Photo credit: Dave DiCello

17-803 Empirical Methods

Bogdan Vasilescu, Institute for Software Research

Diff-in-Diff & CausalImpact

Tuesday, April 13, 2021

Plan for Today

- Diff in diff model
- Causal Impact (see https://youtu.be/GTgZfCltMm8)

Read this if you get a chance –>

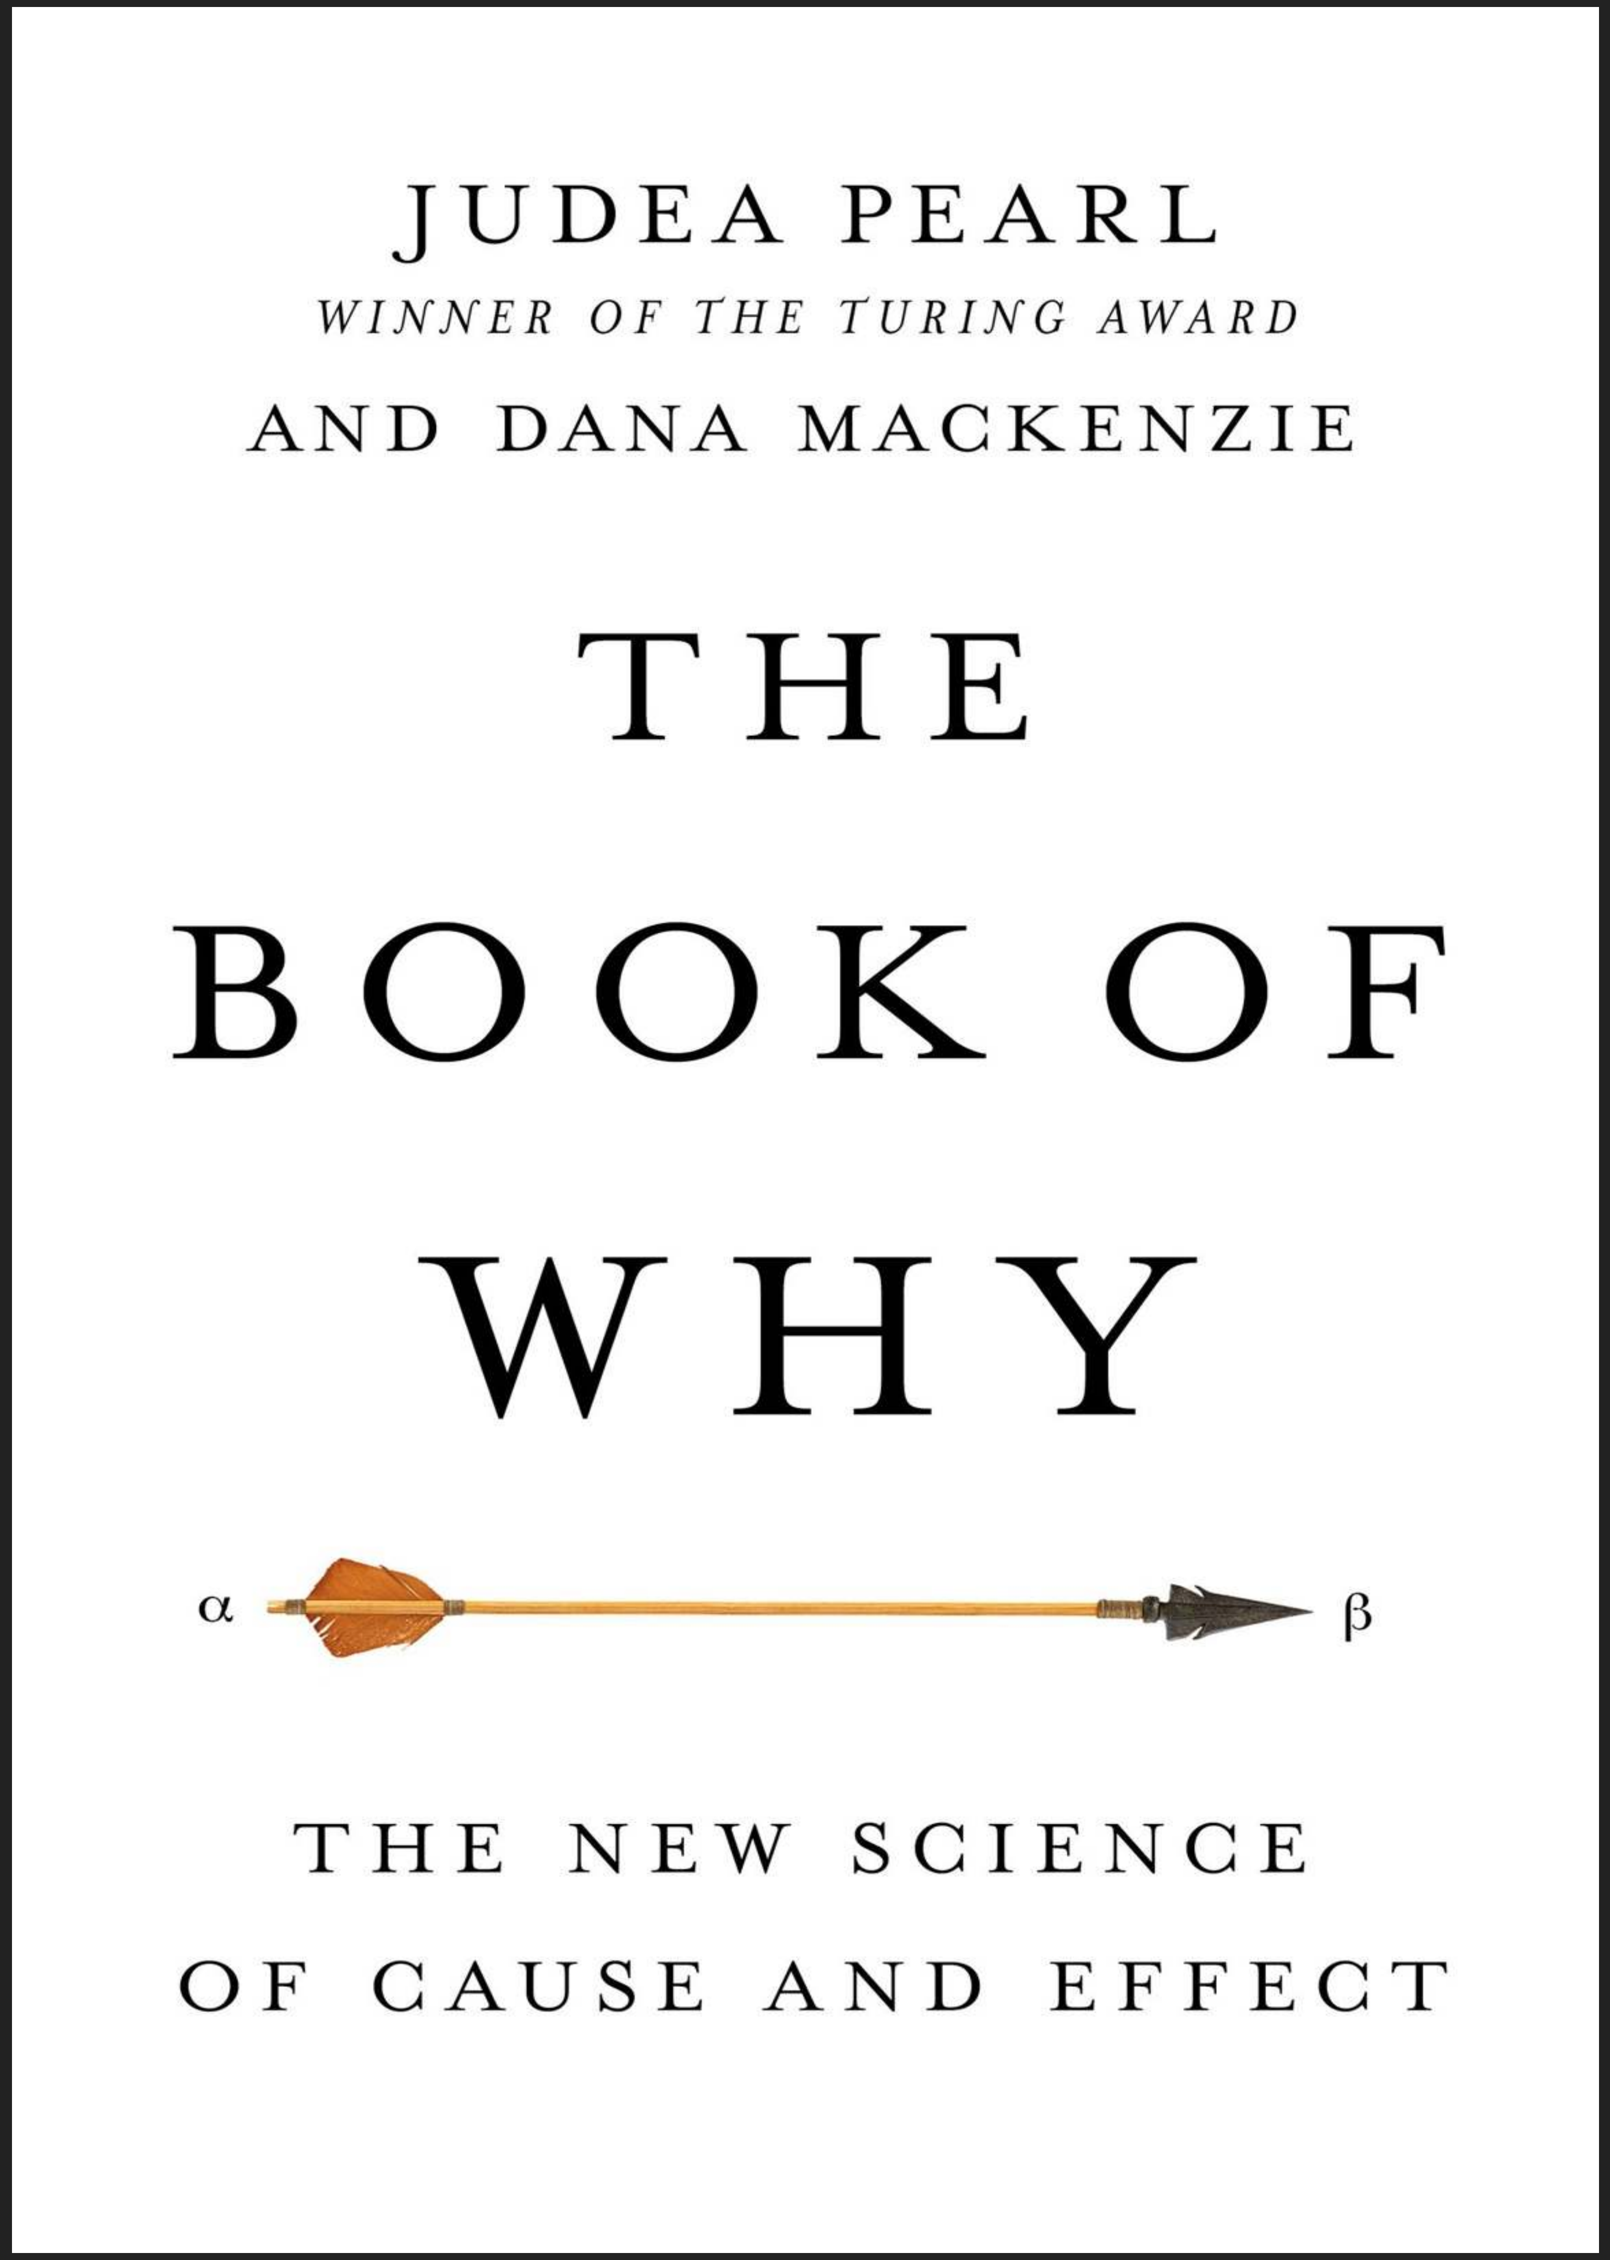

Diff-in-diff

Slides from: Rajendran, P. Causal Inference using Difference in Differences, Causal Impact, and Synthetic Control. https://towardsdatascience.com/causal-inference-using-difference-in-differences-causal-impact-and-synthetic-control-f8639c408268

The Key Concept

- The difference-in-difference (diff-in-diff) is a powerful model which allows us to look at the effect of an intervention by taking into consideration:
  - how a group mean changes before and after a policy intervention (treatment group) AND
  - compare this change with the mean over time of a similar group which did not undergo the treatment (control group).

