiments

You can't always randomly assign people to do things

So let other people (or the government, or nature, or something else) do it for you

Quasi-experiments

Quasi-experiment
A situation where you, as researcher,
did not assign people to treatment/control

External validity 👍 **Selection 👎**

Assignment to treatment is "as if" random

Quasi-experiments vs. DAG adjustment

We did a lot of work with DAGs!
You're good at closing backdoors with matching and IPW

DAGs can work for any kind of observational data,
even without a quasi-experimentalish situation

Quasi-experiments are a little different:
the **context** isolates pathway between treatment and outcome

They're wildly popular in social sciences (especially economics!),
maybe more credible (?) there than just making DAG adjustments

You can still draw a DAG for a quasi-experiment though!

Analyzing quasi-experiments

Difference-in-differences

DiD; DD; diff-in-diff

Regression discontinuity

RD; RDD

Instrumental variables

IV

Interactions & regression

Sliders and switches

Categorical variables

Continuous variables

$$\widehat{\text{Happiness}} = \beta_0 + \beta_1 \text{Life expectancy} + \beta_2 \text{Latin America} + \varepsilon$$

```
model1 <- lm(happiness_score ~ life_expectancy + latin_america,
             data = world_happiness)
tidy(model1)
```

```
## # A tibble: 3 × 5
##   term                          estimate std.error statistic  p.value
##   <chr>                            <dbl>     <dbl>     <dbl>    <dbl>
## 1 (Intercept)                     -2.08    0.537       -3.87 1.61e- 4
## 2 life_expectancy                  0.102   0.00745     13.7  1.95e-28
## 3 latin_americaLatin America       0.623   0.173        3.61 4.17e- 4
```

Life expectancy = continuous / slider

"For every 1-year increase in life expectancy, happiness is associated with a β_1 increase"

Latin America = categorical / switch

"Being in Latin America is associated with a β_2 increase in happiness"

Indicators and interactions

Indicators (dummies)

Change in intercept for specific group

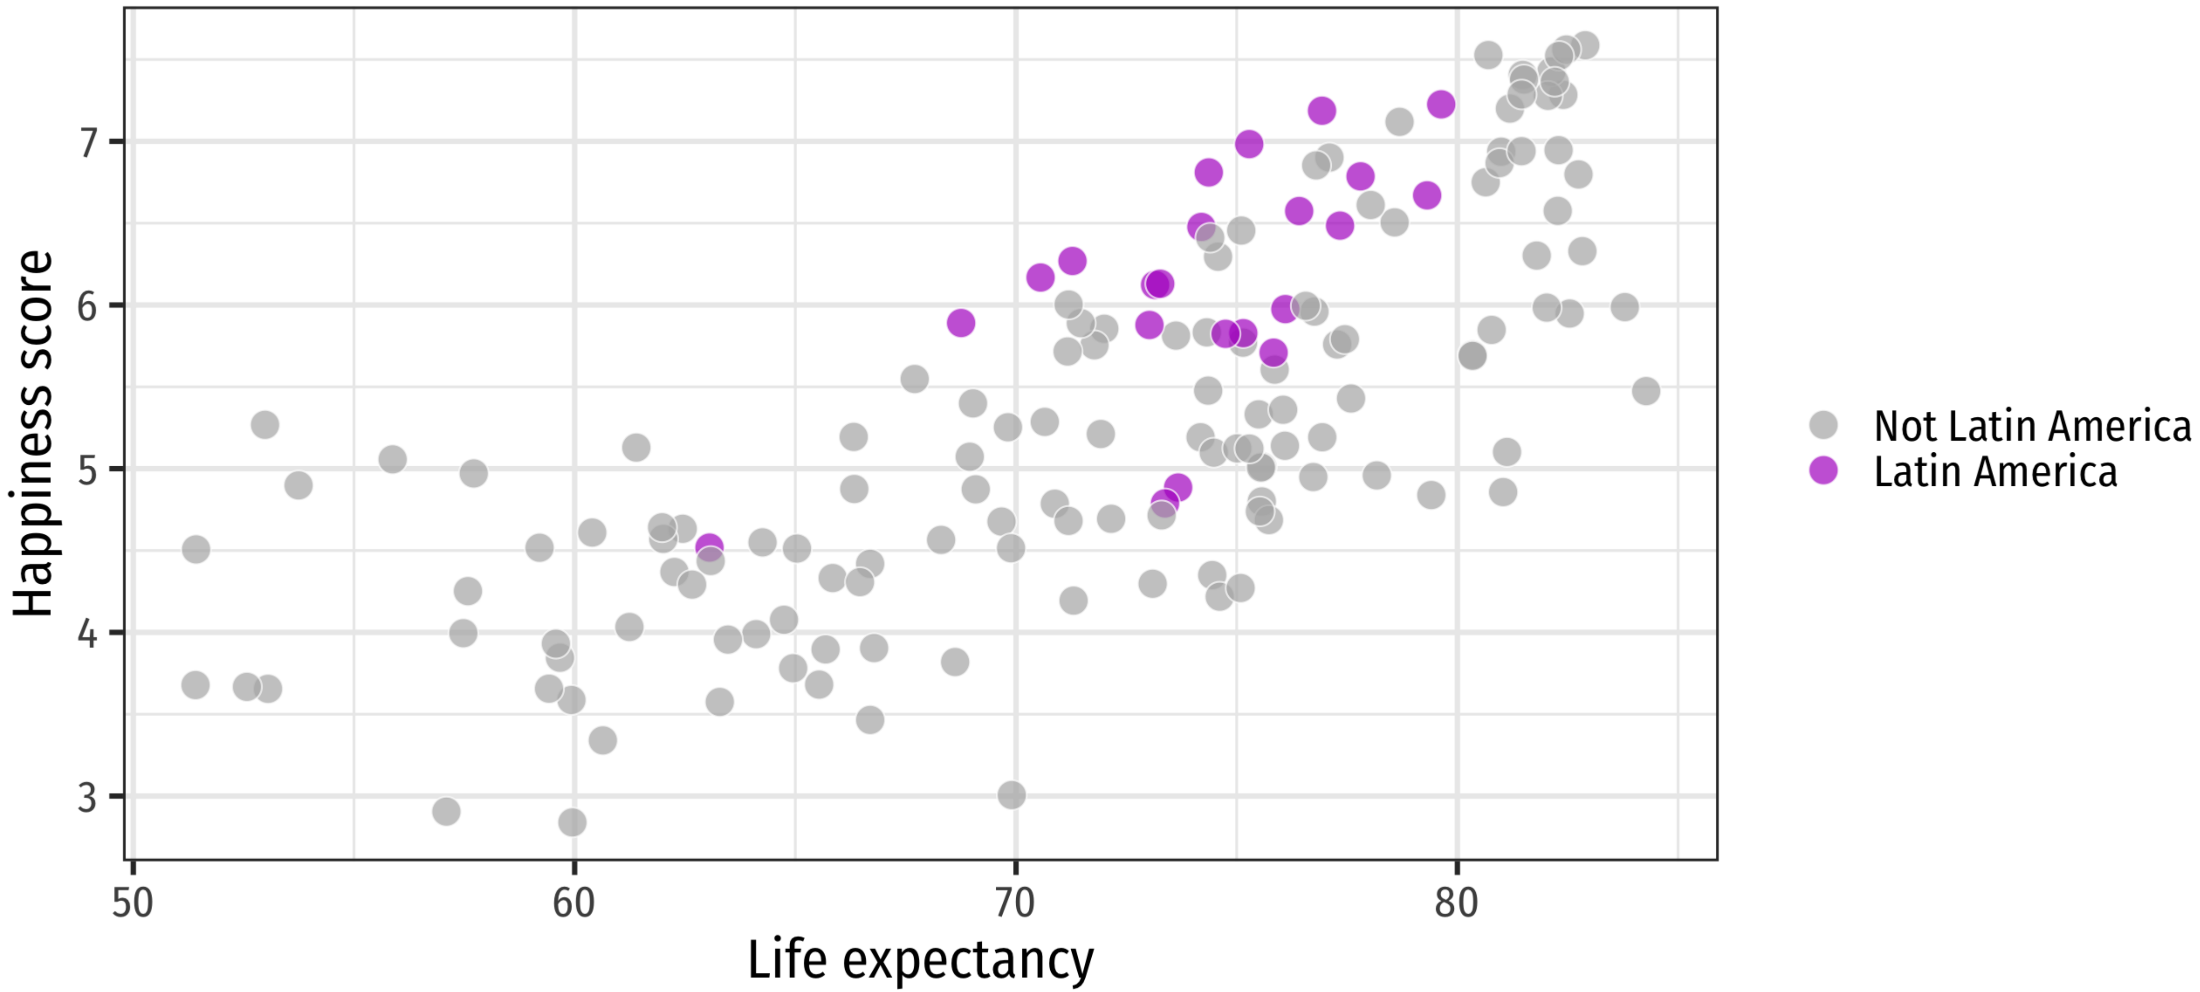
Happiness score
7
6
5
4
3
50
60
70
80
Life expectancy
Not Latin America
Latin America