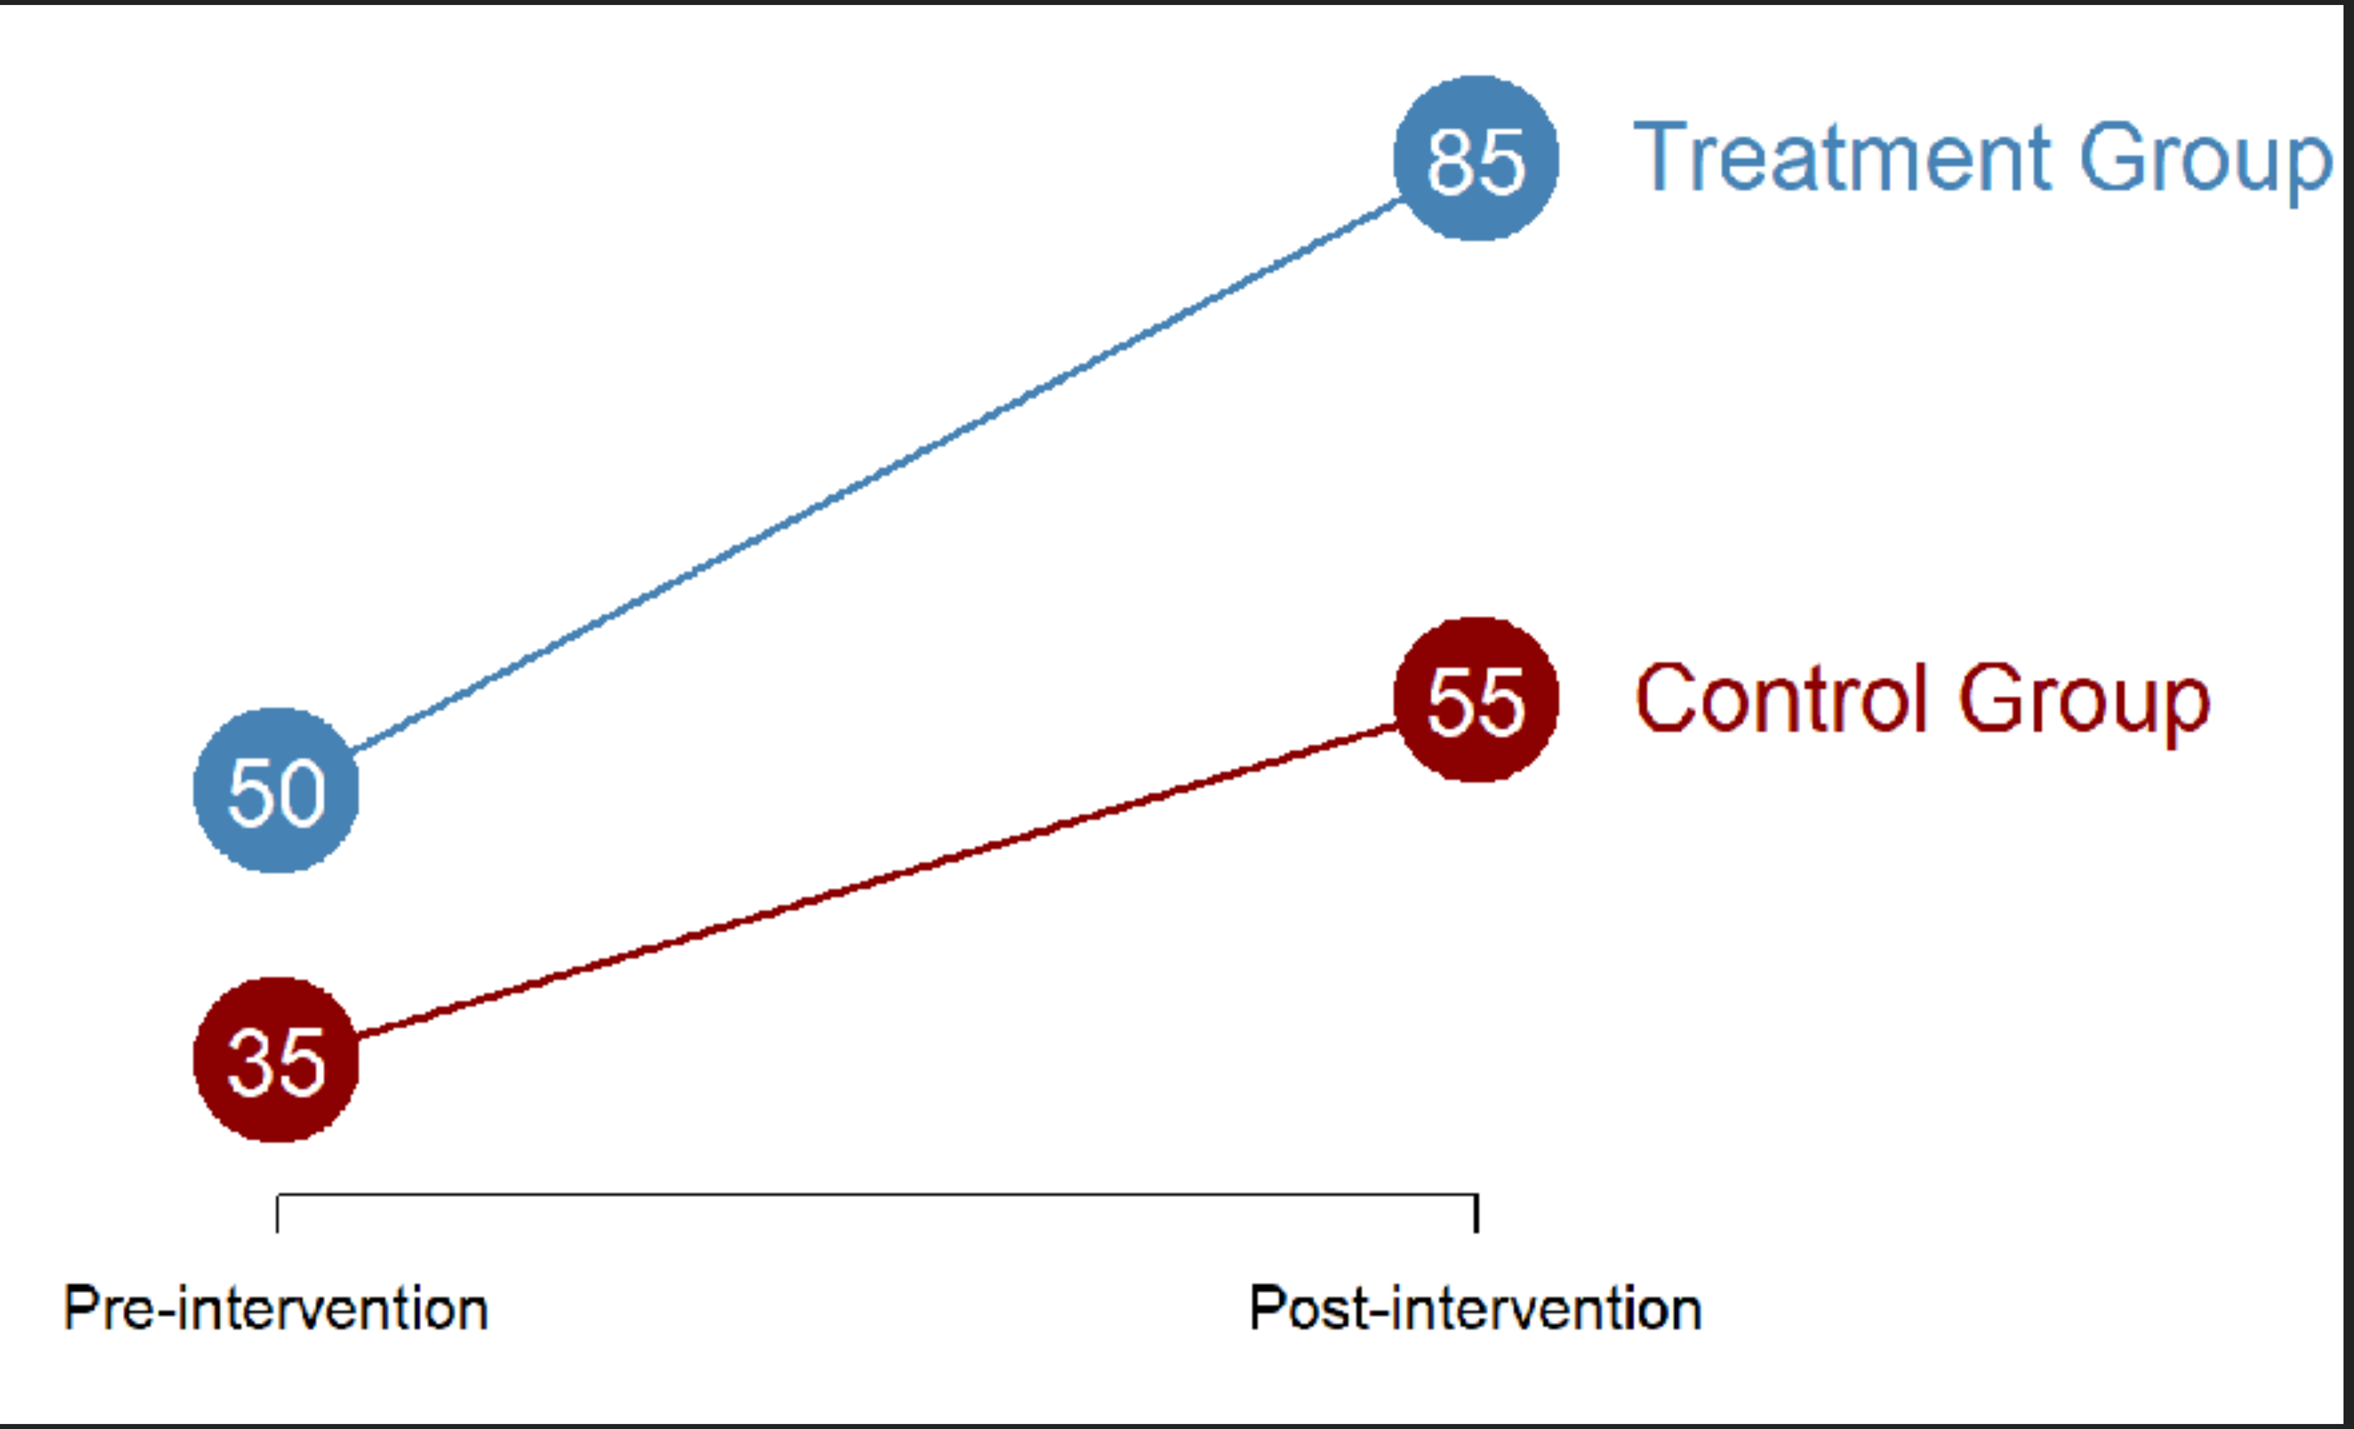

The Key Concept

- For two groups, we observe the avg outcome before & after intervention.
- The diff-in-diff estimator is the difference of their mean differences:
  - Diff-in-Diff estimate =

    (Treatment_post - Treatment_pre) -
    (Control_post - Control_pre)
- For example, the difference-in-difference based on the figure:
  - (85 - 50) - (55 - 35) = 15

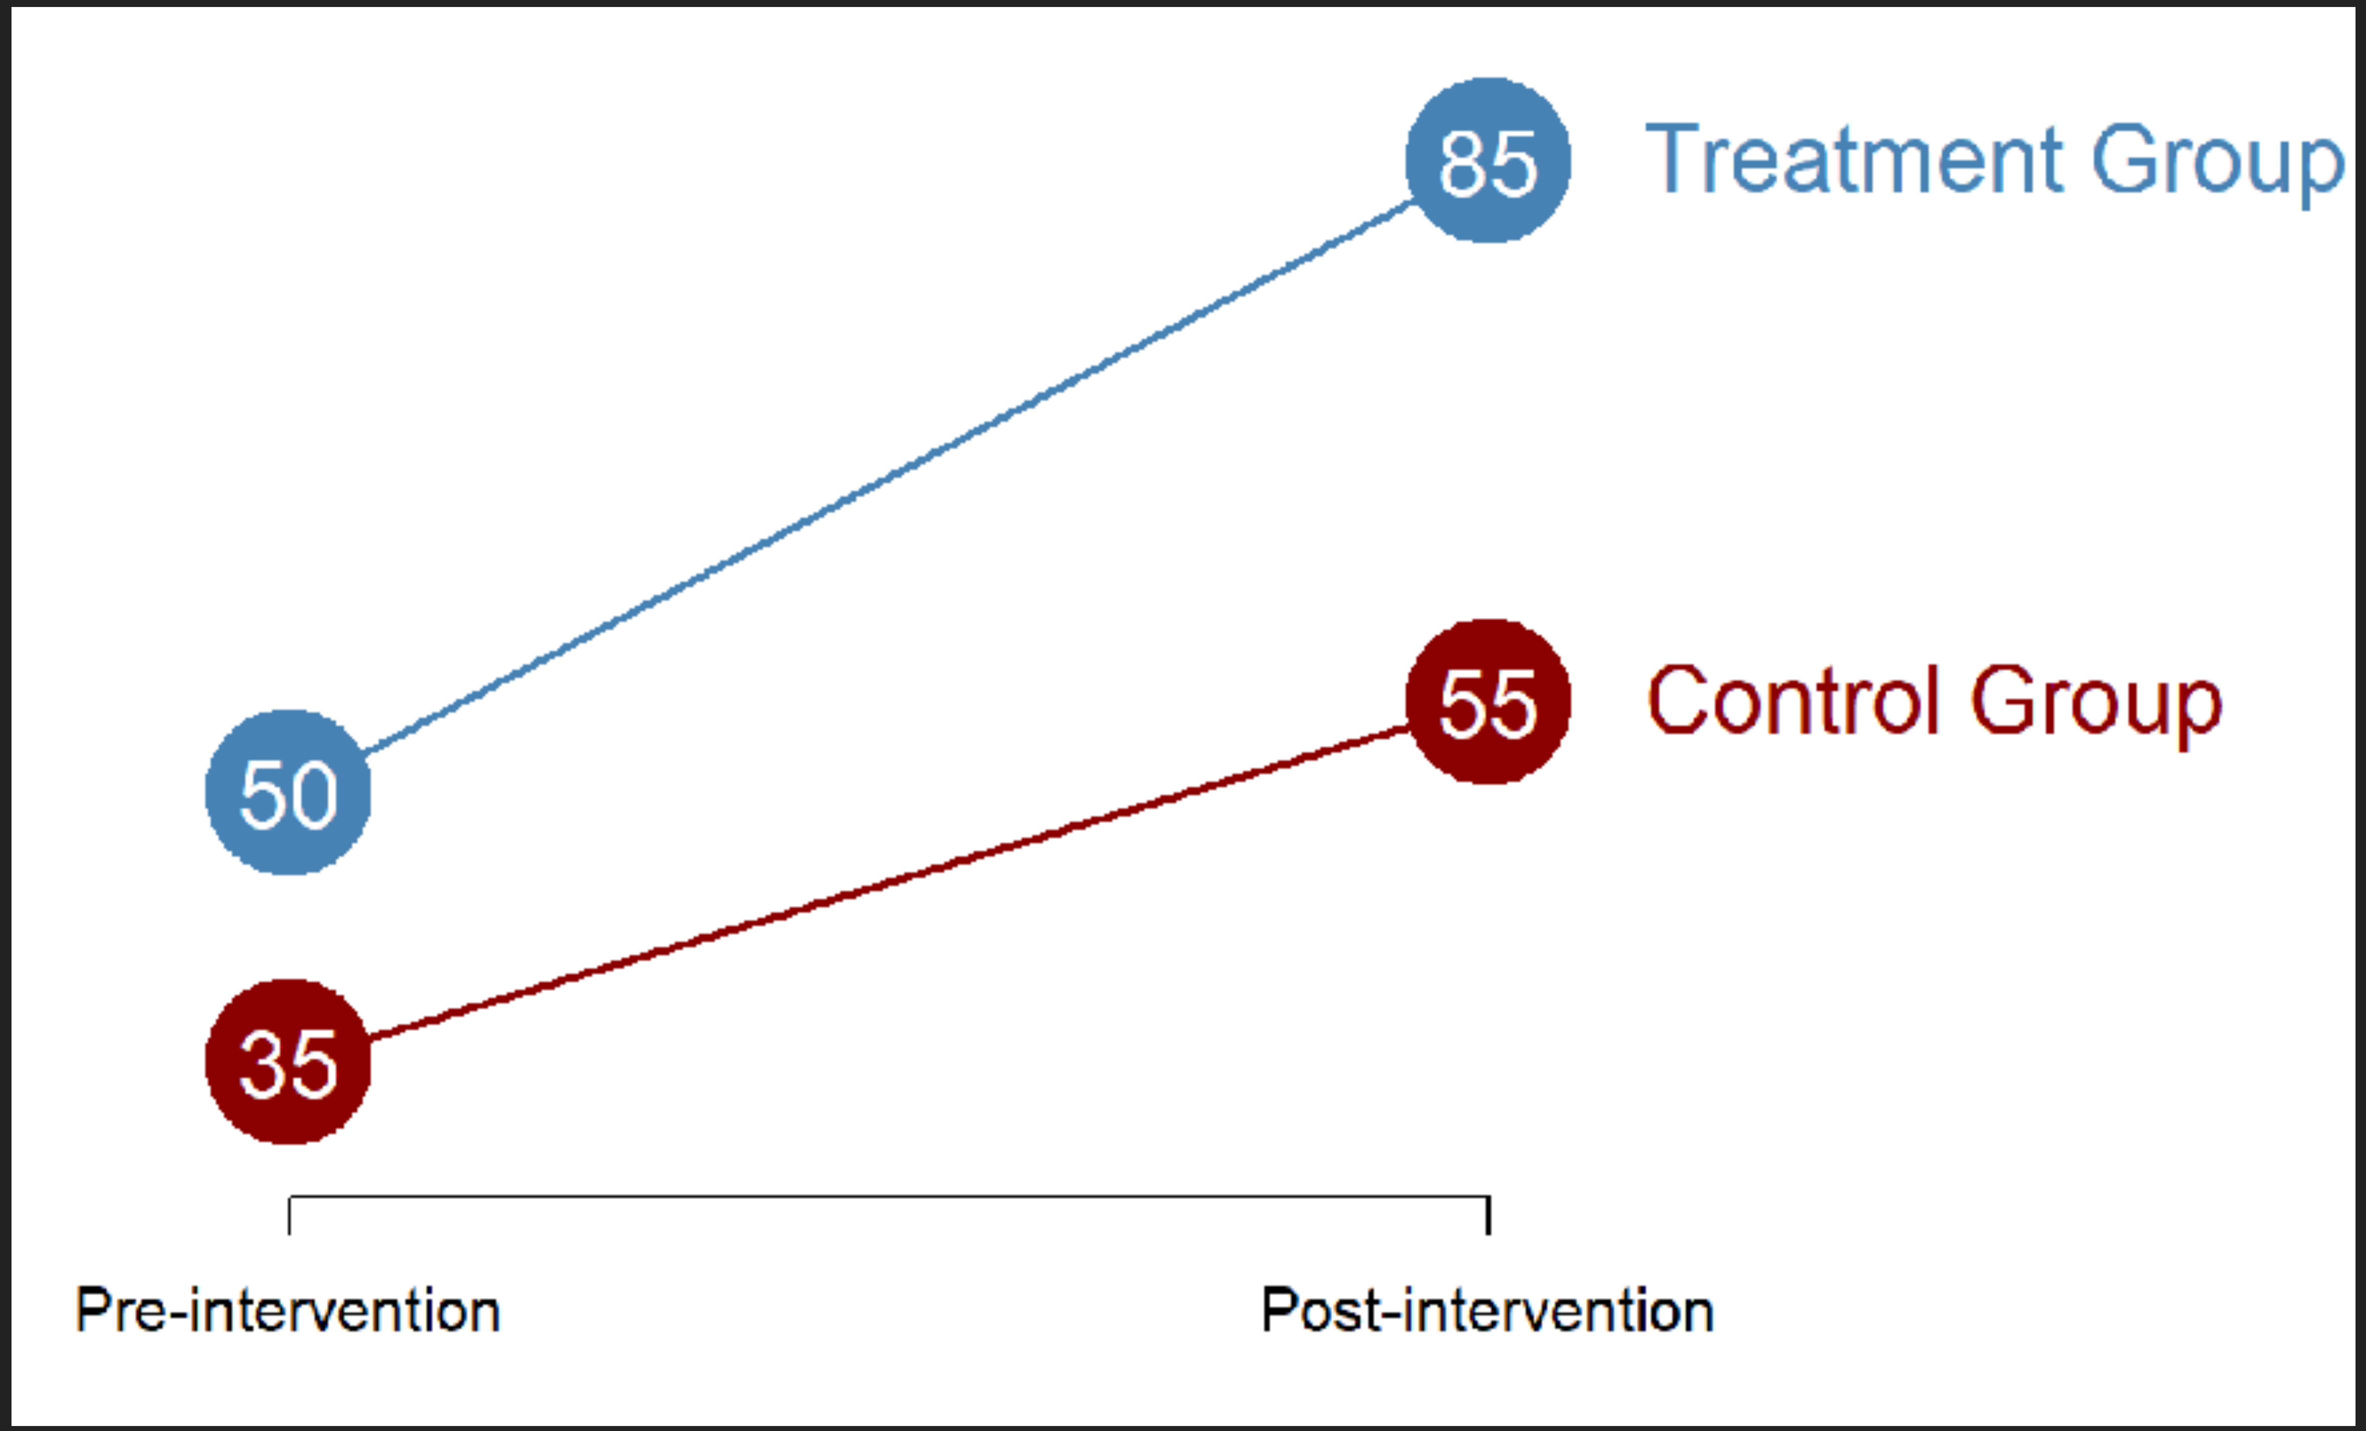

The Statistical Model

- $Y = \beta 0 + \beta 1 * Treatment + \beta 2 * Post + \beta 3 * Treatment * Post + e$

Subject	Outcome	Treatment	Post	Treatment * Post
1	74	1	1	1
1	46	1	0	0
2	96	1	1	0
2	54	1	0	1
3	50	0	1	0
3	30	0	0	0
4	60	0	1	0
4	40	0	0	0