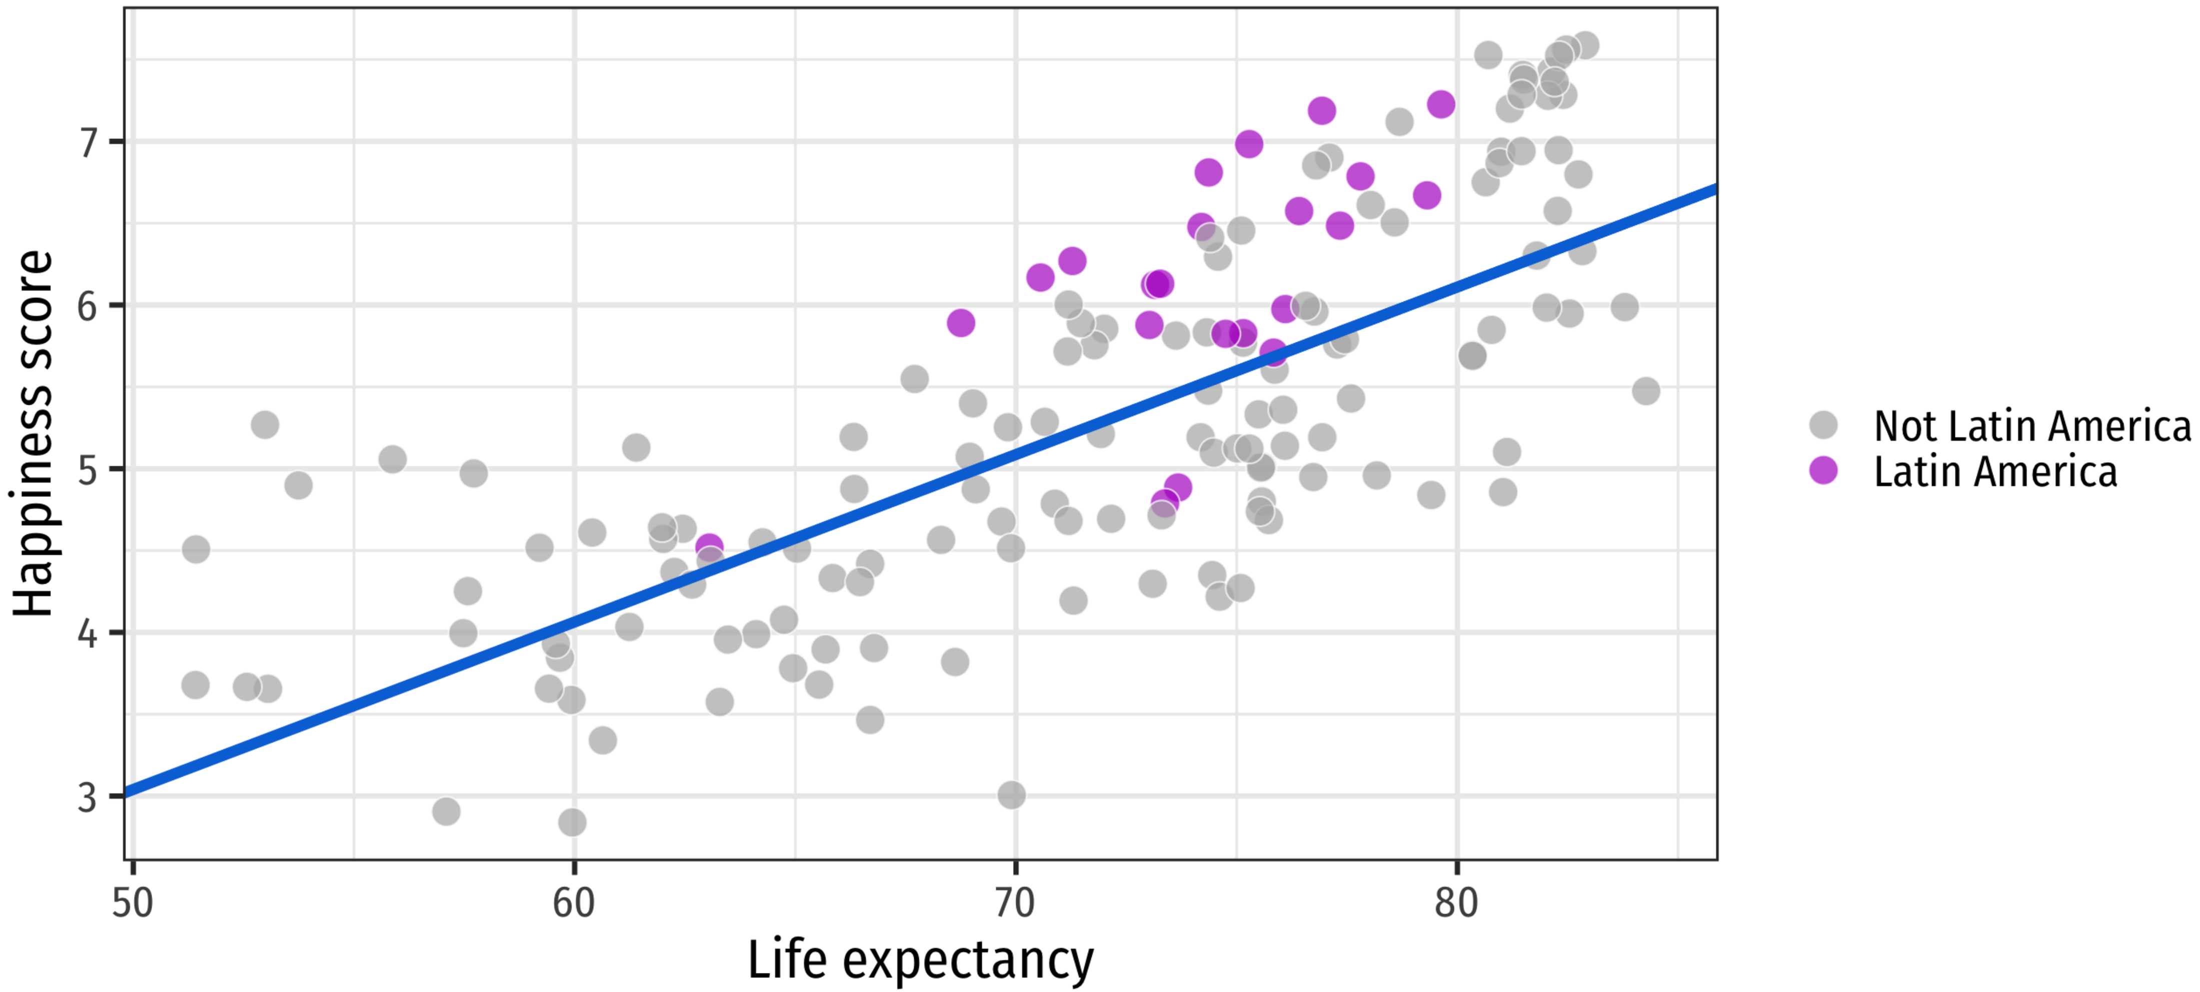

World slope = 0.102

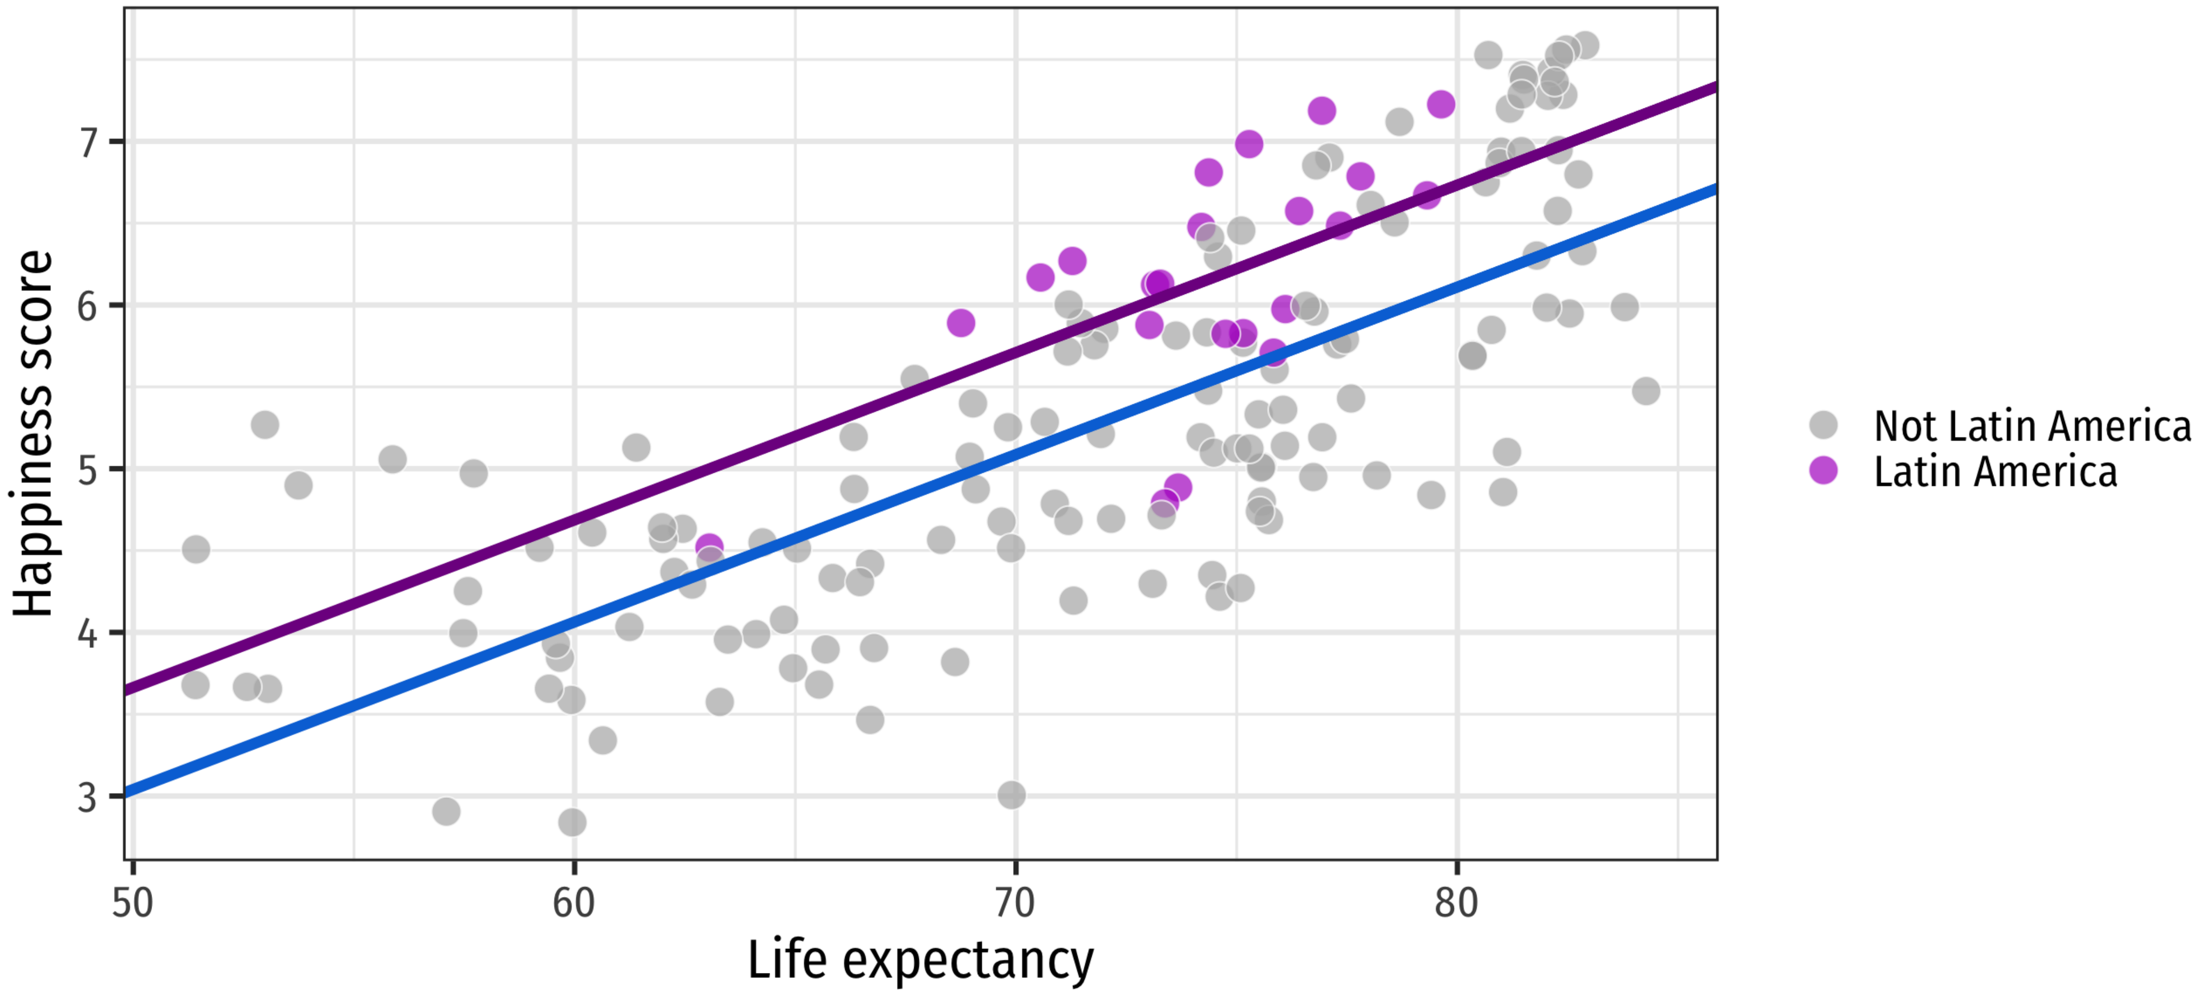

Latin America intercept shifted up 0.62; line has same slope as world (0.102)

$$\begin{aligned}\widehat{\text{Happiness}} =& \beta_0 + \beta_1 \text{Life expectancy} + \beta_2 \text{Latin America} + \\ & \beta_3 (\text{Life expectancy} \times \text{Latin America}) + \varepsilon\end{aligned}$$

```
model2 <- lm(happiness_score ~ life_expectancy + latin_america +
               (life_expectancy * latin_america), data = world_happiness)
tidy(model2)
```

```
## # A tibble: 4 × 5
##   term                                 estimate std.error stati…¹  p.value
##   <chr>                                   <dbl>     <dbl>   <dbl>    <dbl>
## 1 (Intercept)                           -2.02     0.545    -3.70  2.98e- 4
## 2 life_expectancy                        0.102    0.00757  13.4   1.65e-27
## 3 latin_americaLatin America            -1.52     3.36     -0.450 6.53e- 1
## 4 life_expectancy:latin_americaLatin America   0.0288    0.0453    0.637 5.25e- 1
## # … with abbreviated variable name ¹statistic
```

"In Latin America, for every 1-year increase in life expectancy, happiness is associated with a β_1 + β_3 increase *and* the intercept is β_2 lower"

Indicators and interactions

Indicators (dummies)

Change in intercept for specific group

Interactions

Change in slope for specific group

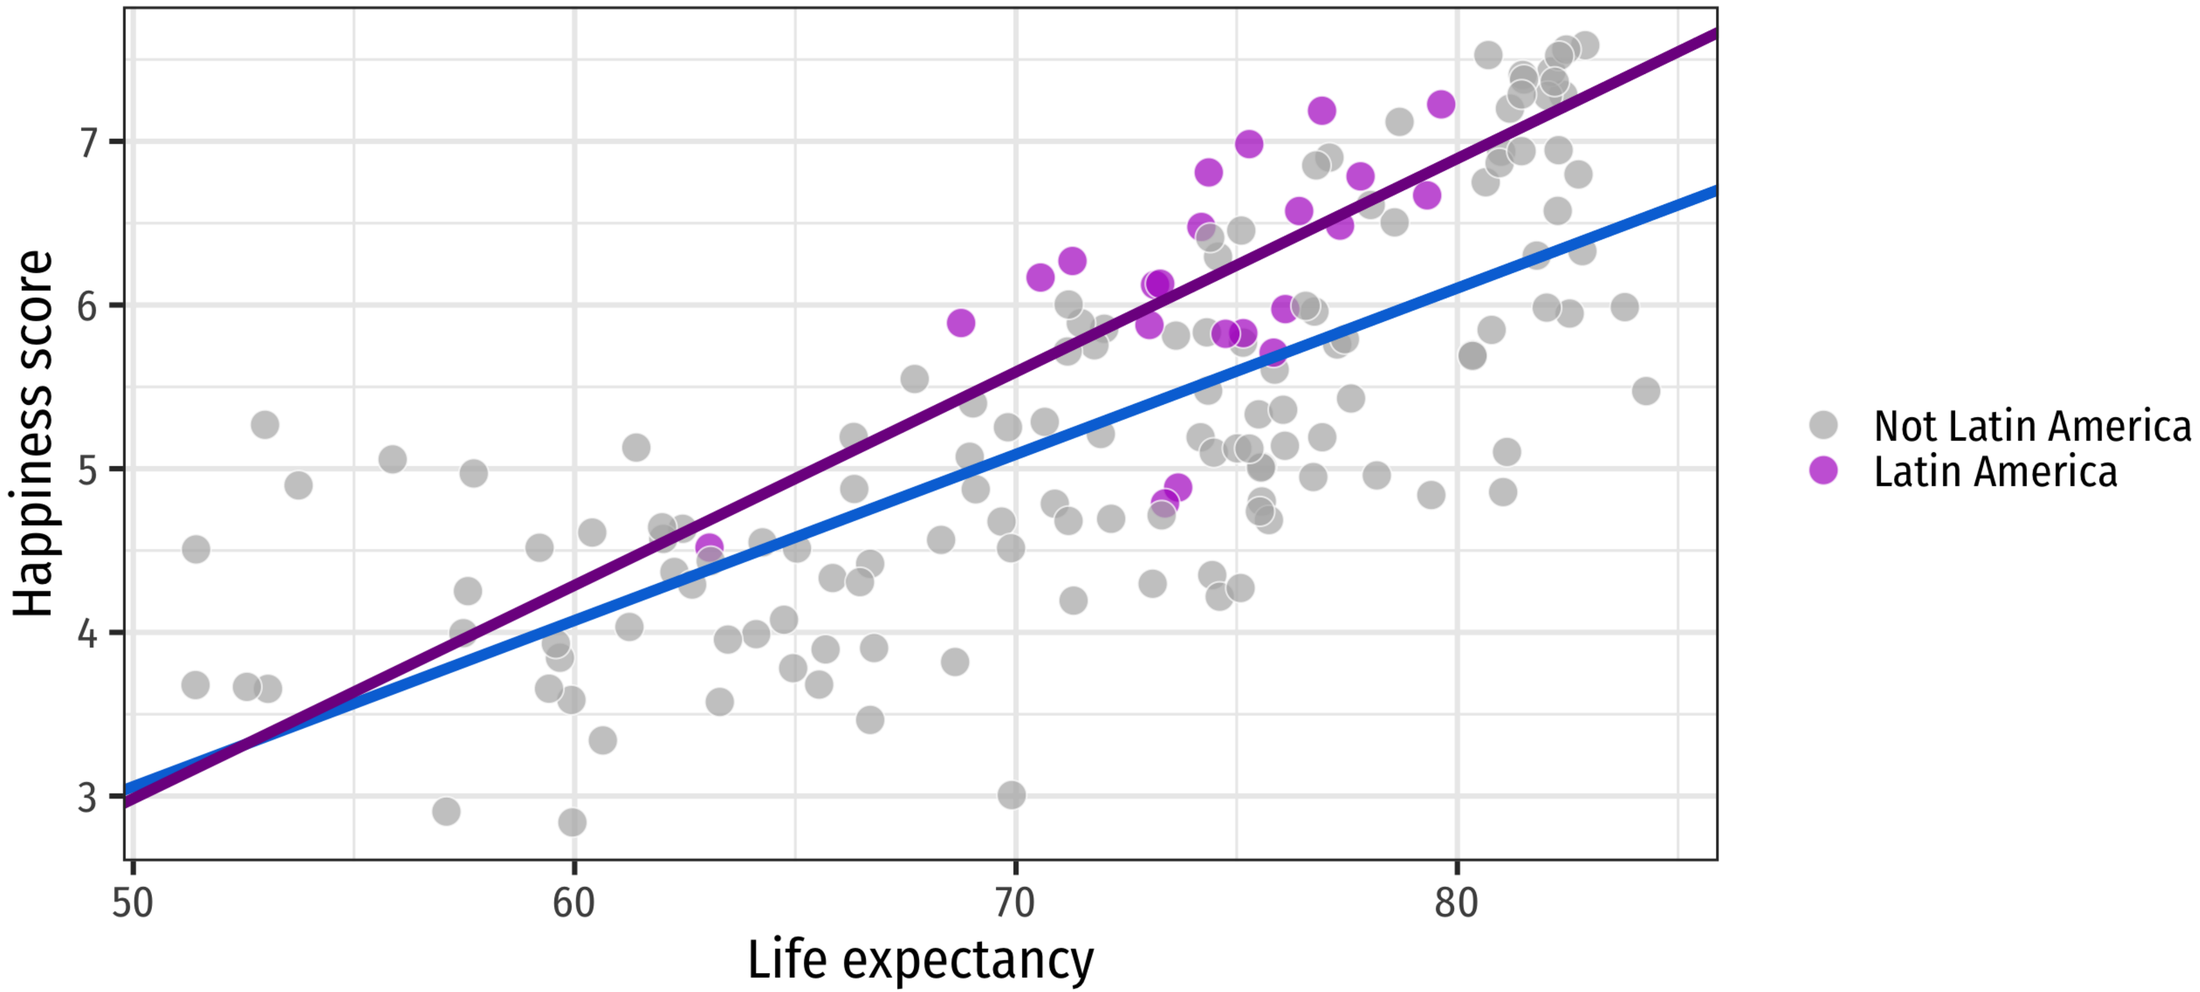

Latin America slope is 0.029 + 0.102 = 0.13; different from rest of the world

Interactions

What would happen if you ran this?

```
model3 <- lm(happiness_score ~ (life_expectancy * latin_america),
             data = world_happiness)
```

```
## # A tibble: 4 × 5
##   term                                     estimate std.error stati…¹  p.value
##   <chr>                                       <dbl>     <dbl>   <dbl>    <dbl>
## 1 (Intercept)                               -2.02     0.545    -3.70  2.98e- 4
## 2 life_expectancy                            0.102    0.00757  13.4   1.65e-27
## 3 latin_americaLatin America                -1.52     3.36     -0.450 6.53e- 1
## 4 life_expectancy:latin_americaLatin America  0.0288   0.0453    0.637 5.25e- 1
## # … with abbreviated variable name ¹statistic
```

It still works!

Both terms have to be in the model; R will add them for you if you leave them out

Interactions

What would happen if you ran this?

```r
model4 <- lm(happiness_score ~ life_expectancy * region, # region has multiple categories
             data = world_happiness)
```

```
## # A tibble: 14 × 5
##    term                                      estim…¹ std.e…² stati…³ p.value
##    <chr>                                       <dbl>   <dbl>   <dbl>   <dbl>
##  1 (Intercept)                               -2.81    2.05   -1.37   1.73e-1
##  2 life_expectancy                            0.112   0.0271  4.12   6.33e-5
##  3 regionEurope & Central Asia               -2.78    2.76   -1.01   3.16e-1
##  4 regionLatin America & Caribbean           -0.724   3.72   -0.195  8.46e-1
##  5 regionMiddle East & North Africa          -3.13    3.14   -0.997  3.21e-1
##  6 regionNorth America                        2.88   23.2     0.124  9.01e-1
##  7 regionSouth Asia                           4.98    5.54    0.898  3.71e-1
##  8 regionSub-Saharan Africa                   6.33    2.48    2.55   1.18e-2
##  9 life_expectancy:regionEurope & Central Asia  0.0367  0.0361  1.02   3.11e-1
## 10 life_expectancy:regionLatin America & Caribb…  0.0187  0.0497  0.376  7.07e-1
## 11 life_expectancy:regionMiddle East & North Af…  0.0410  0.0419  0.978  3.30e-1
## 12 life_expectancy:regionNorth America       -0.0221  0.288  -0.0767 9.39e-1
## 13 life_expectancy:regionSouth Asia          -0.0768  0.0790 -0.972  3.33e-1
## 14 life_expectancy:regionSub-Saharan Africa  -0.101   0.0354 -2.84   5.12e-3
## # … with abbreviated variable names ¹estimate, ²std.error, ³statistic
```

Changes in slopes and intercepts for each region

General idea of interactions

The *additional* change that happens when combining two explanatory variables

Life expectancy effect

Latin America effect

Additional life expectancy effect in Latin America

Is there a discount when combining cheese and chili?