The Coefficients

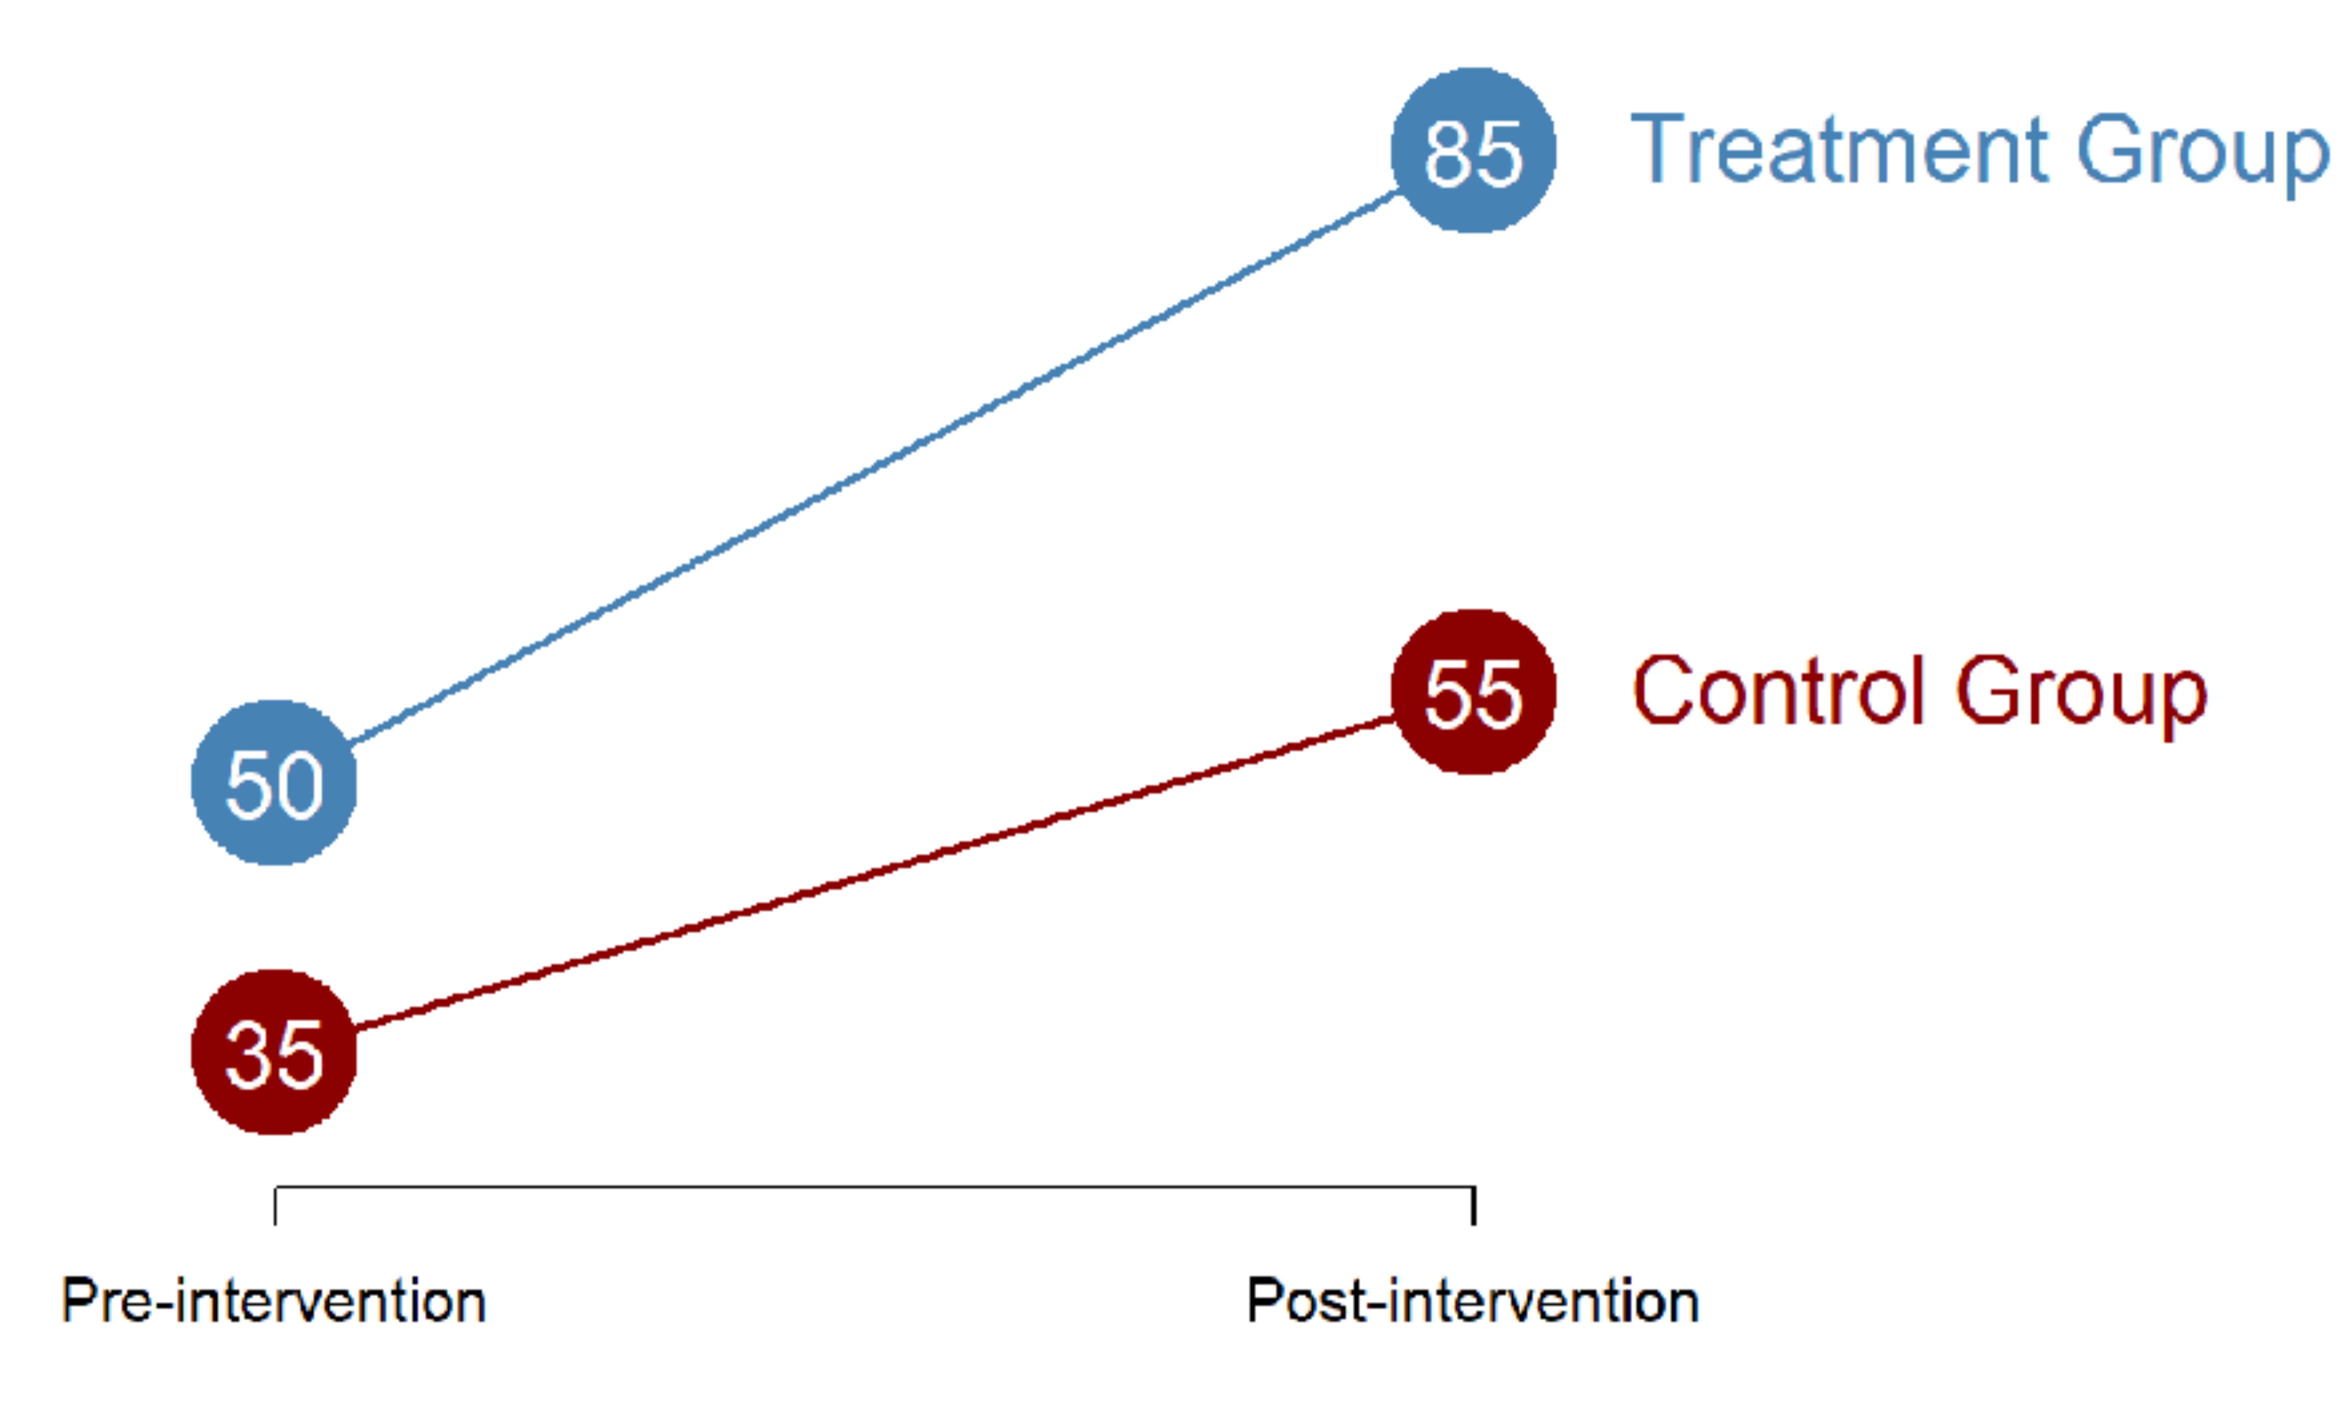


```
##
## Call:
## lm(formula = Y ~ T + P + T * P)
##
## Coefficients:
##              Estimate Std. Error   t value Pr(>|t|)
## (Intercept) 3.500e+01  1.954e-14 1.791e+15   <2e-16 ***
## T           1.500e+01  2.764e-14 5.427e+14   <2e-16 ***
## P           2.000e+01  2.764e-14 7.236e+14   <2e-16 ***
## T:P         1.500e+01  3.909e-14 3.837e+14   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 2.764e-14 on 4 degrees of freedom
## Multiple R-squared:      1,  Adjusted R-squared:      1
## F-statistic: 1.151e+30 on 3 and 4 DF,  p-value: < 2.2e-16
```

The Coefficients

```
##
## Call:
## lm(formula = Y ~ T + P + T * P)
##
## Coefficients:
##              Estimate Std. Error   t value Pr(>|t|)
## (Intercept) 3.500e+01  1.954e-14 1.791e+15   <2e-16 ***
## T           1.500e+01  2.764e-14 5.427e+14   <2e-16 ***
## P           2.000e+01  2.764e-14 7.236e+14   <2e-16 ***
## T:P         1.500e+01  3.909e-14 3.837e+14   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 2.764e-14 on 4 degrees of freedom
## Multiple R-squared:      1,  Adjusted R-squared:      1
## F-statistic: 1.151e+30 on 3 and 4 DF,  p-value: < 2.2e-16
```

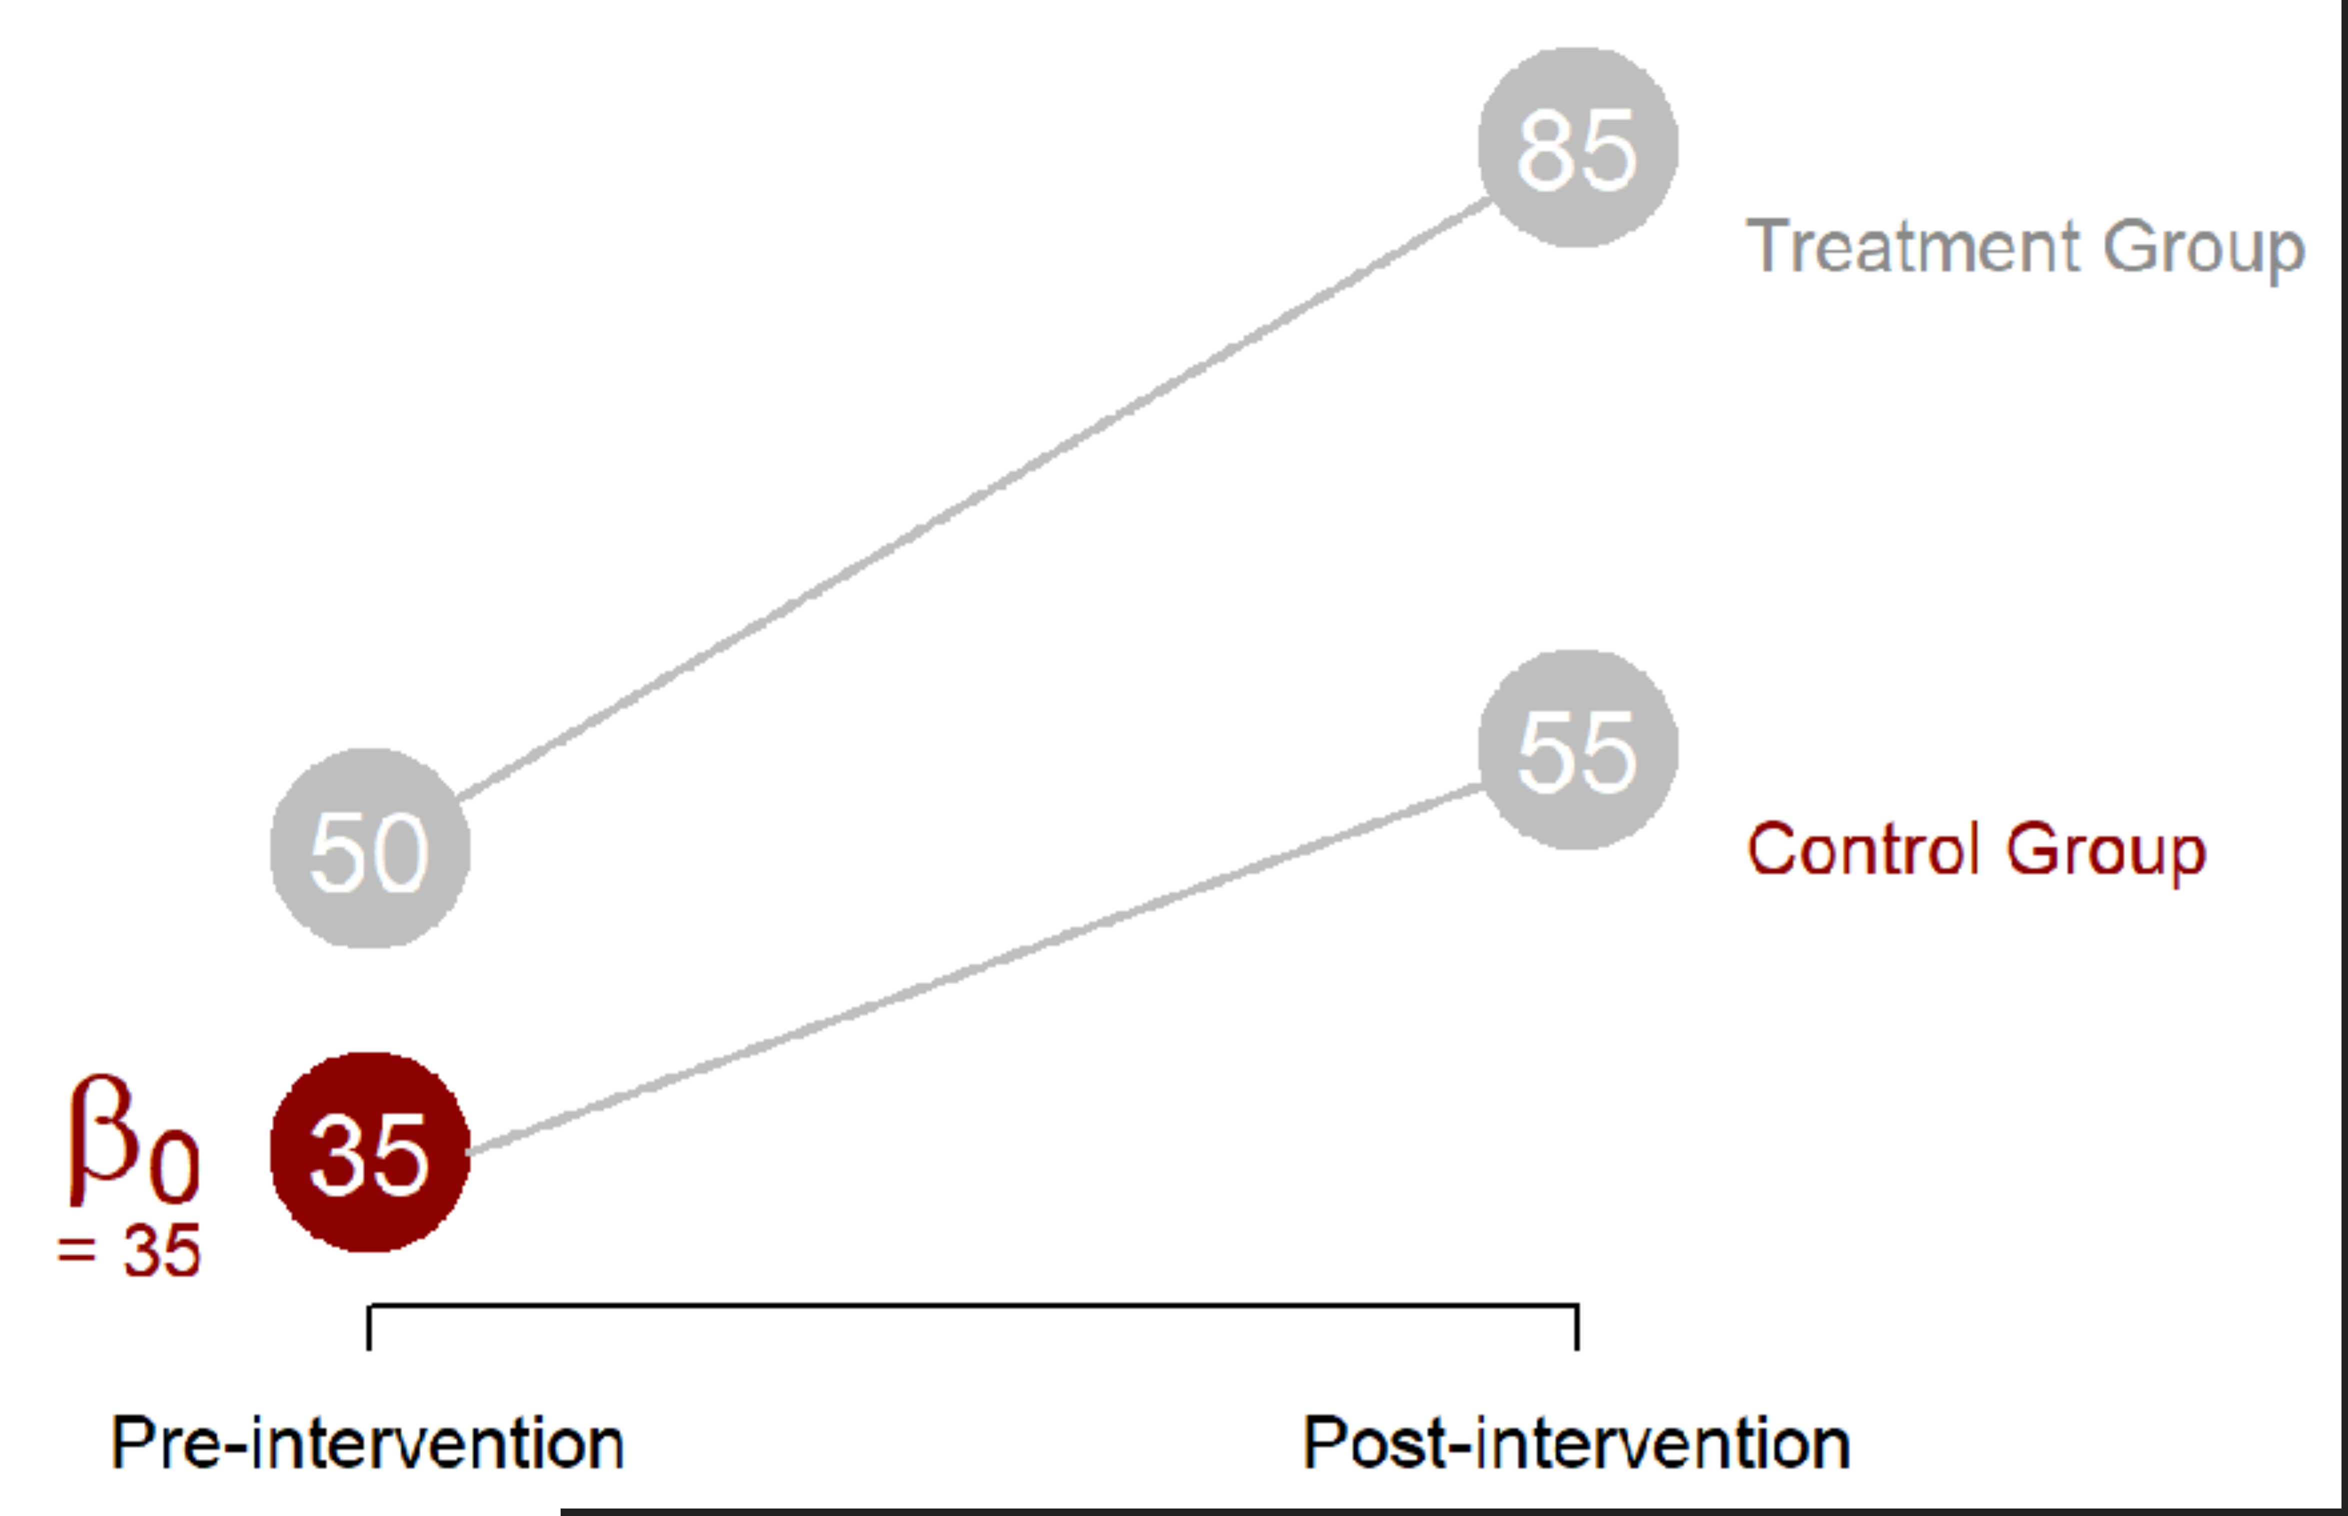

The Coefficients

```
##
## Call:
## lm(formula = Y ~ T + P + T * P)
##
## Coefficients:
##              Estimate Std. Error   t value Pr(>|t|)
## (Intercept) 3.500e+01  1.954e-14 1.791e+15   <2e-16 ***
## T           1.500e+01  2.764e-14 5.427e+14   <2e-16 ***
## P           2.000e+01  2.764e-14 7.236e+14   <2e-16 ***
## T:P         1.500e+01  3.909e-14 3.837e+14   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 2.764e-14 on 4 degrees of freedom
## Multiple R-squared:      1,  Adjusted R-squared:      1
## F-statistic: 1.151e+30 on 3 and 4 DF,  p-value: < 2.2e-16
```