What is the cheese effect?

What is the chili effect?

What is the chili × cheese effect?

Two wrongs make a right


I ♥ federalism
(for the natural experiments)

Raising the minimum wage

What happens if you raise the minimum wage?

Economic theory says there should be fewer jobs

New Jersey in 1992

$\$4.25 \longrightarrow \$5.05$

Before vs. after

Average # of jobs per fast food restaurant in NJ

$\text{New Jersey}_{\text{Before change}} = 20.44$

$\text{New Jersey}_{\text{After change}} = 21.03$

$\Delta = 0.59$

Is this the causal effect?

Treatment vs. control

Average # of jobs per fast food restaurant

$\text{Pennsylvania}_{\text{After change}} = 21.17$

$\text{New Jersey}_{\text{After change}} = 21.03$

$\Delta = -0.14$

Is this the causal effect?

Problems

Comparing only before/after

You're only looking at the treatment group!

Impossible to know if change happened because of treatment or just naturally

Comparing only treatment/control

You're only looking at post-treatment values

Impossible to know if change happened because of natural growth

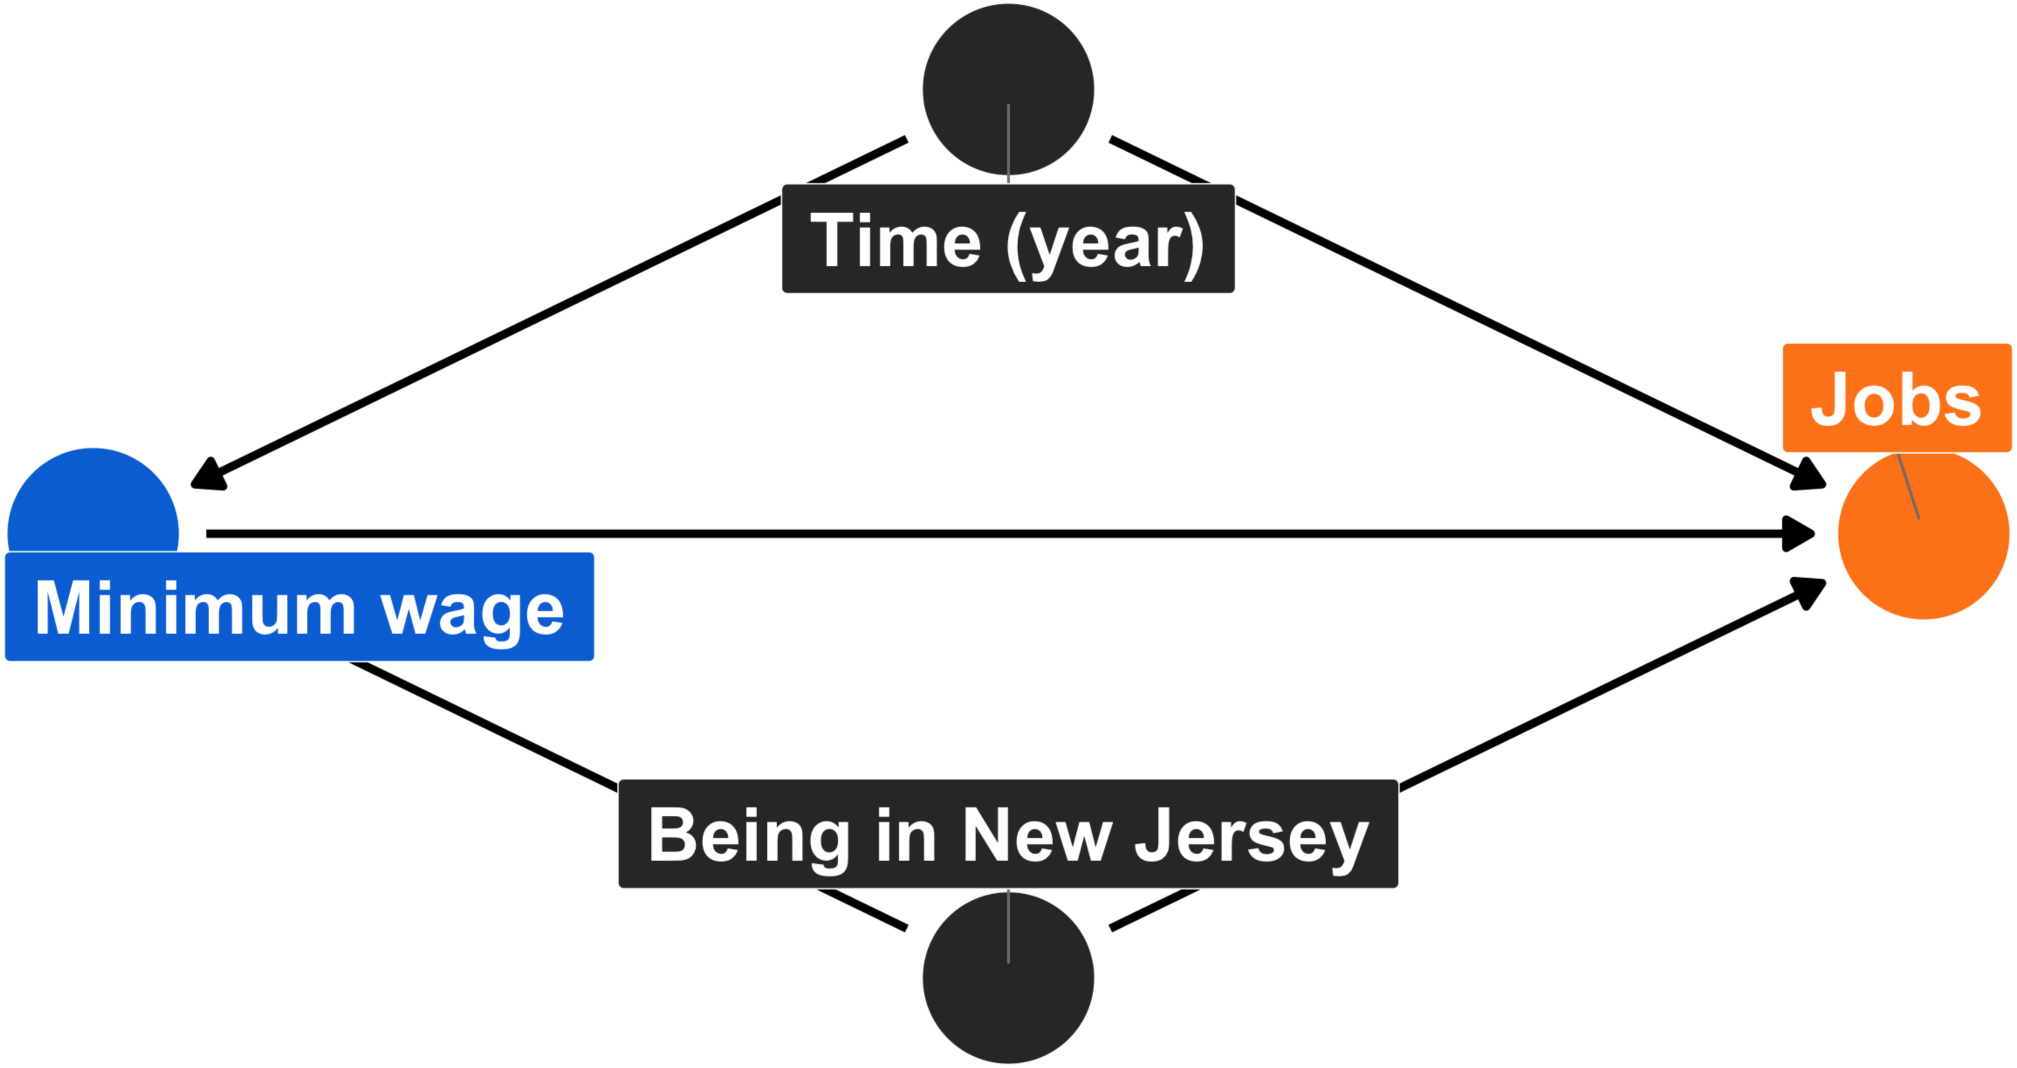
Time (year)
Jobs
Minimum wage
Being in New Jersey

	Pre mean	Post mean	
Control	**A** (never treated)	**B** (never treated)	
Treatment	**C** (not yet treated)	**D** (treated)	

	Pre mean	Post mean	∆ (post − pre)
Control	**A** (never treated)	**B** (never treated)	**B − A**
Treatment	**C** (not yet treated)	**D** (treated)	**D − C**

∆ (post − pre) = within-unit growth

	Pre mean	Post mean	
Control	**A** (never treated)	**B** (never treated)	
Treatment	**C** (not yet treated)	**D** (treated)	
∆ (treatment − control)	**C − A**	**D − B**	

∆ (treatment − control) = across-group growth

	Pre mean	Post mean	Δ (post − pre)
Control	**A** (never treated)	**B** (never treated)	**B − A**
Treatment	**C** (not yet treated)	**D** (treated)	**D − C**
Δ (treatment − control)	**C − A**	**D − B**	**(D − C) − (B − A)** *or* **(D − B) − (C − A)**

$\Delta_{\text{within units}} - \Delta_{\text{within groups}}$ =
Difference-in-differences =
causal effect!

$$
\begin{aligned}
\mathrm{DD} = \; & \left(\bar{x}_{\text{treatment, post}} - \bar{x}_{\text{treatment, pre}}\right) \\
& - \left(\bar{x}_{\text{control, post}} - \bar{x}_{\text{control, pre}}\right)
\end{aligned}
$$

	Pre mean	Post mean	Δ (post – pre)
Pennsylvania	**23.33** A	**21.17** B	**-2.16** B – A
New Jersey	**20.44** C	**21.03** D	**0.59** D – C
Δ **(NJ – PA)**	**-2.89** C – A	**-0.14** D – B	**(0.59) – (–2.16) =** **2.76**

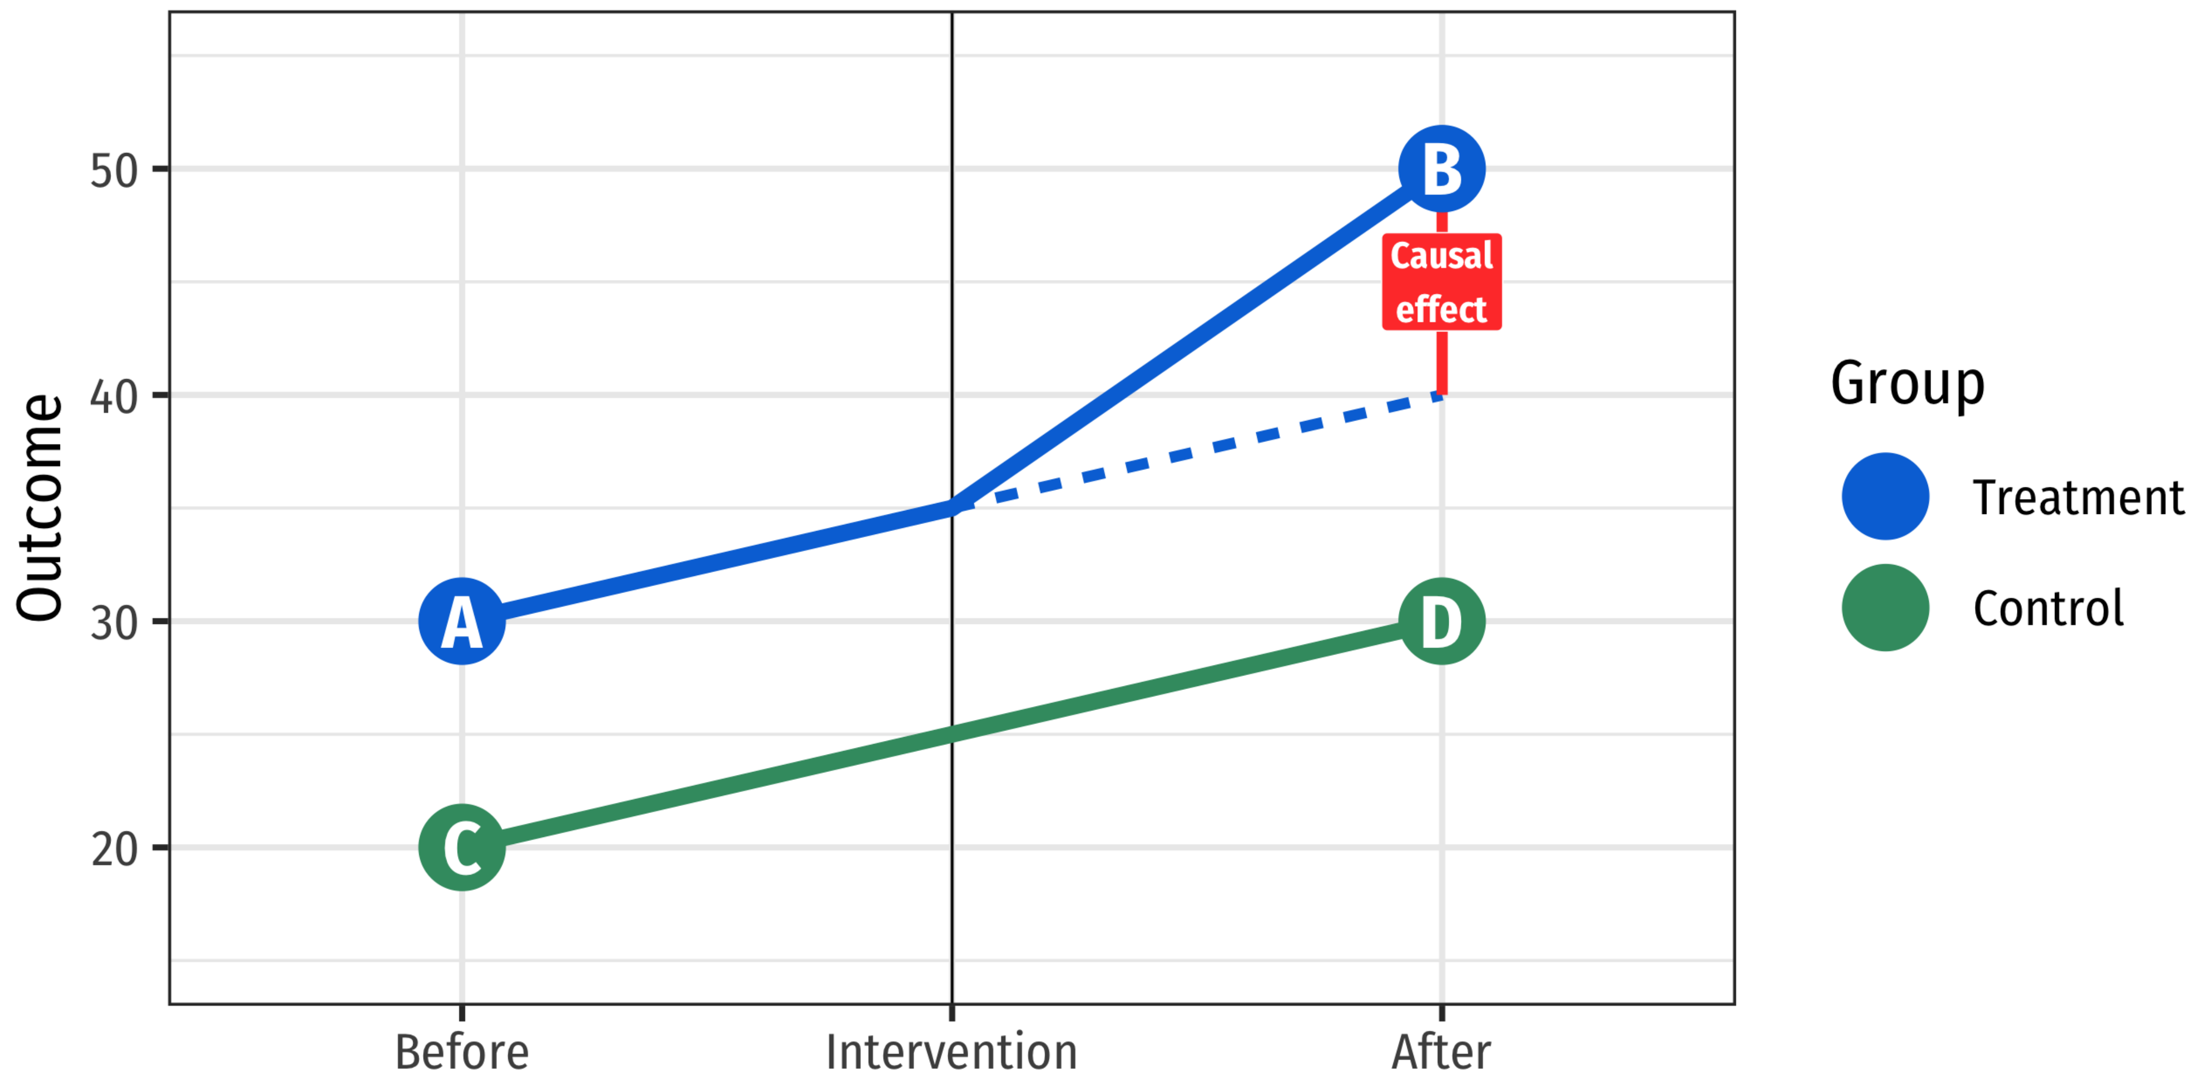

50
40
30
20
Outcome
B
Causal effect
A
C
D
Before
Intervention
After
Group
Treatment
Control

An easier way?

Finding all the group means is tedious!

What if there are other backdoors to worry about?

Regression to the rescue!

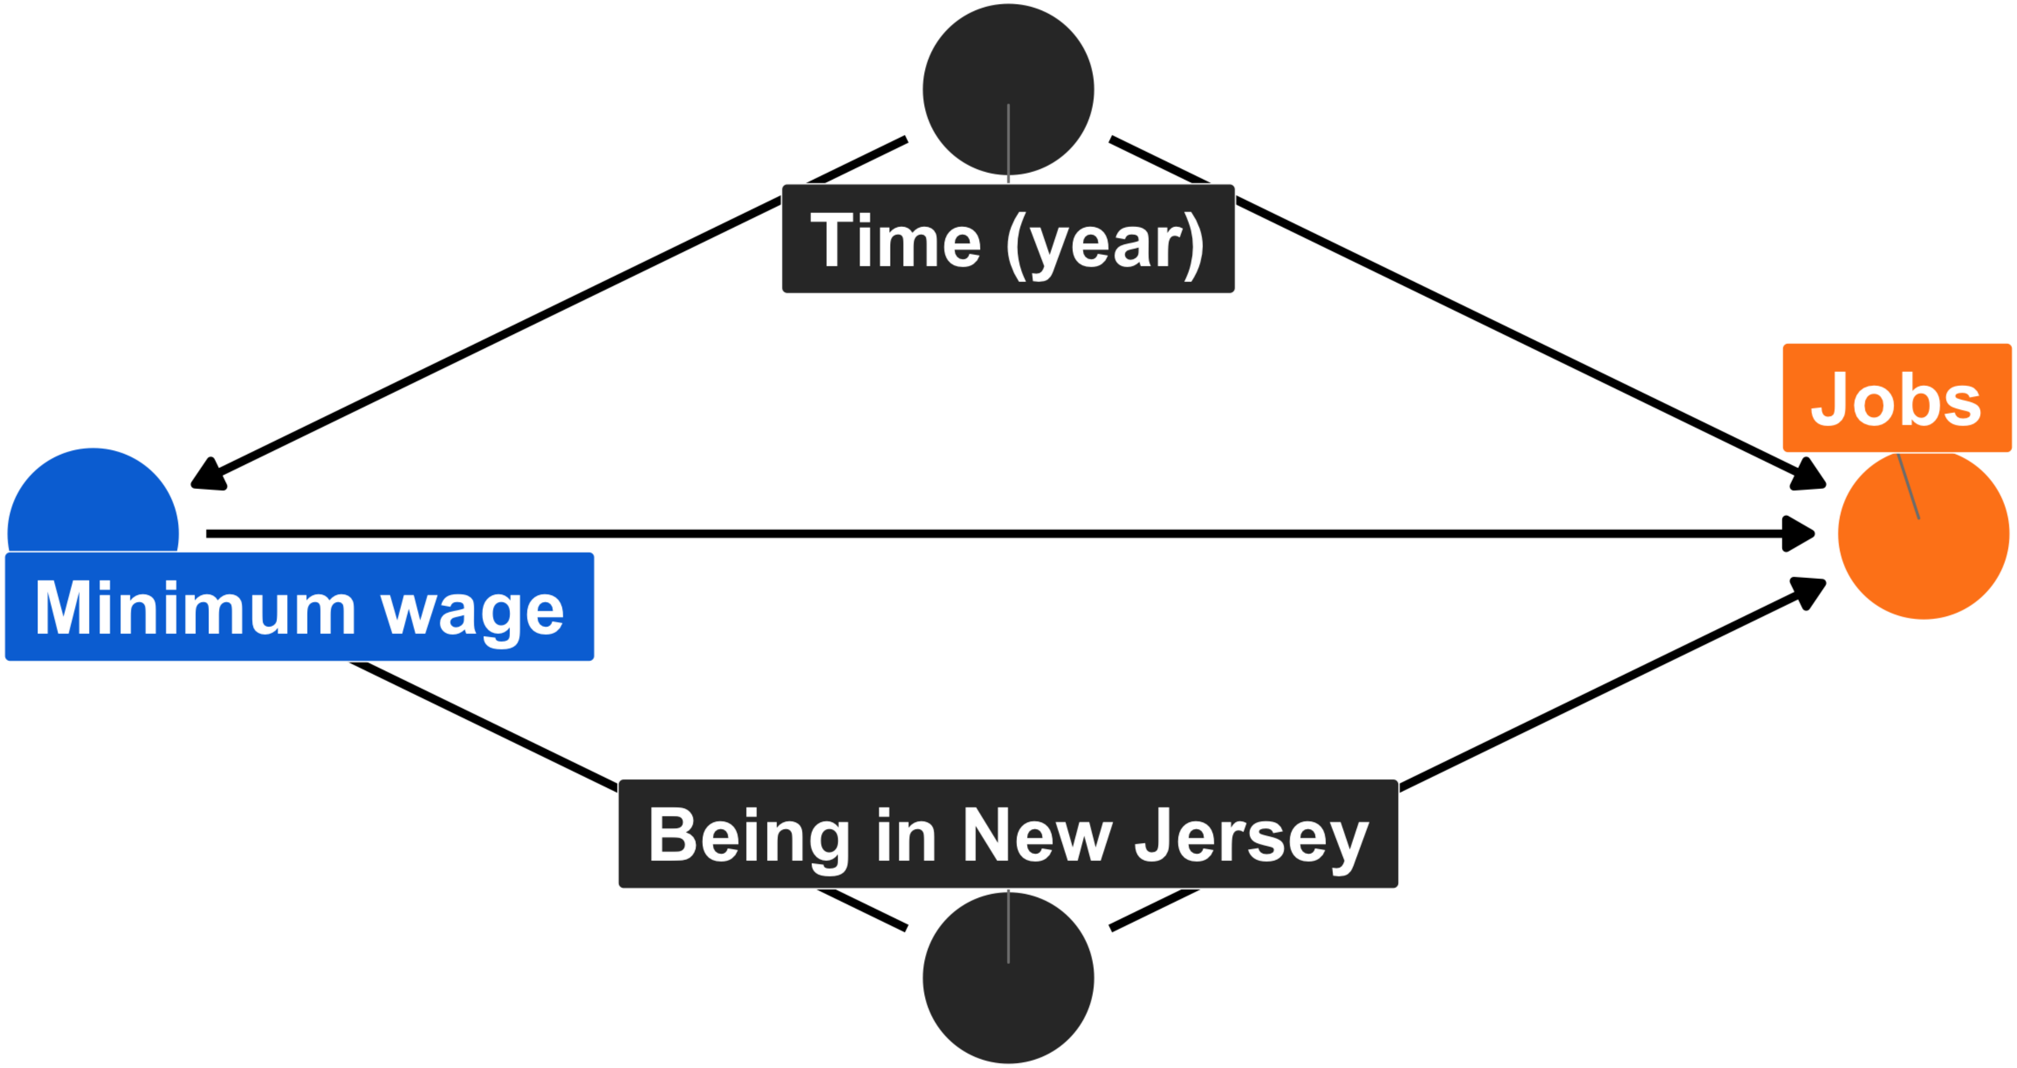
Time (year)
Jobs
Minimum wage
Being in New Jersey

$$
\begin{aligned}
Y_{it} = \; & \alpha + \beta \, \text{Group}_i + \gamma \, \text{Time}_t + \\
& \delta \, (\text{Group}_i \times \text{Time}_t) + \varepsilon_{it}
\end{aligned}
$$

```r
model <- lm(outcome ~ group + time + (group * time))
```

Group = 1 or TRUE if treatment

Time = 1 or TRUE if after

$$
\begin{aligned}
Y_{it} \; = \; & \alpha + \beta \, \text{Group}_i + \gamma \, \text{Time}_t + \\
& \delta \, (\text{Group}_i \times \text{Time}_t) + \varepsilon_{it}
\end{aligned}
$$

```
model <- lm(outcome ~ group + time + (group * time))
```

α = Mean of control, pre-treatment

β = Increase in outcome across groups

γ = Increase in outcome over time within units

δ = Difference in differences!

$$Y_{it} = \alpha + \beta\,\text{Group}_i + \gamma\,\text{Time}_t + \delta\,(\text{Group}_i \times \text{Time}_t) + \varepsilon_{it}$$

	Pre mean	Post mean	Δ (post − pre)
Control	α	α + γ	γ
Treatment	α + β	α + β + γ + δ	γ + δ
Δ (trtmt − ctrl)	β	β + δ	δ


```
hotdogs
```

```
## # A tibble: 4 × 3
##   price cheese chili
##   <dbl> <lgl>  <lgl>
## 1  2    FALSE  FALSE
## 2  2.35 TRUE   FALSE
## 3  2.35 FALSE  TRUE
## 4  2.7  TRUE   TRUE
```

```
model_hotdogs <-
  lm(price ~ cheese + chili +
       cheese * chili,
     data = hotdogs)
```

```
tidy(model_hotdogs)
```

```
## # A tibble: 4 × 2
##   term                 estimate
##   <chr>                   <dbl>
## 1 (Intercept)              2
## 2 cheeseTRUE               0.35
## 3 chiliTRUE                0.35
## 4 cheeseTRUE:chiliTRUE     0
```

Gotta catch'em all! Pokémon GO and physical activity among young adults: difference in differences study

Katherine B Howe,[1,2] Christian Suharlim,[3] Peter Ueda,[4,5] Daniel Howe, Ichiro Kawachi,[2] Eric B Rimm[1,6,7]

[1]Department of Epidemiology, Harvard TH Chan School of Public Health, Boston, MA, USA
[2]Department of Social and Behavioral Sciences, Harvard TH Chan School of Public Health, Boston, MA, USA
[3]Center for Health and Decision Science, Department of Health Policy and Management, Harvard TH Chan School of Public Health, Boston, MA, USA
[4]Department of Global Health and Population, Harvard TH Chan School of Public Health, Boston, MA, USA
[5]Clinical Epidemiology Unit, Department of Medicine, Solna,

ABSTRACT

OBJECTIVE
To estimate the effect of playing Pokémon GO on the number of steps taken daily up to six weeks after installation of the game.

DESIGN
Cohort study using online survey data.

PARTICIPANTS
Survey participants of Amazon Mechanical Turk (n=1182) residing in the United States, aged 18 to 35 years and using iPhone 6 series smartphones.

MAIN OUTCOME MEASURES
Number of daily steps taken each of the four weeks before and six weeks after installation of Pokémon

CONCLUSIONS
Pokémon GO was associated with an increase in the daily number of steps after installation of the game. The association was, however, moderate and no longer observed after six weeks.

Introduction

Pokémon GO is an augmented reality game in which players search real world locations for cartoon characters appearing on their smartphone screen. Since its launch in July 2016, the game has been downloaded over 500 million times worldwide.

Games that incentivise exercise might have the potential to promote and sustain physical activity habits.[1 2] Because walking is encouraged while playing,

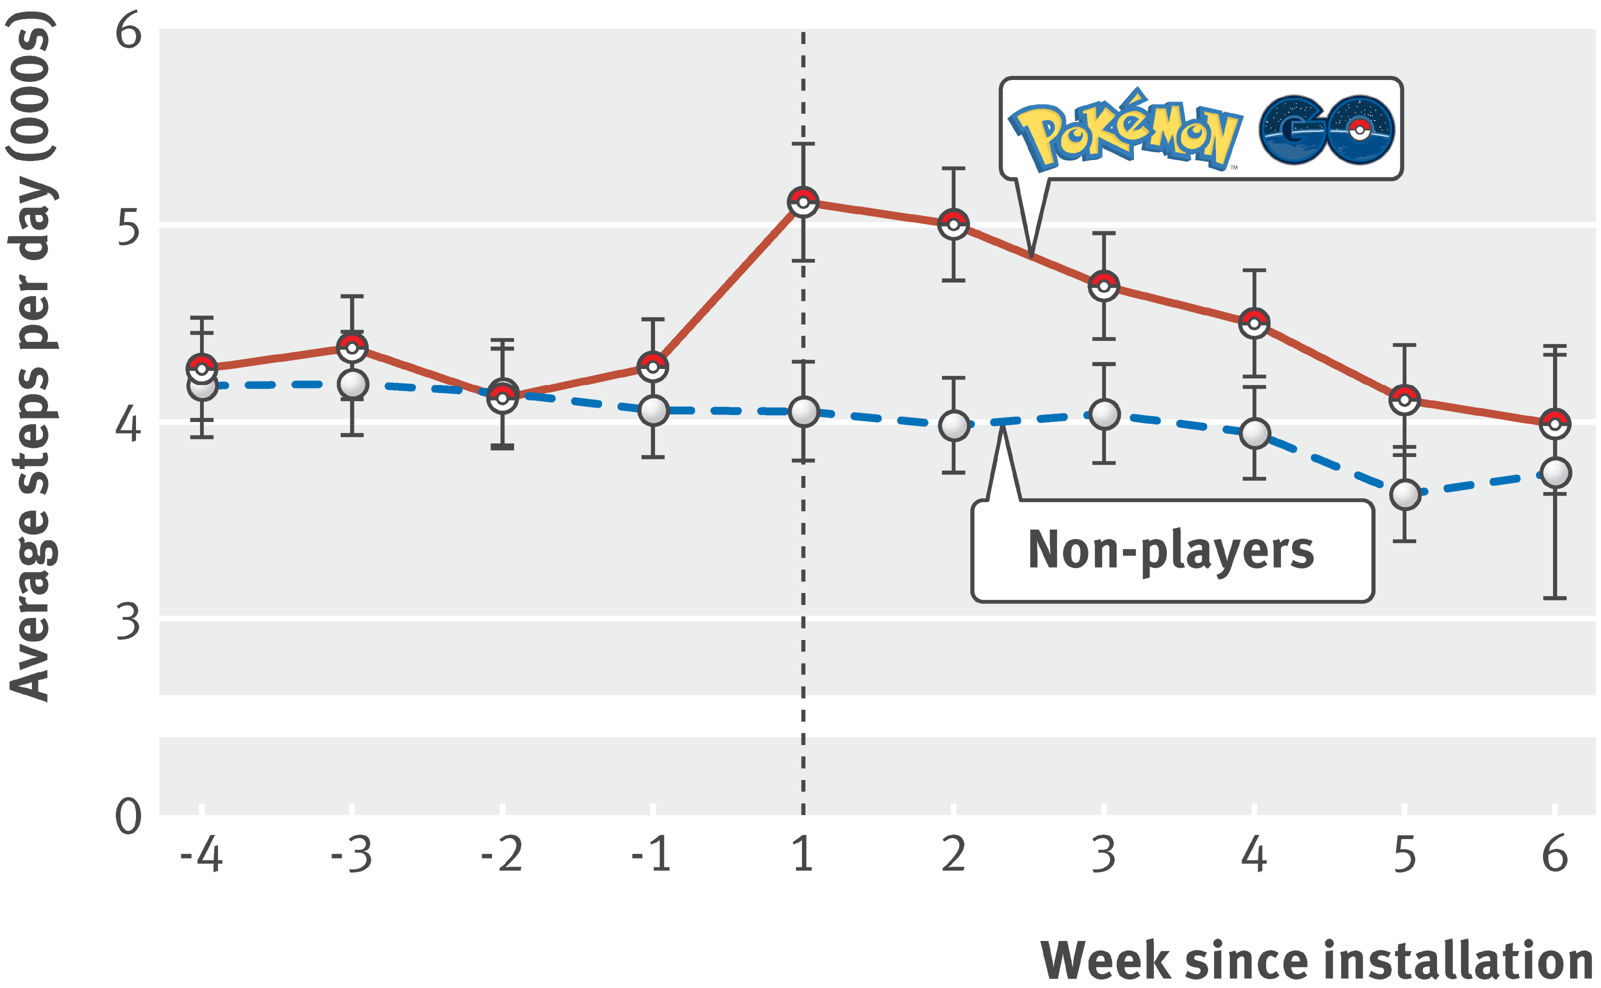
Average steps per day (000s)
6
5
4
3
0
-4
-3
-2
-1
1
2
3
4
5
6
Week since installation
Pokémon GO
Non-players