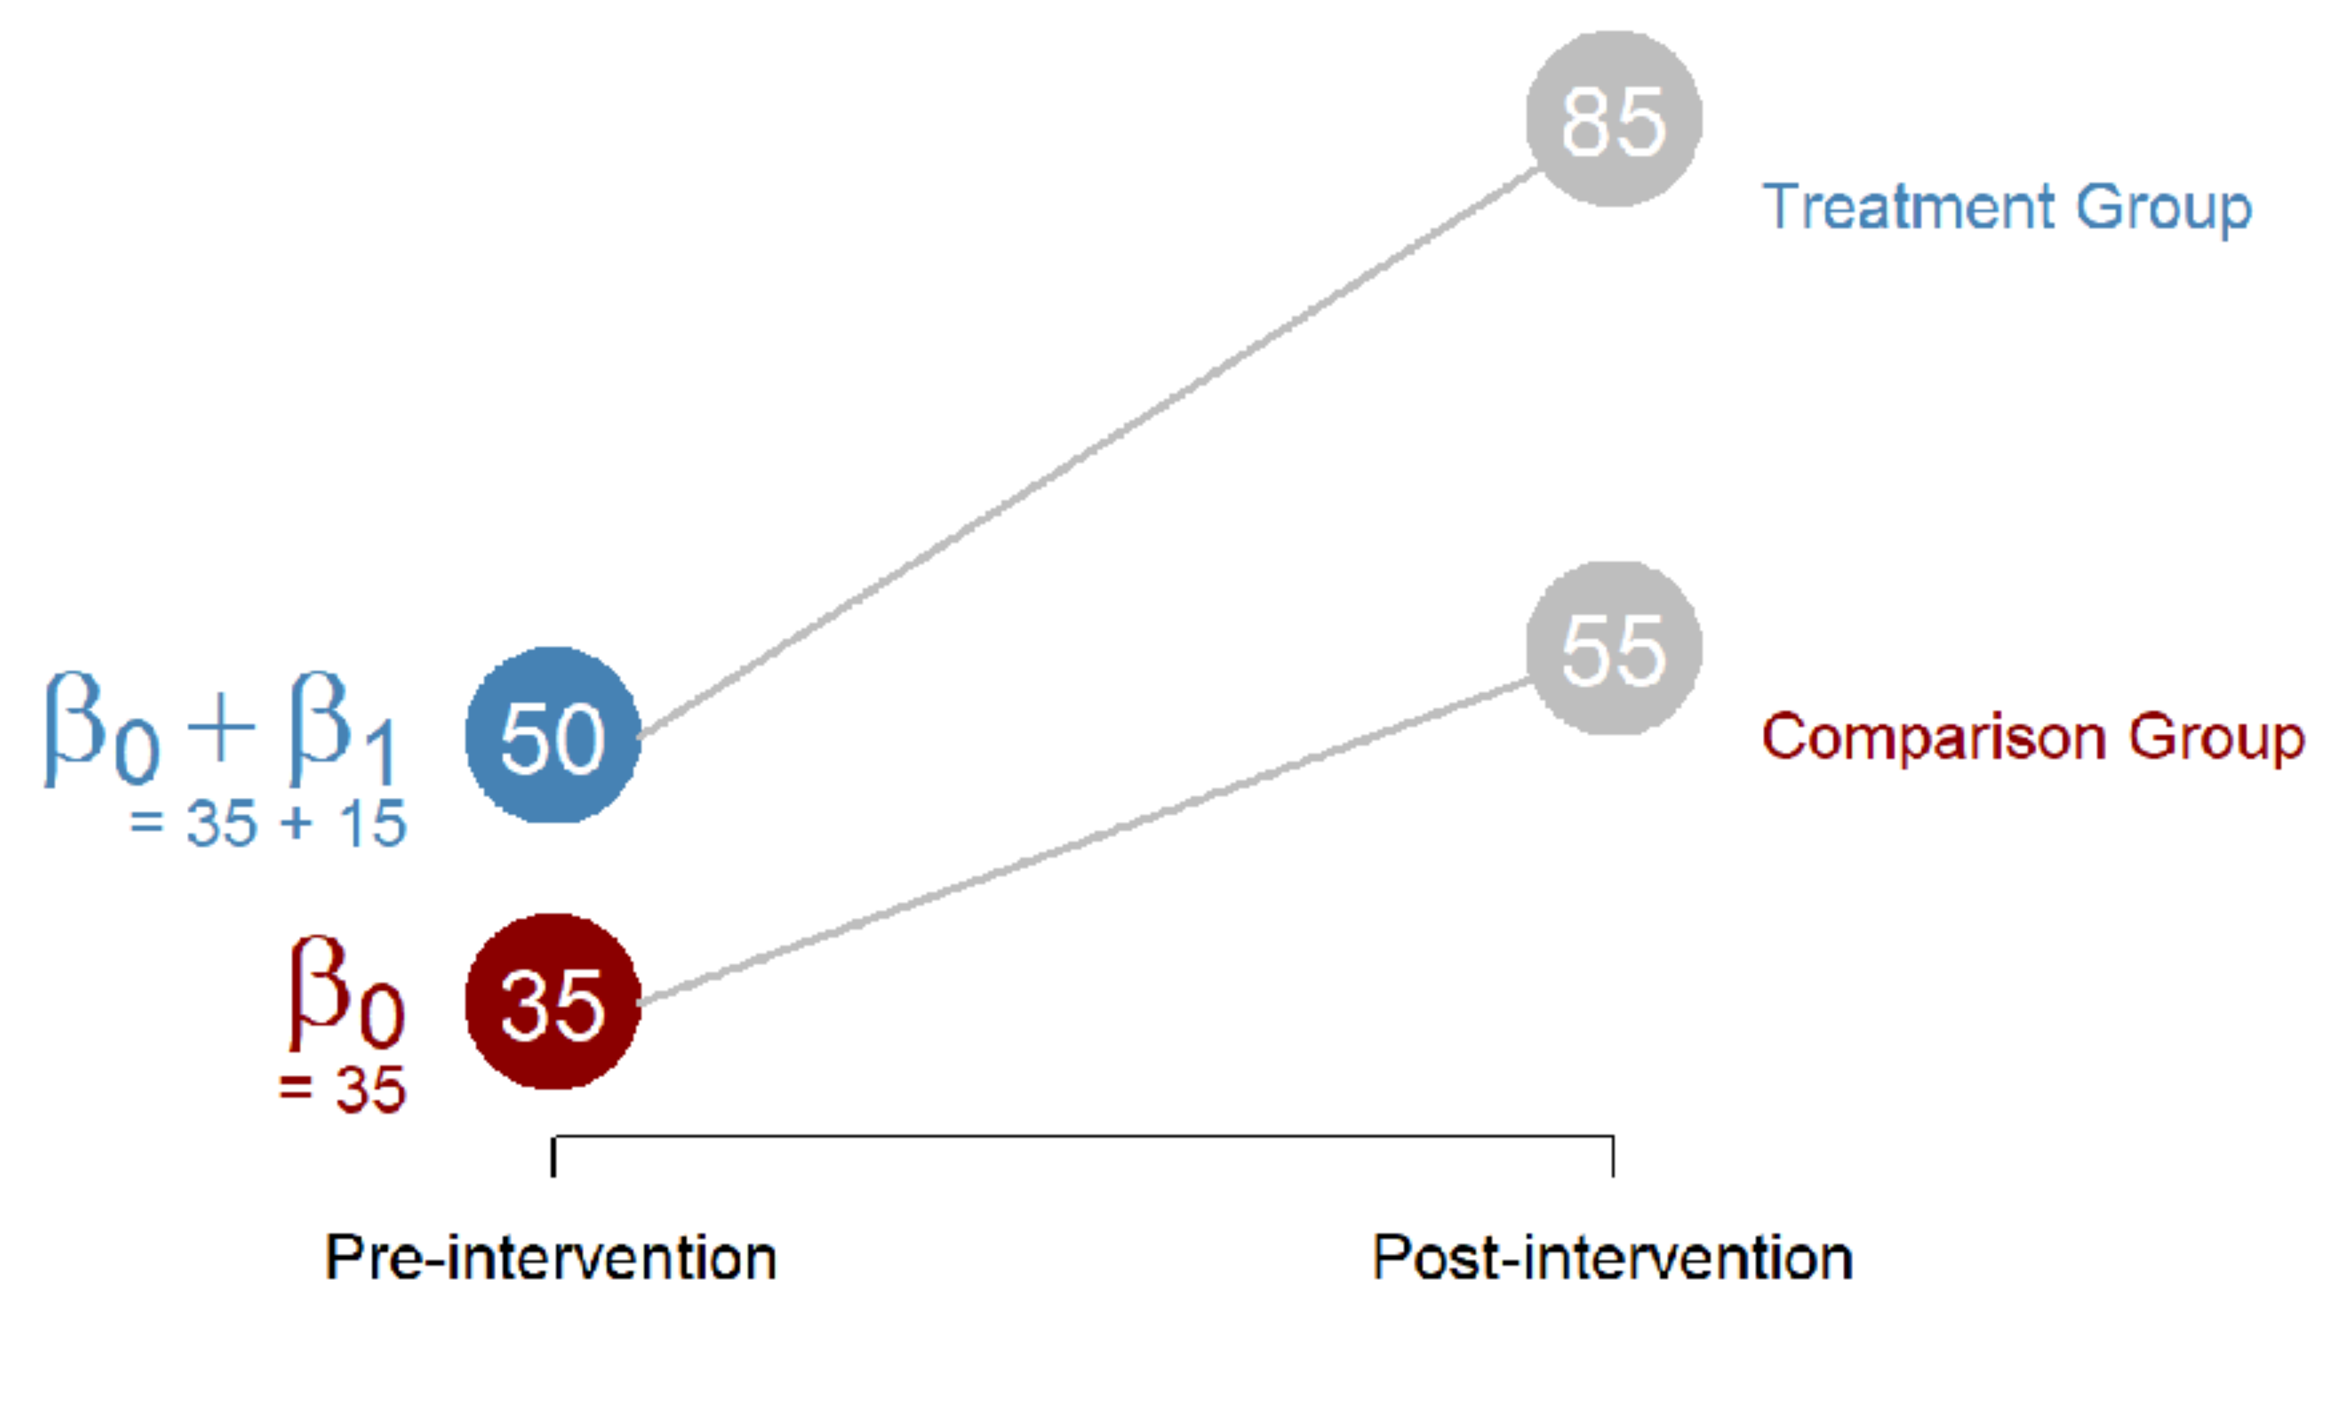

The Coefficients

```
##
## Call:
## lm(formula = Y ~ T + P + T * P)
##
## Coefficients:
##               Estimate Std. Error   t value Pr(>|t|)
## (Intercept) 3.500e+01  1.954e-14 1.791e+15   <2e-16 ***
## T           1.500e+01  2.764e-14 5.427e+14   <2e-16 ***
## P           2.000e+01  2.764e-14 7.236e+14   <2e-16 ***
## T:P         1.500e+01  3.909e-14 3.837e+14   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 2.764e-14 on 4 degrees of freedom
## Multiple R-squared:      1,  Adjusted R-squared:      1
## F-statistic: 1.151e+30 on 3 and 4 DF,  p-value: < 2.2e-16
```