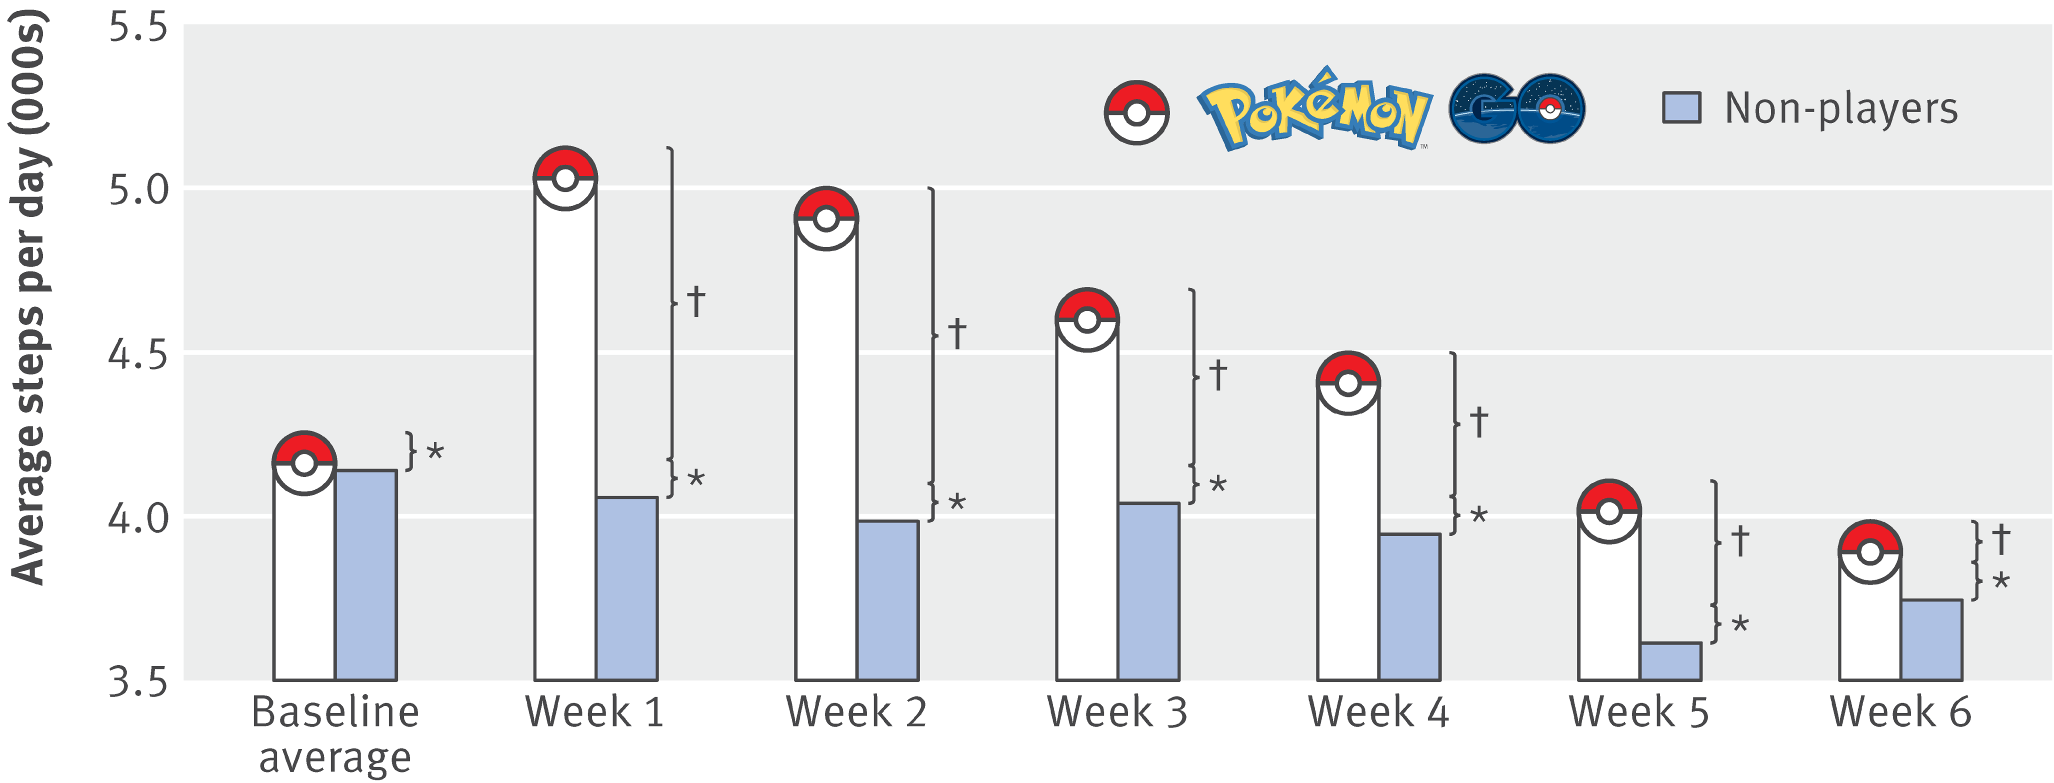

* Baseline difference 4 week average	† Difference	† Difference	† Difference	† Difference	† Difference	† Difference
114 (-212 to 440)	955 (697 to 1213)	906 (647 to 1164)	544 (280 to 808)	446 (169 to 722)	381 (43 to 720)	130 (-593 to 853)

Diff-in-diff assumptions

Assumptions

Parallel trends

Treatment and control groups might have different values at first, but we assume that the treatment group would have changed like the control group in the absence of treatment