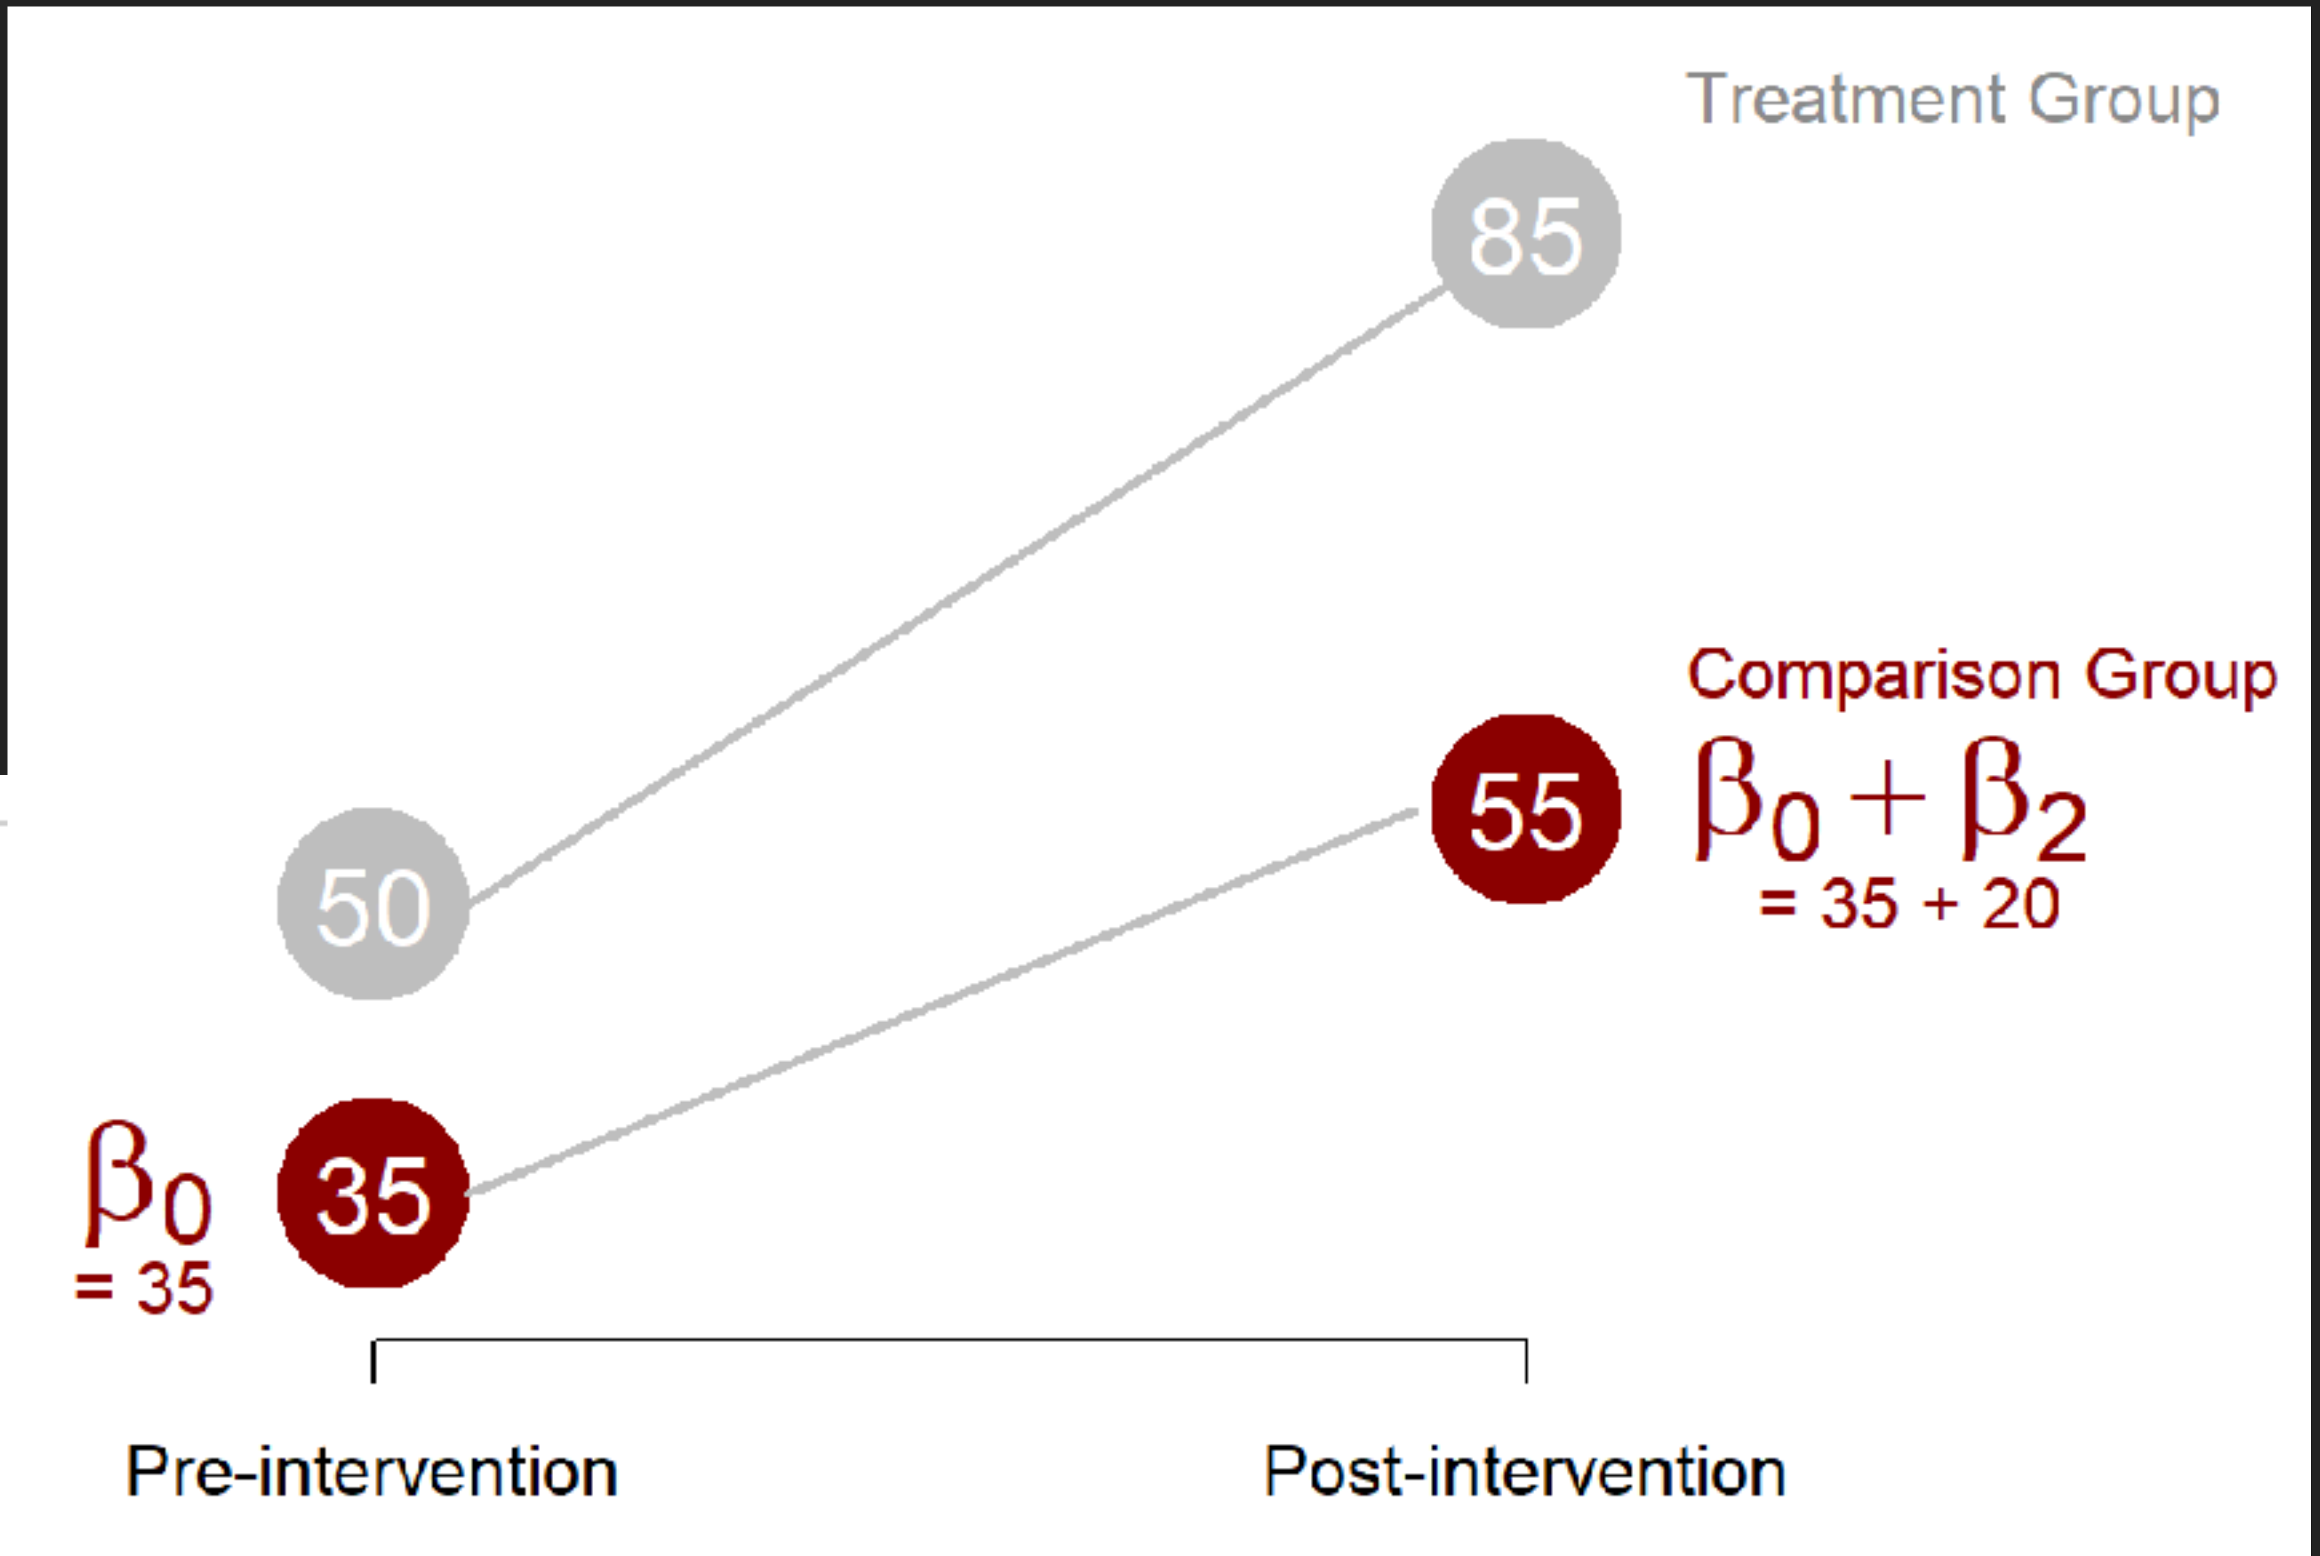

The Coefficients

```
##
## Call:
## lm(formula = Y ~ T + P + T * P)
##
## Coefficients:
##              Estimate Std. Error   t value Pr(>|t|)
## (Intercept) 3.500e+01  1.954e-14 1.791e+15   <2e-16 ***
## T           1.500e+01  2.764e-14 5.427e+14   <2e-16 ***
## P           2.000e+01  2.764e-14 7.236e+14   <2e-16 ***
## T:P         1.500e+01  3.909e-14 3.837e+14   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 2.764e-14 on 4 degrees of freedom
## Multiple R-squared:      1,  Adjusted R-squared:      1
## F-statistic: 1.151e+30 on 3 and 4 DF,  p-value: < 2.2e-16
```