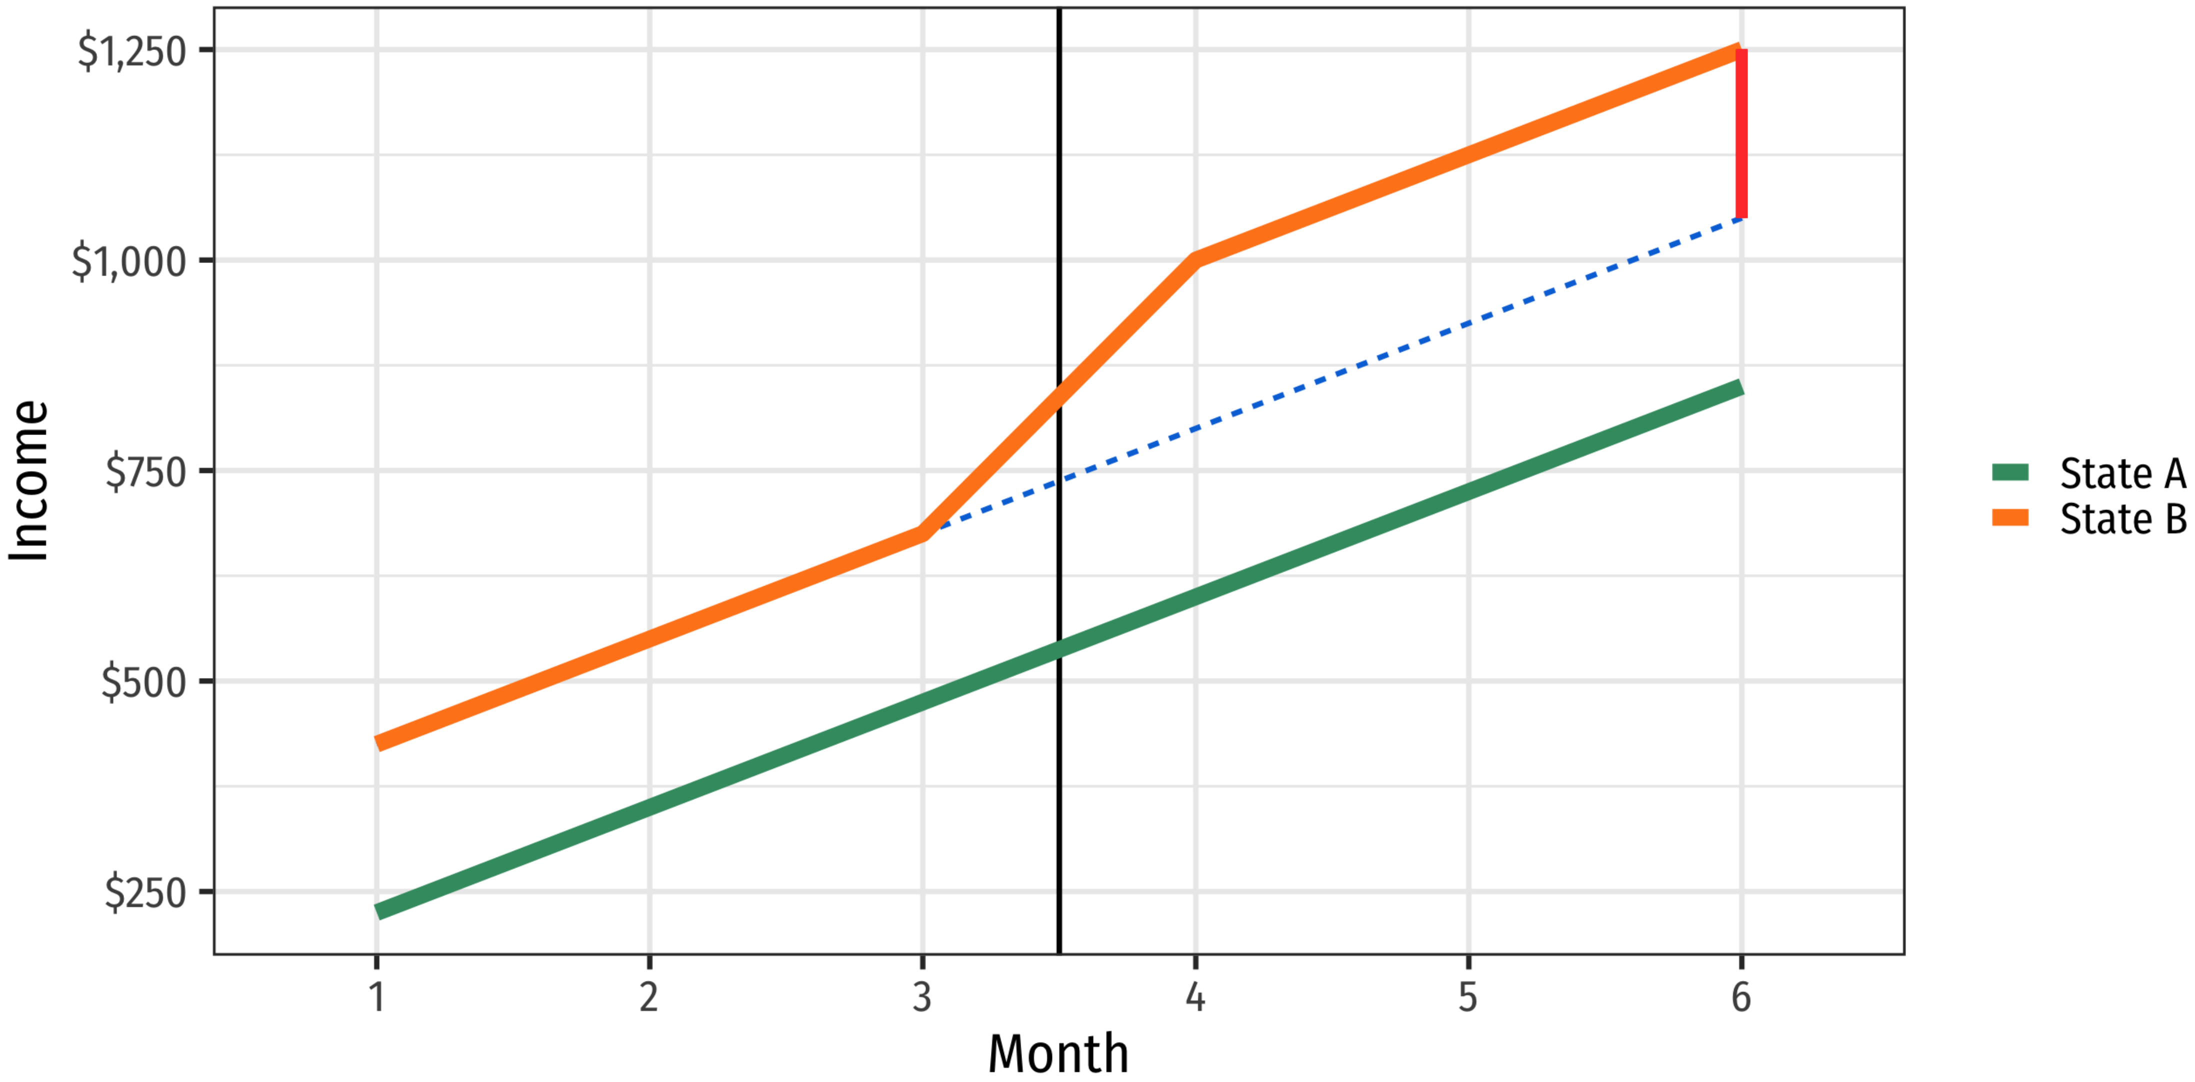

Income
$1,250
$1,000
$750
$500
$250
1
2
3
4
5
6
Month
State A
State B

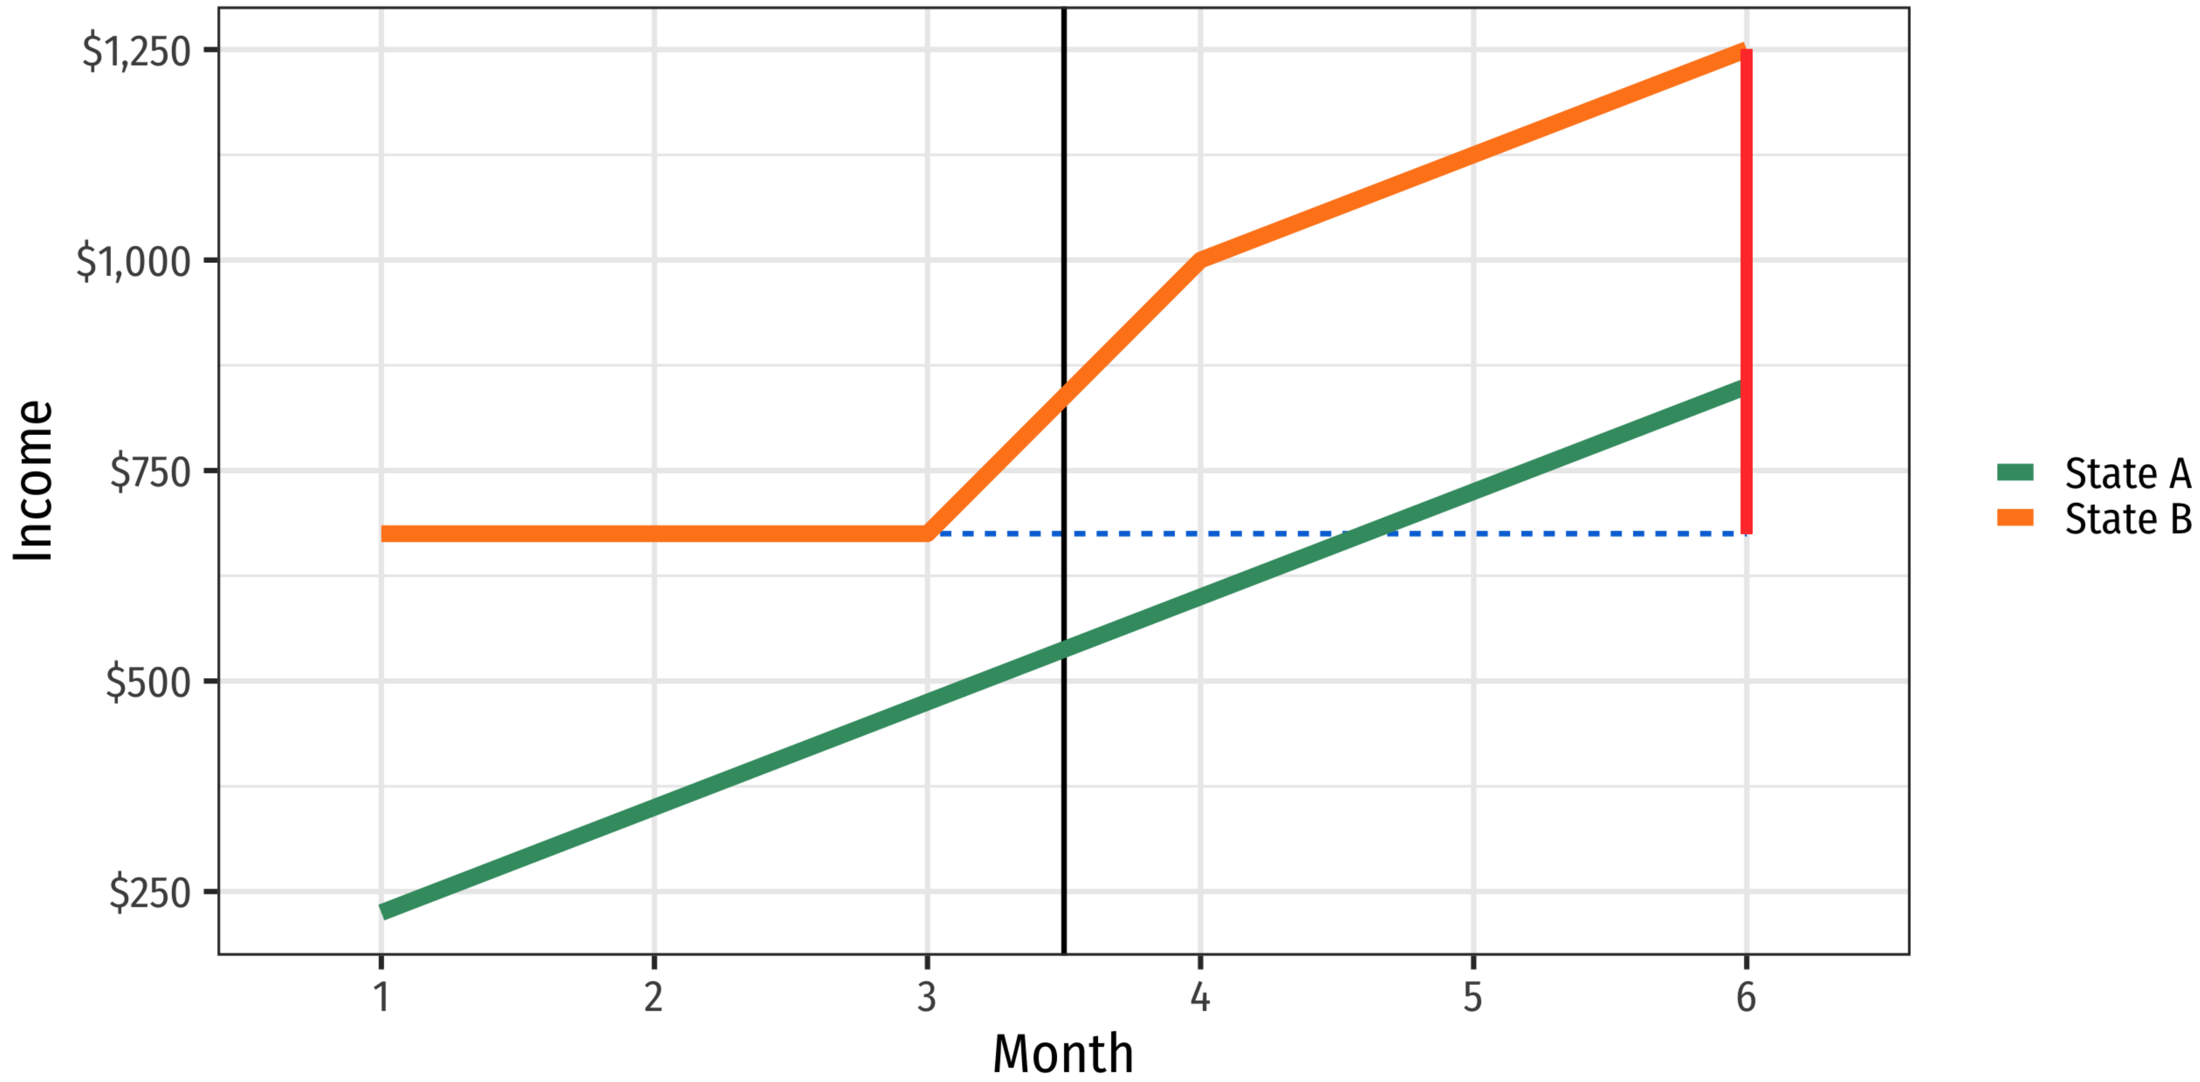
Income
$1,250
$1,000
$750
$500
$250
1
2
3
4
5
6
Month
State A
State B

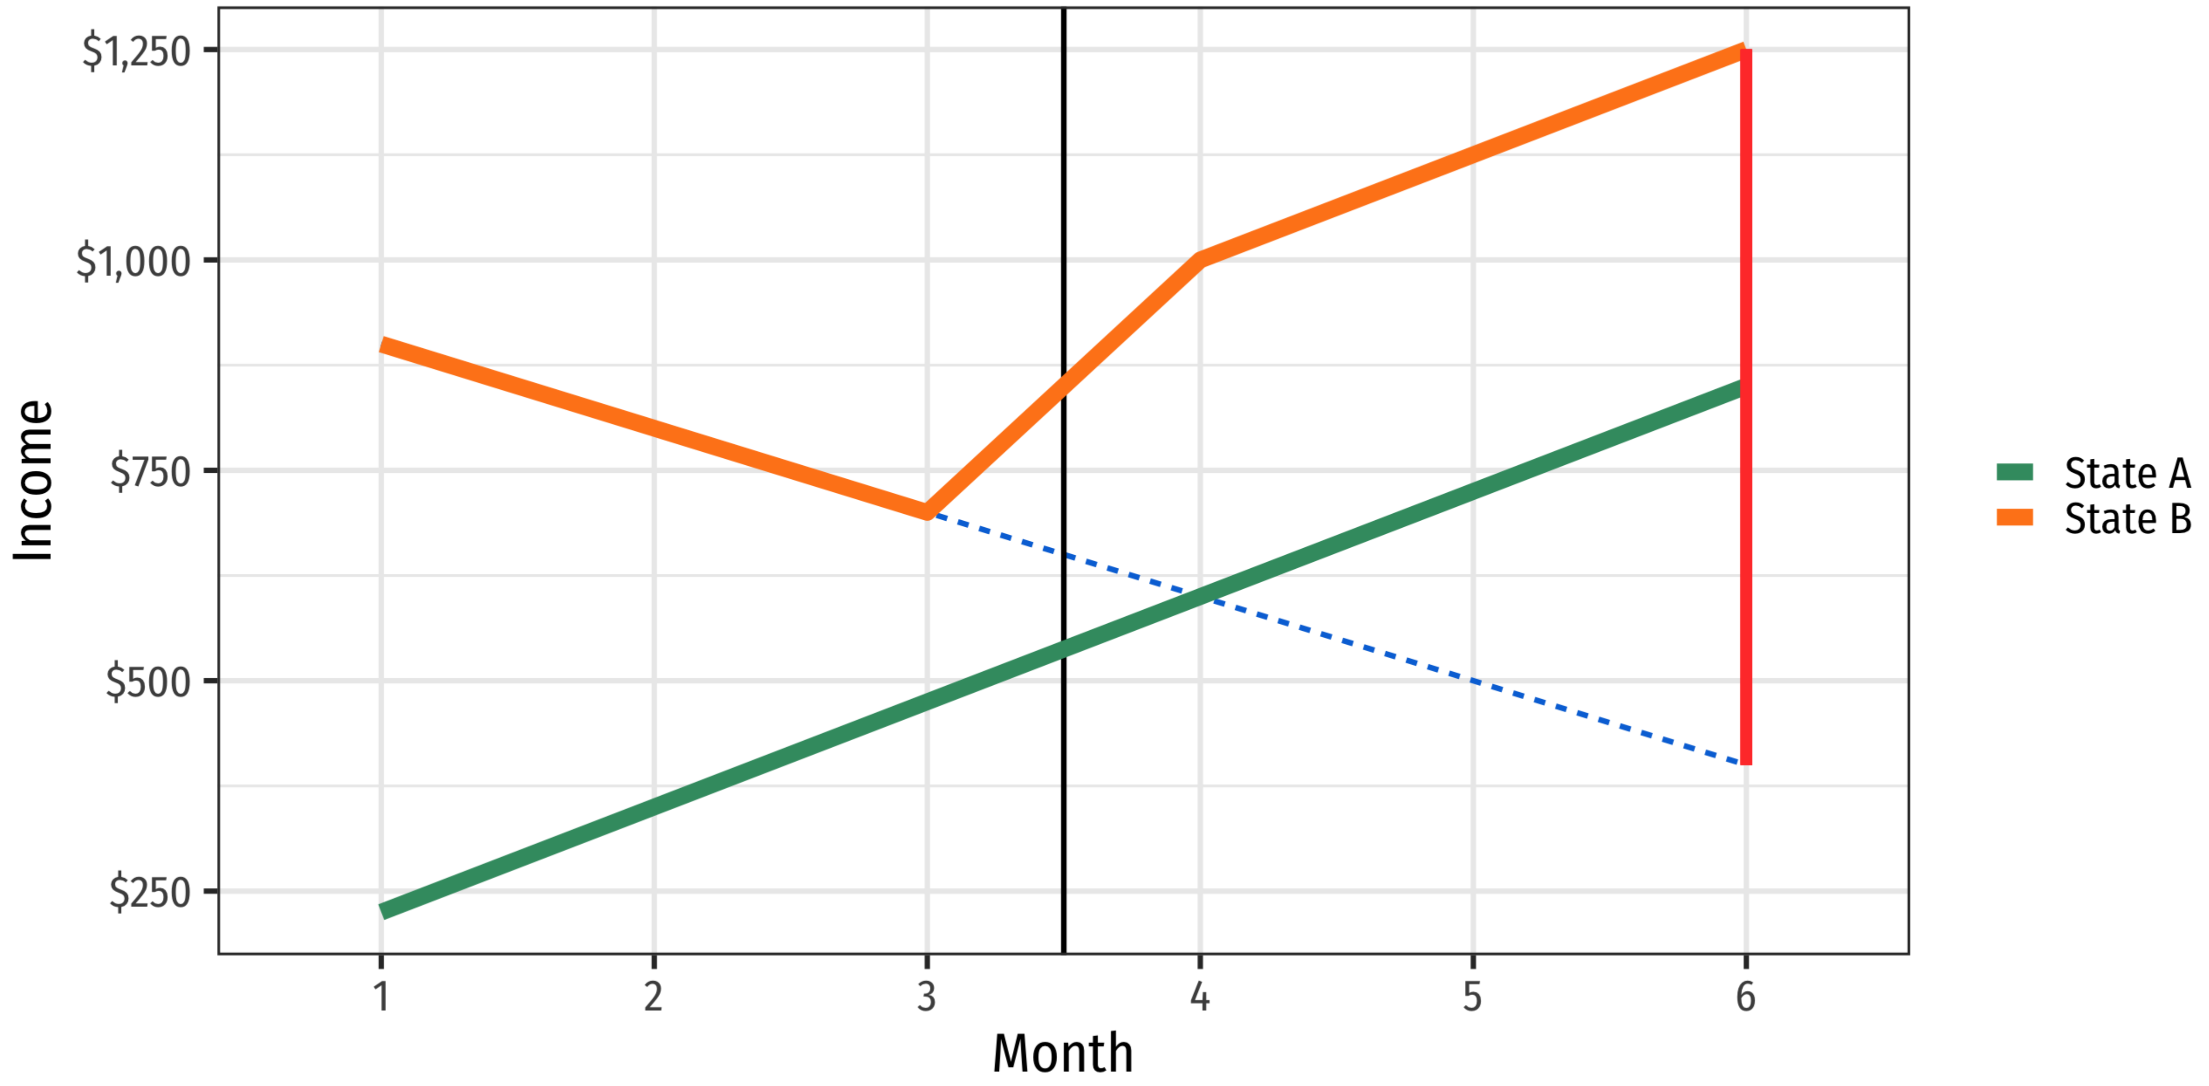

Income
$1,250
$1,000
$750
$500
$250
1
2
3
4
5
6
Month
State A
State B

Assumptions

Parallel trends

Check by pretending the treatment happened earlier; if there's an effect, there's likely an underlying trend

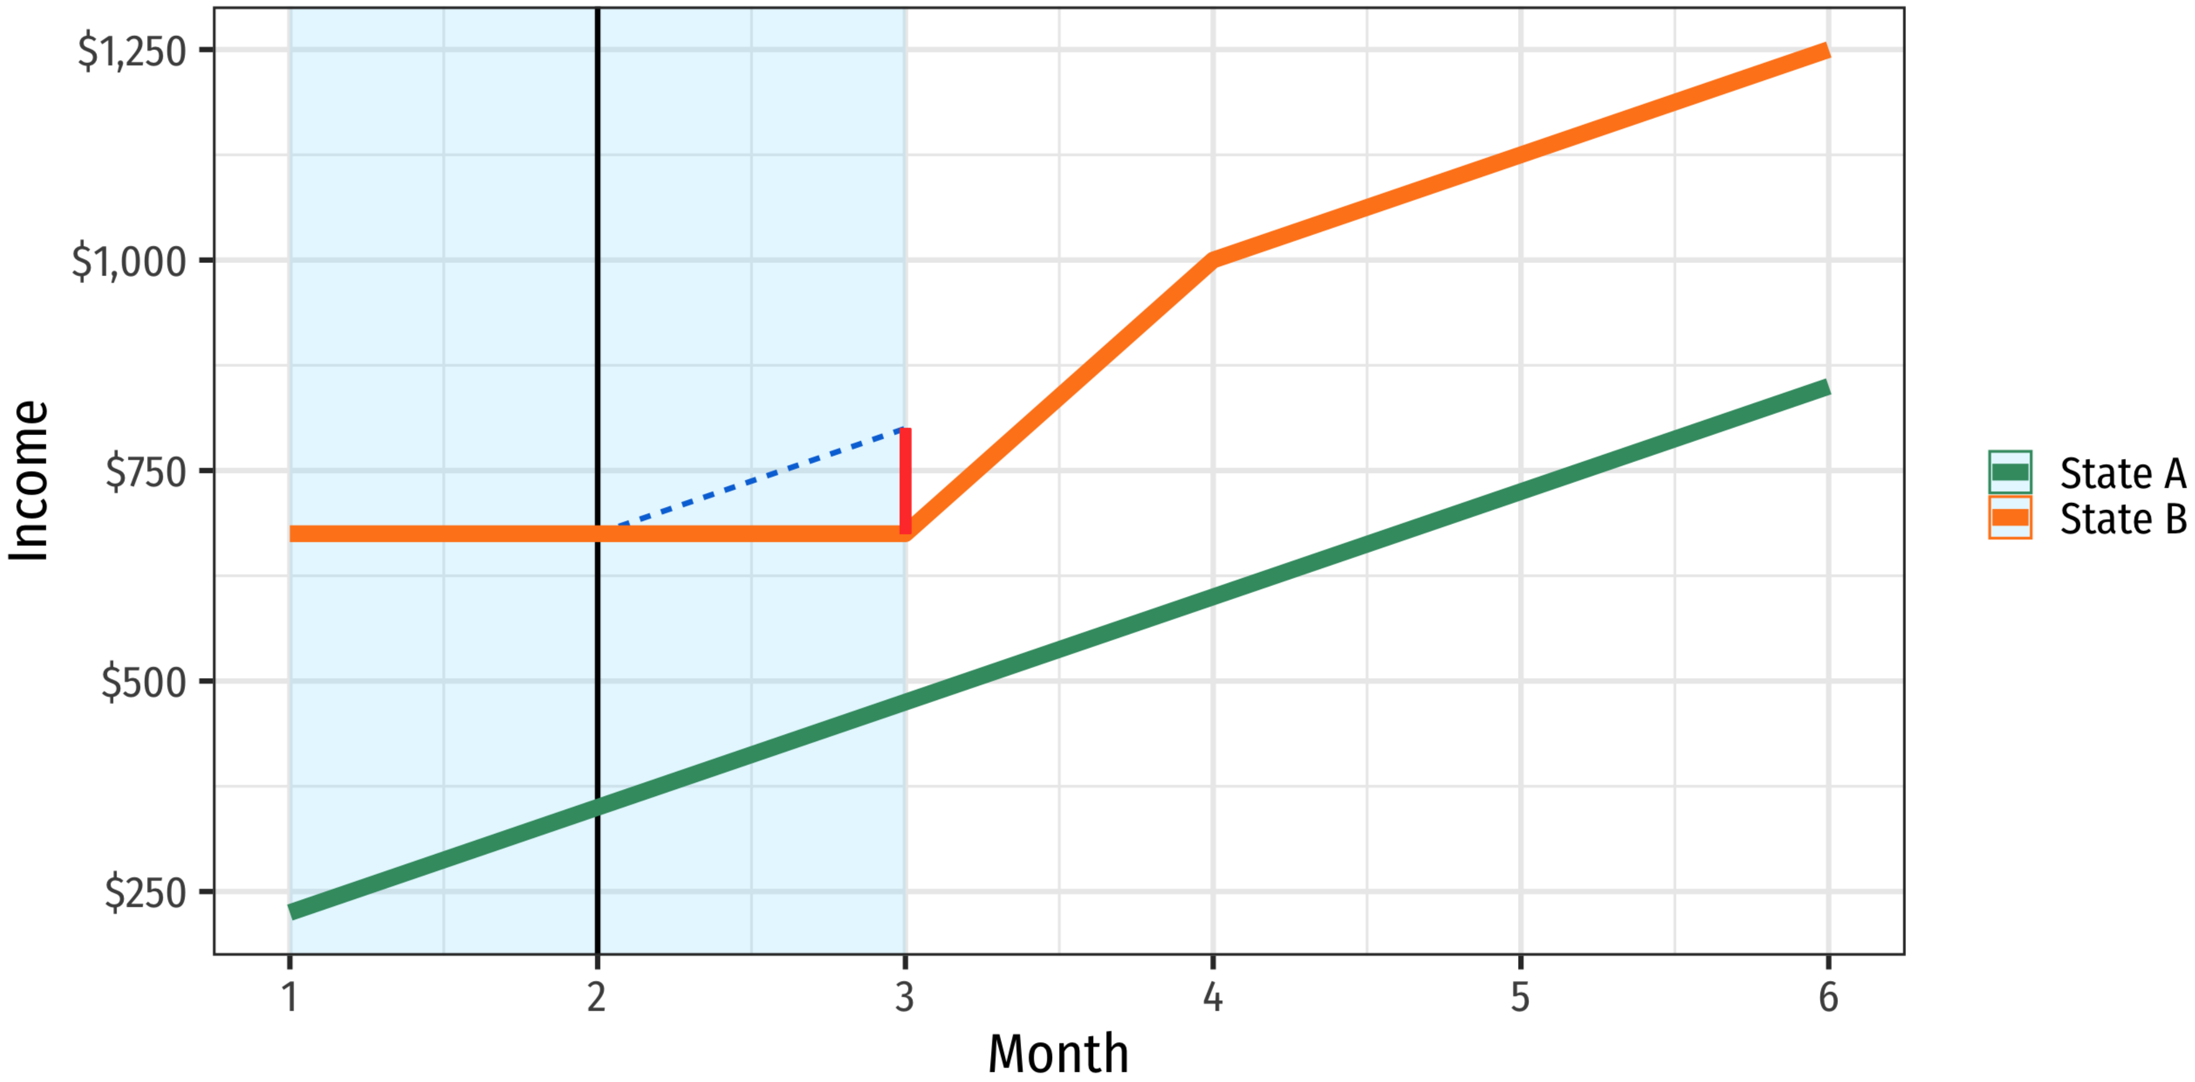

Income
$1,250
$1,000
$750
$500
$250
1
2
3
4
5
6
Month
State A
State B

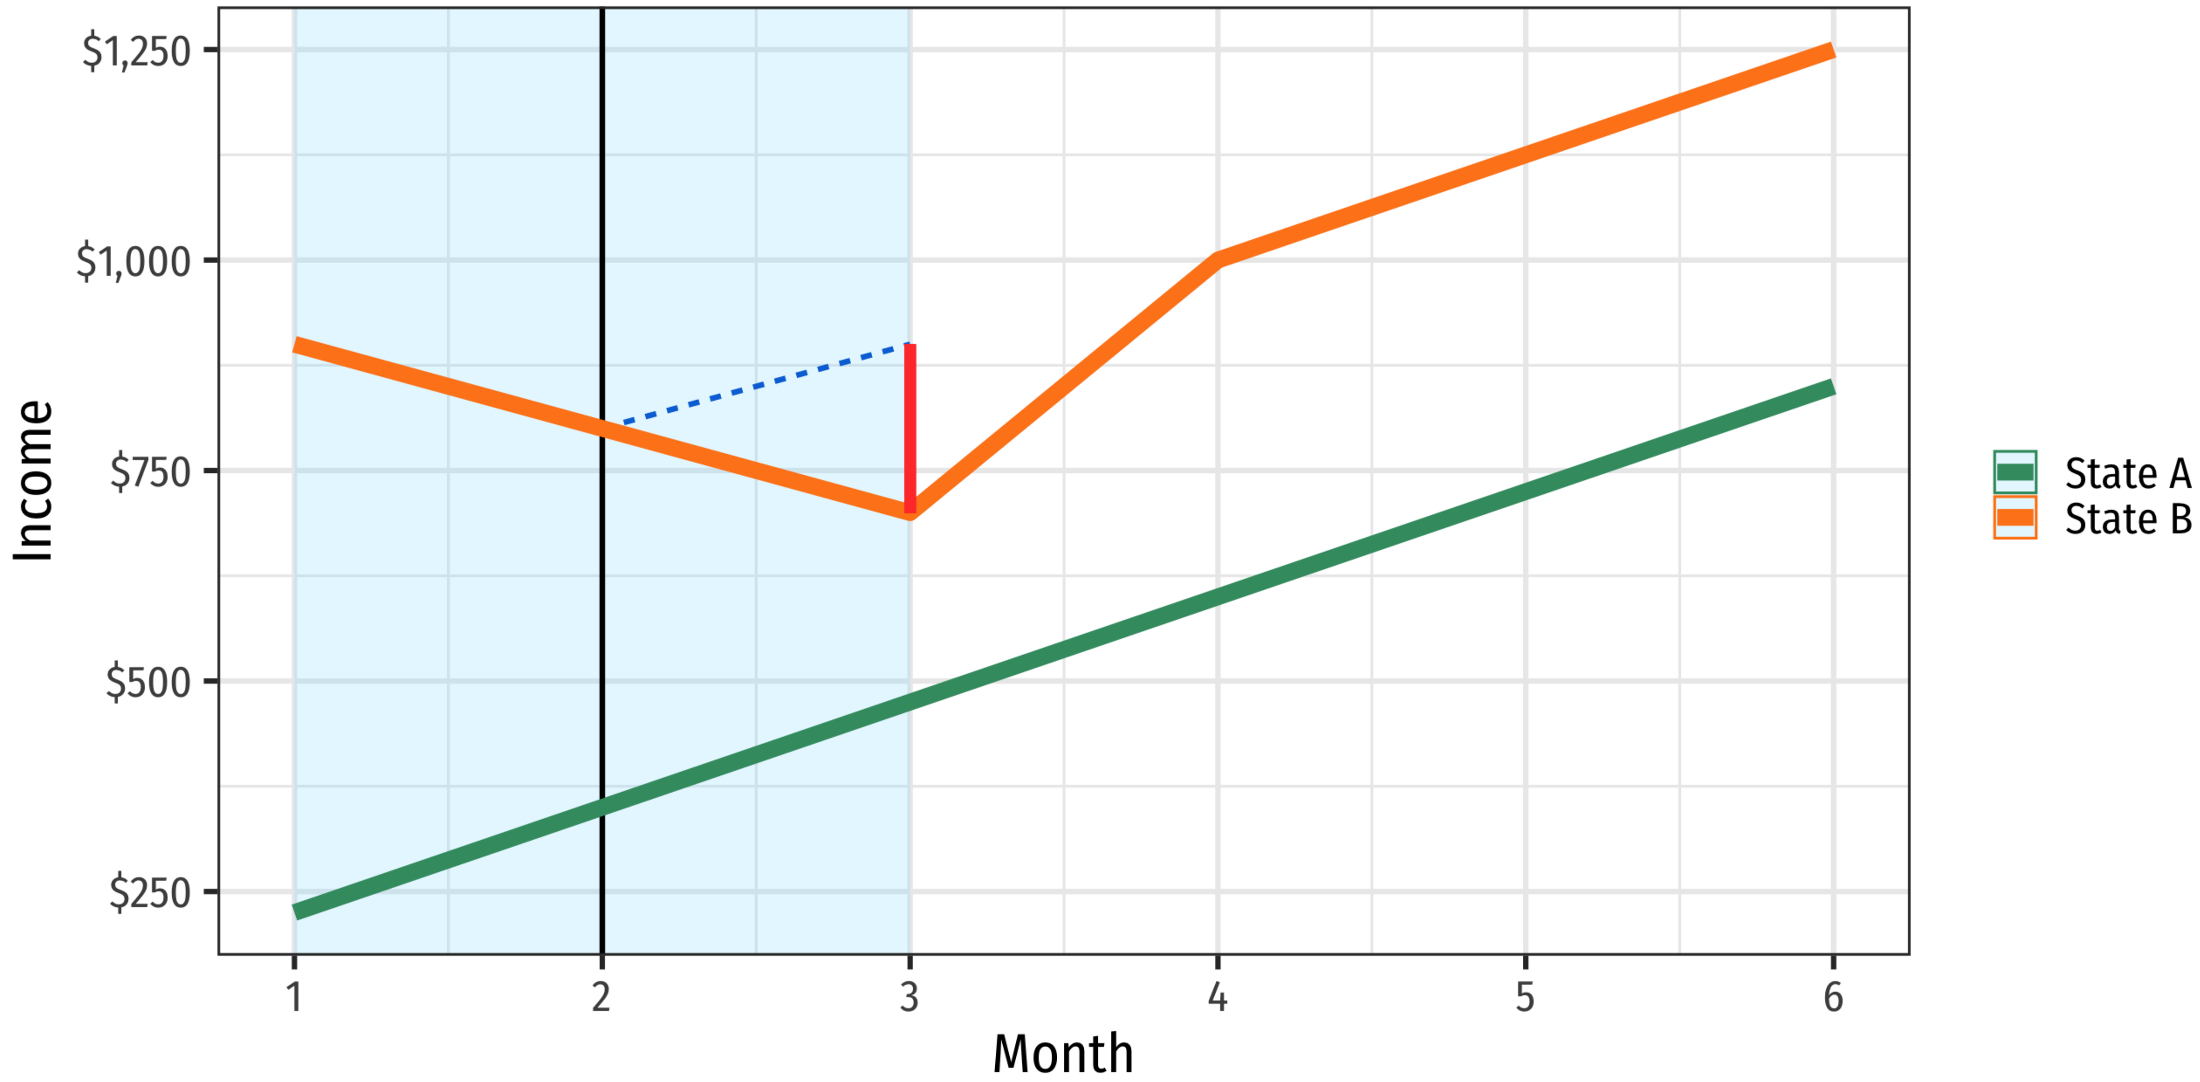

$1,250
$1,000
$750
$500
$250
Income
1
2
3
4
5
6
Month
State A
State B

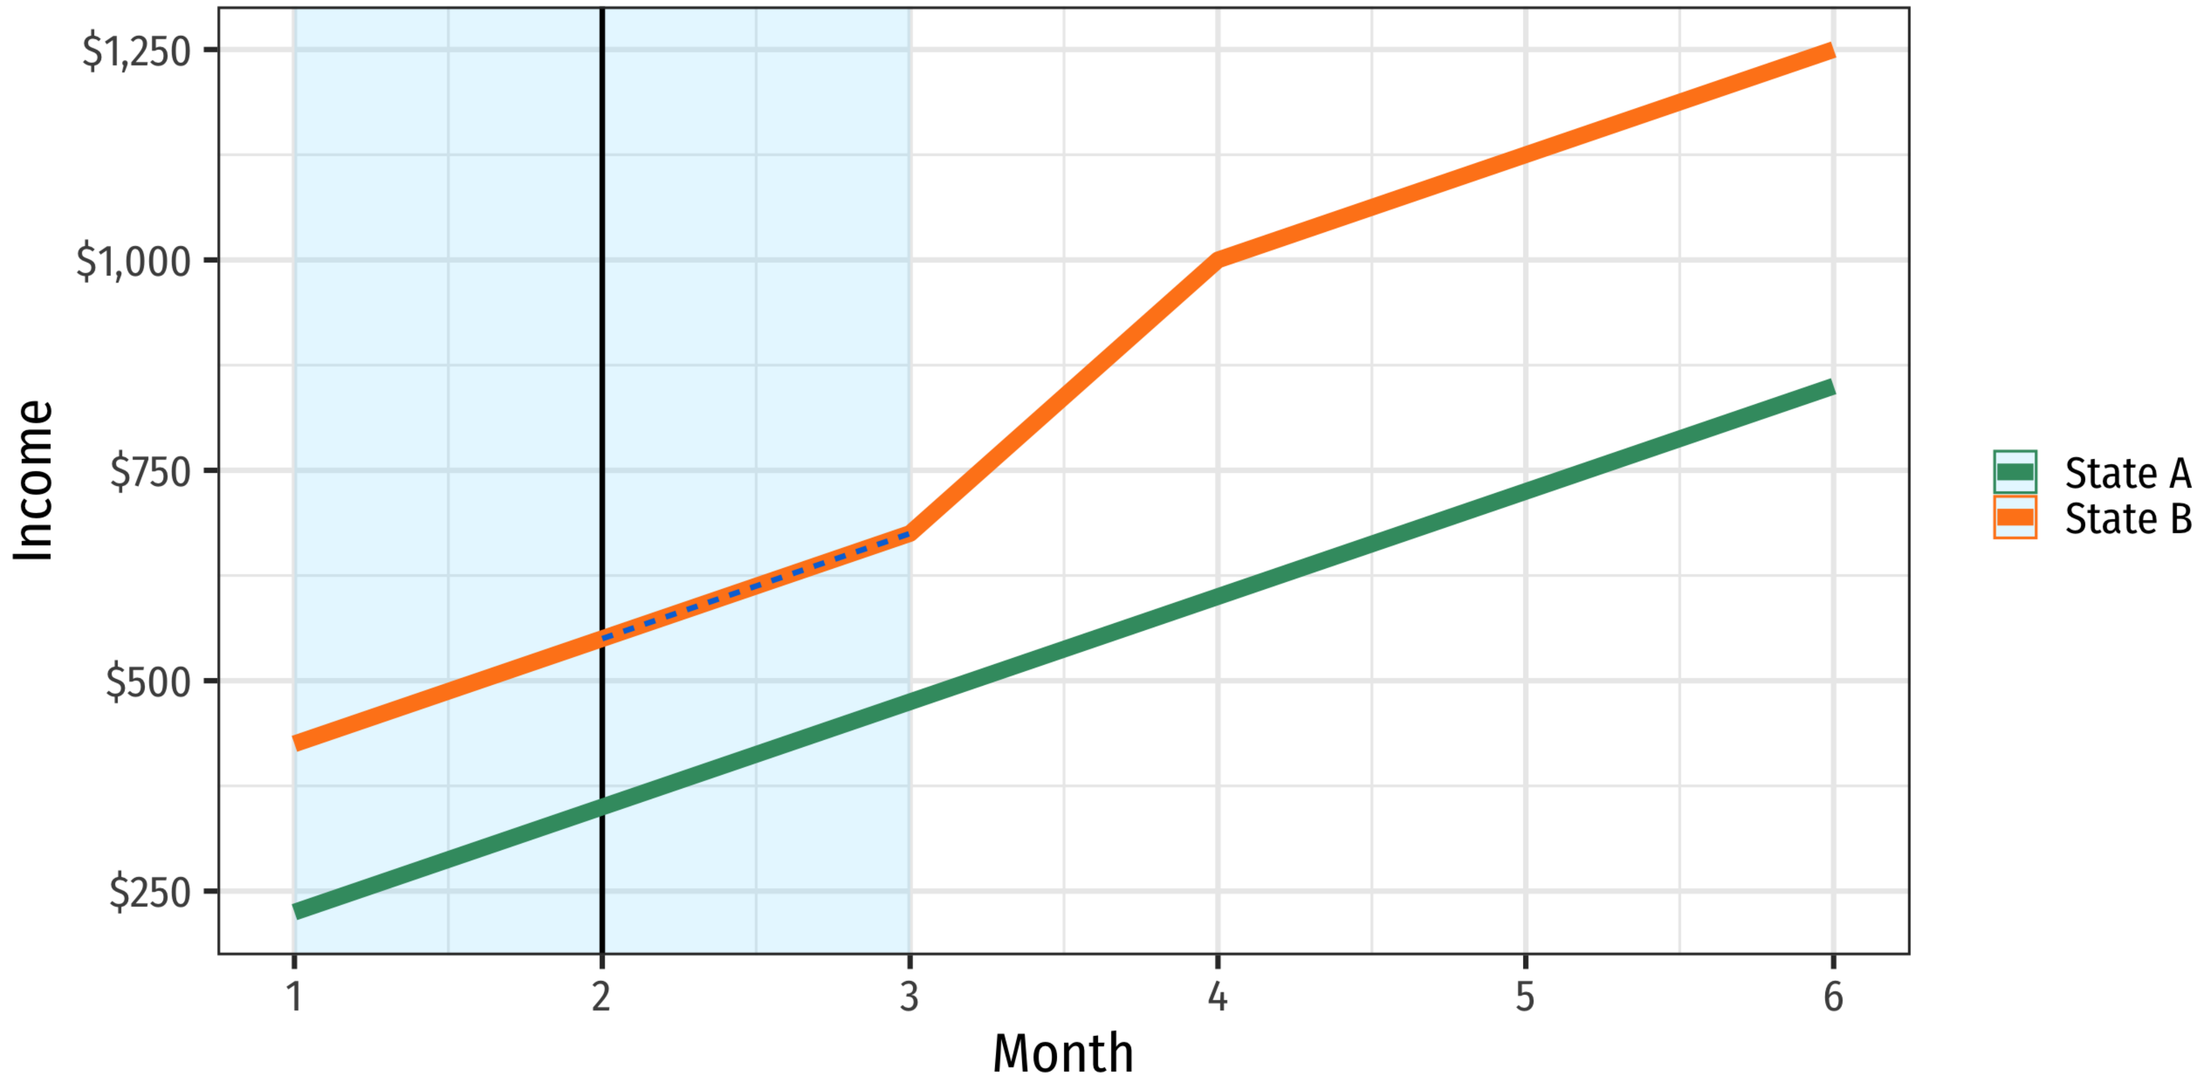
$1,250
$1,000
$750
$500
$250
Income
1
2
3
4
5
6
Month
State A
State B

Assumptions

Treatment timing

Units often receive treatment at different times, which can distort your estimate!

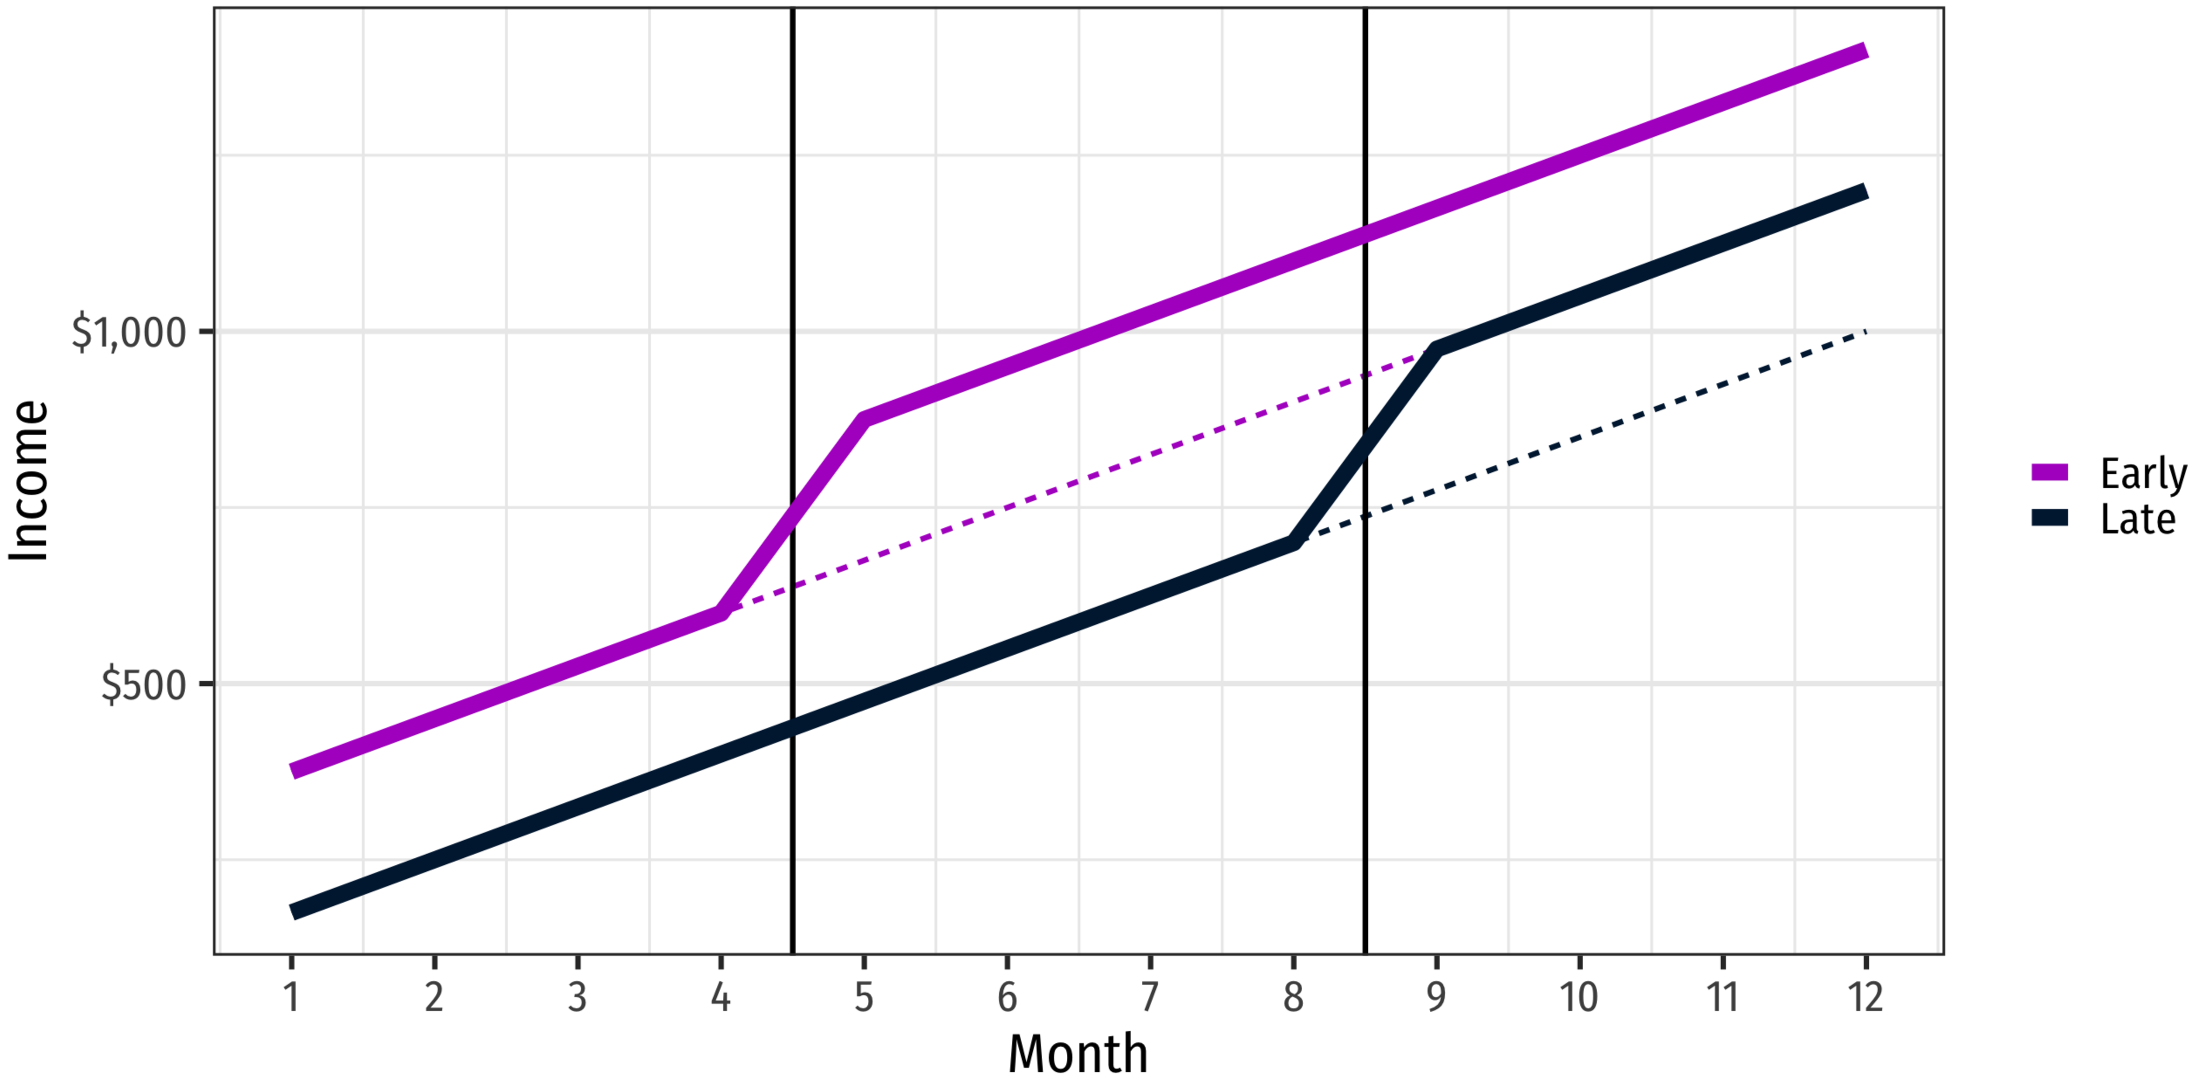
Income
$1,000
$500
1
2
3
4
5
6
7
8
9
10
11
12
Month
Early
Late

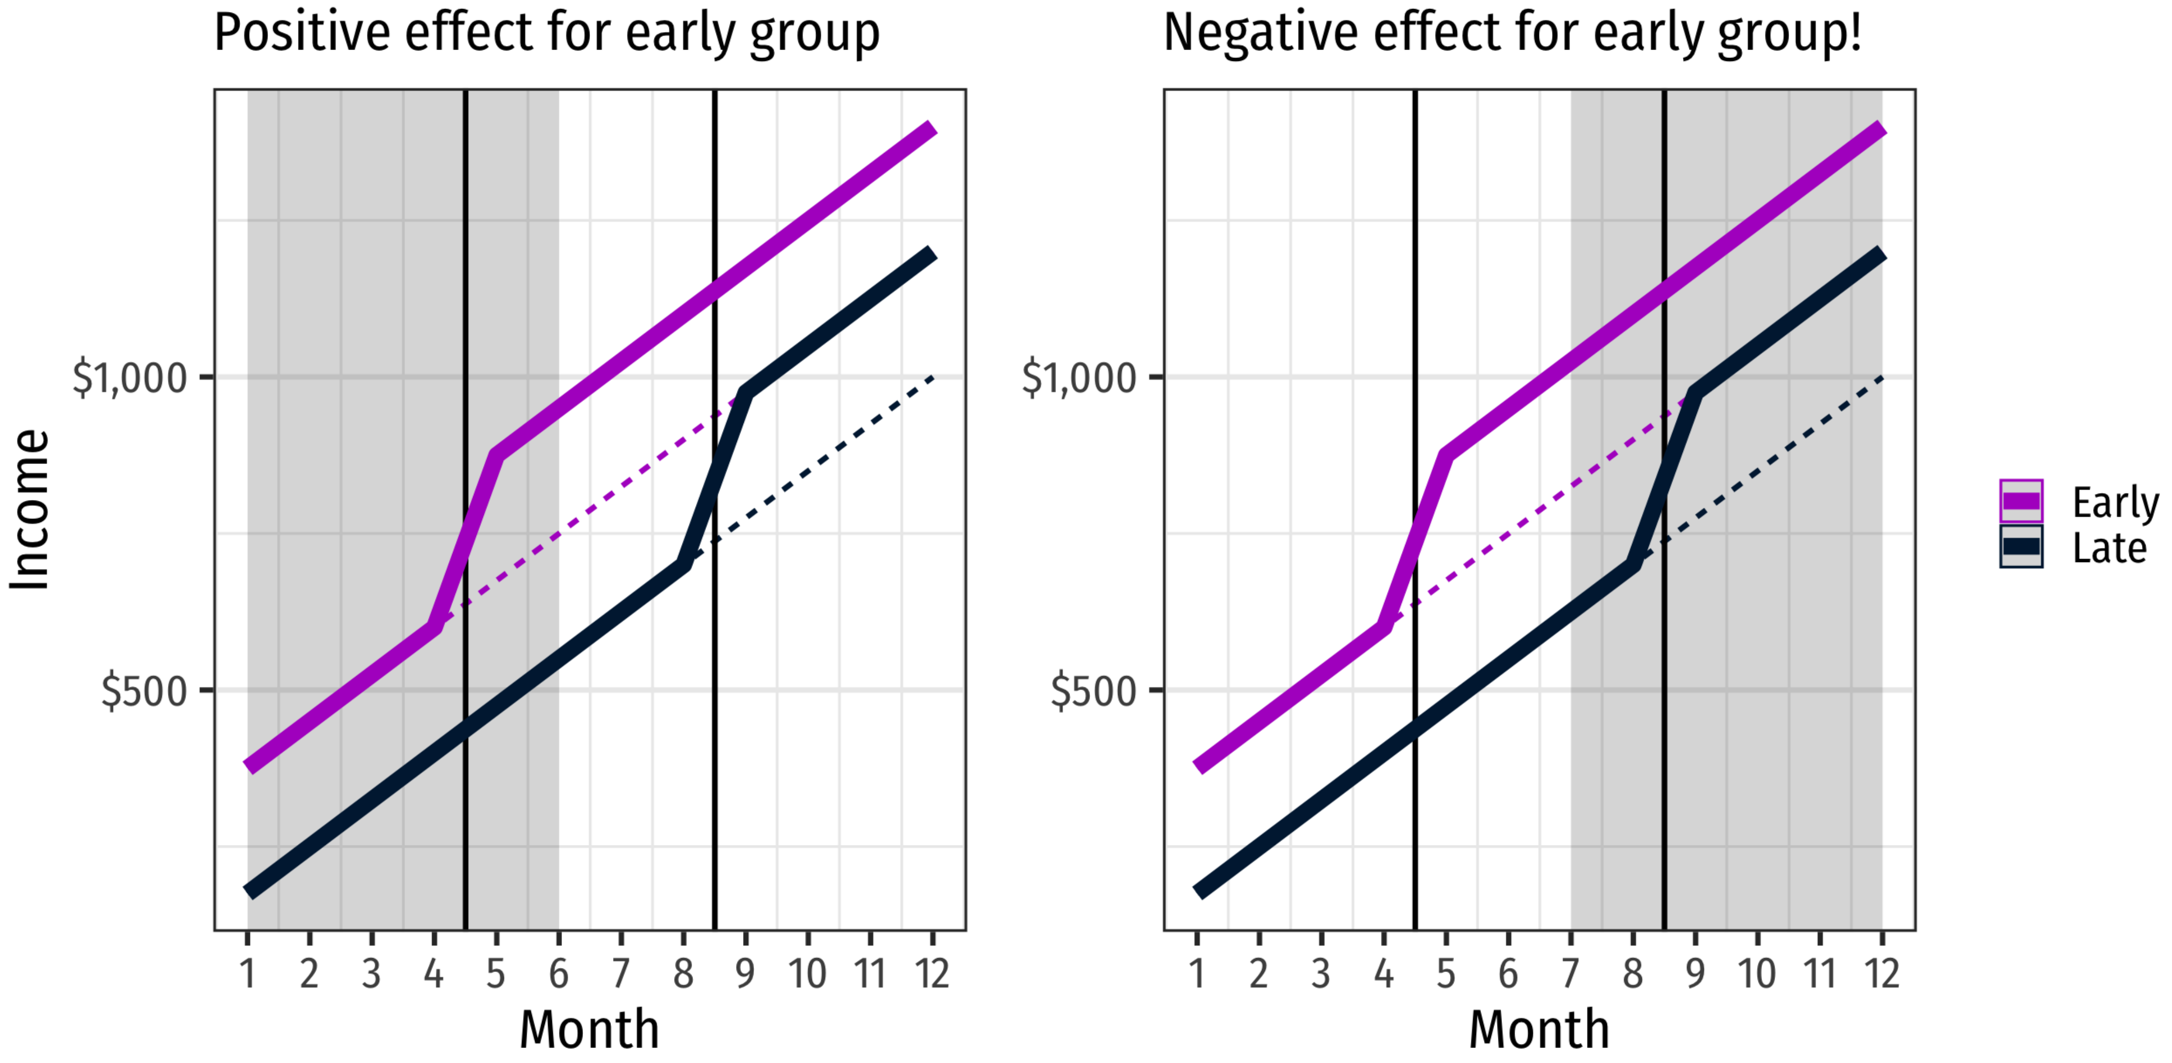
Positive effect for early group
Negative effect for early group!
Income
$1,000
$500
1 2 3 4 5 6 7 8 9 10 11 12
Month
Early
Late

Assumptions

You can check how big of an issue this is with Goodman-Bacon decomposition

R package: `bacondecomp`

DIFFERENCE-IN-DIFFERENCES WITH VARIATION IN TREATMENT TIMING*

Andrew Goodman-Bacon

July 2019

Abstract: The canonical difference-in-differences (DD) estimator contains two time periods, "pre" and "post", and two groups, "treatment" and "control". Most DD applications, however, exploit variation across groups of units that receive treatment at different times. This paper shows that the general estimator equals a weighted average of all possible two-group/two-period DD estimators in the data. This defines the DD estimand and identifying assumption, a generalization of common trends. I discuss how to interpret DD estimates and propose a new balance test. I show how to decompose the difference between two specifications, and provide a new analysis of models that include time-varying controls.