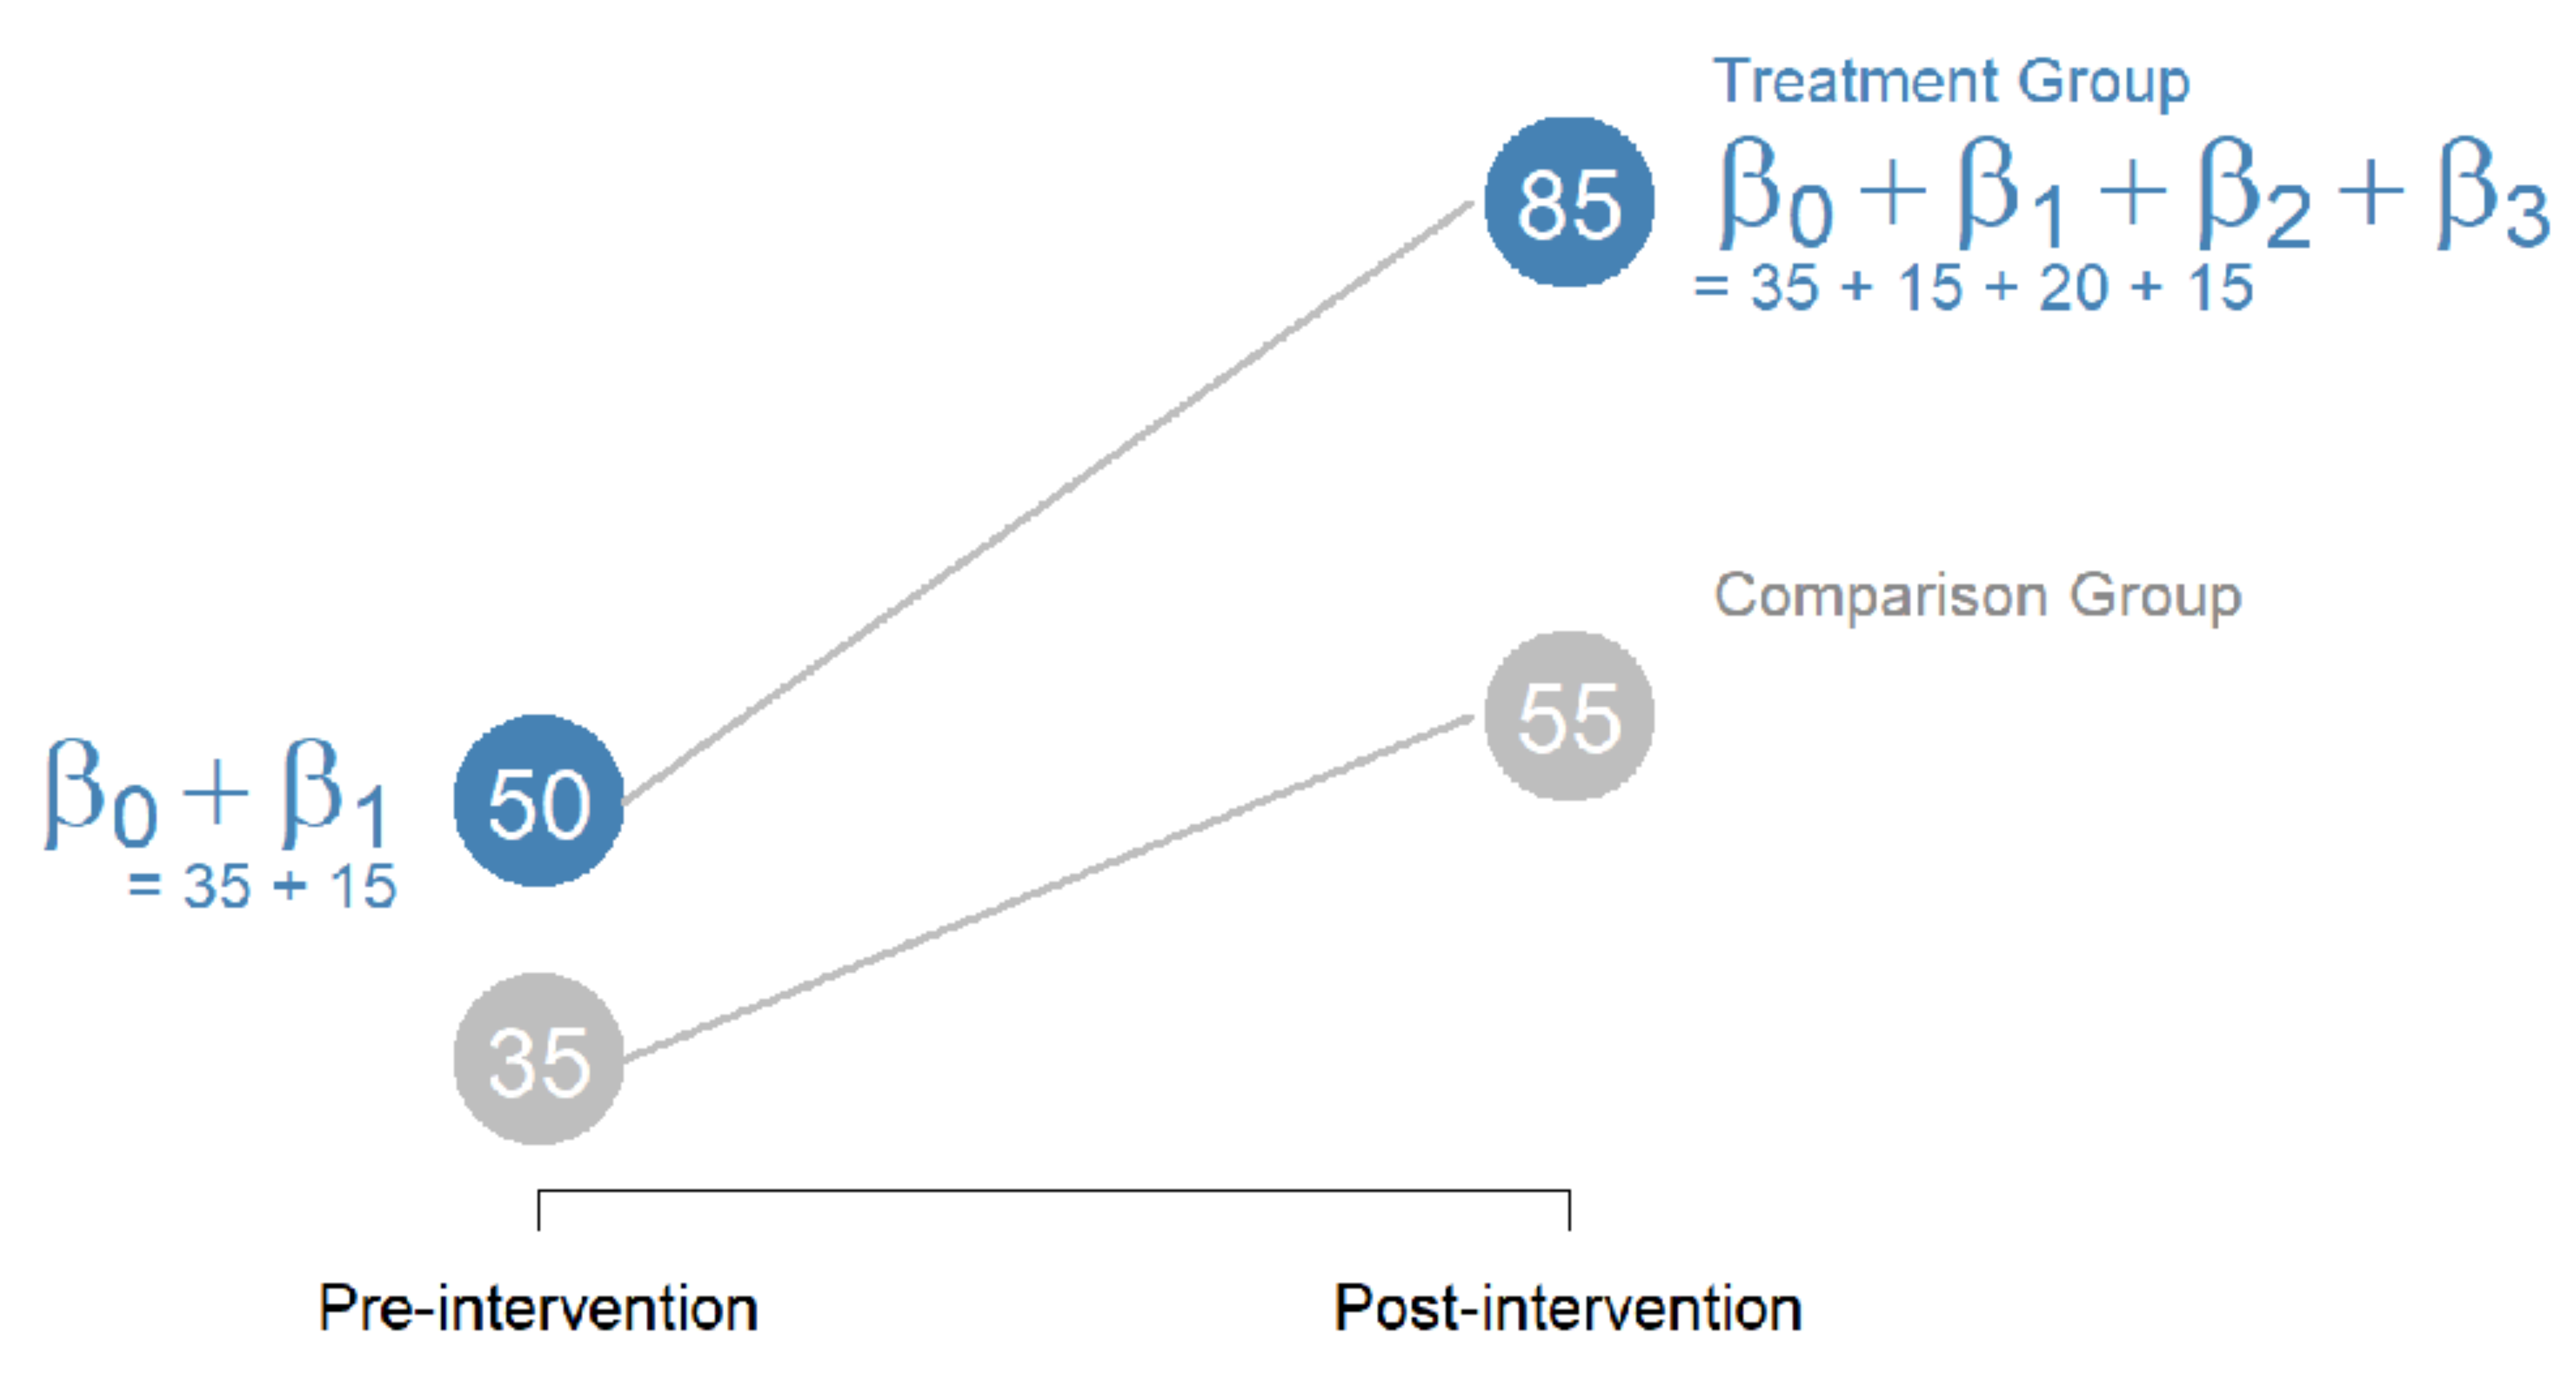

The Coefficients

```
##
## Call:
## lm(formula = Y ~ T + P + T * P)
##
## Coefficients:
##              Estimate Std. Error   t value Pr(>|t|)
## (Intercept) 3.500e+01  1.954e-14 1.791e+15   <2e-16 ***
## T           1.500e+01  2.764e-14 5.427e+14   <2e-16 ***
## P           2.000e+01  2.764e-14 7.236e+14   <2e-16 ***
## T:P         1.500e+01  3.909e-14 3.837e+14   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 2.764e-14 on 4 degrees of freedom
## Multiple R-squared:      1,  Adjusted R-squared:      1
## F-statistic: 1.151e+30 on 3 and 4 DF,  p-value: < 2.2e-16
```

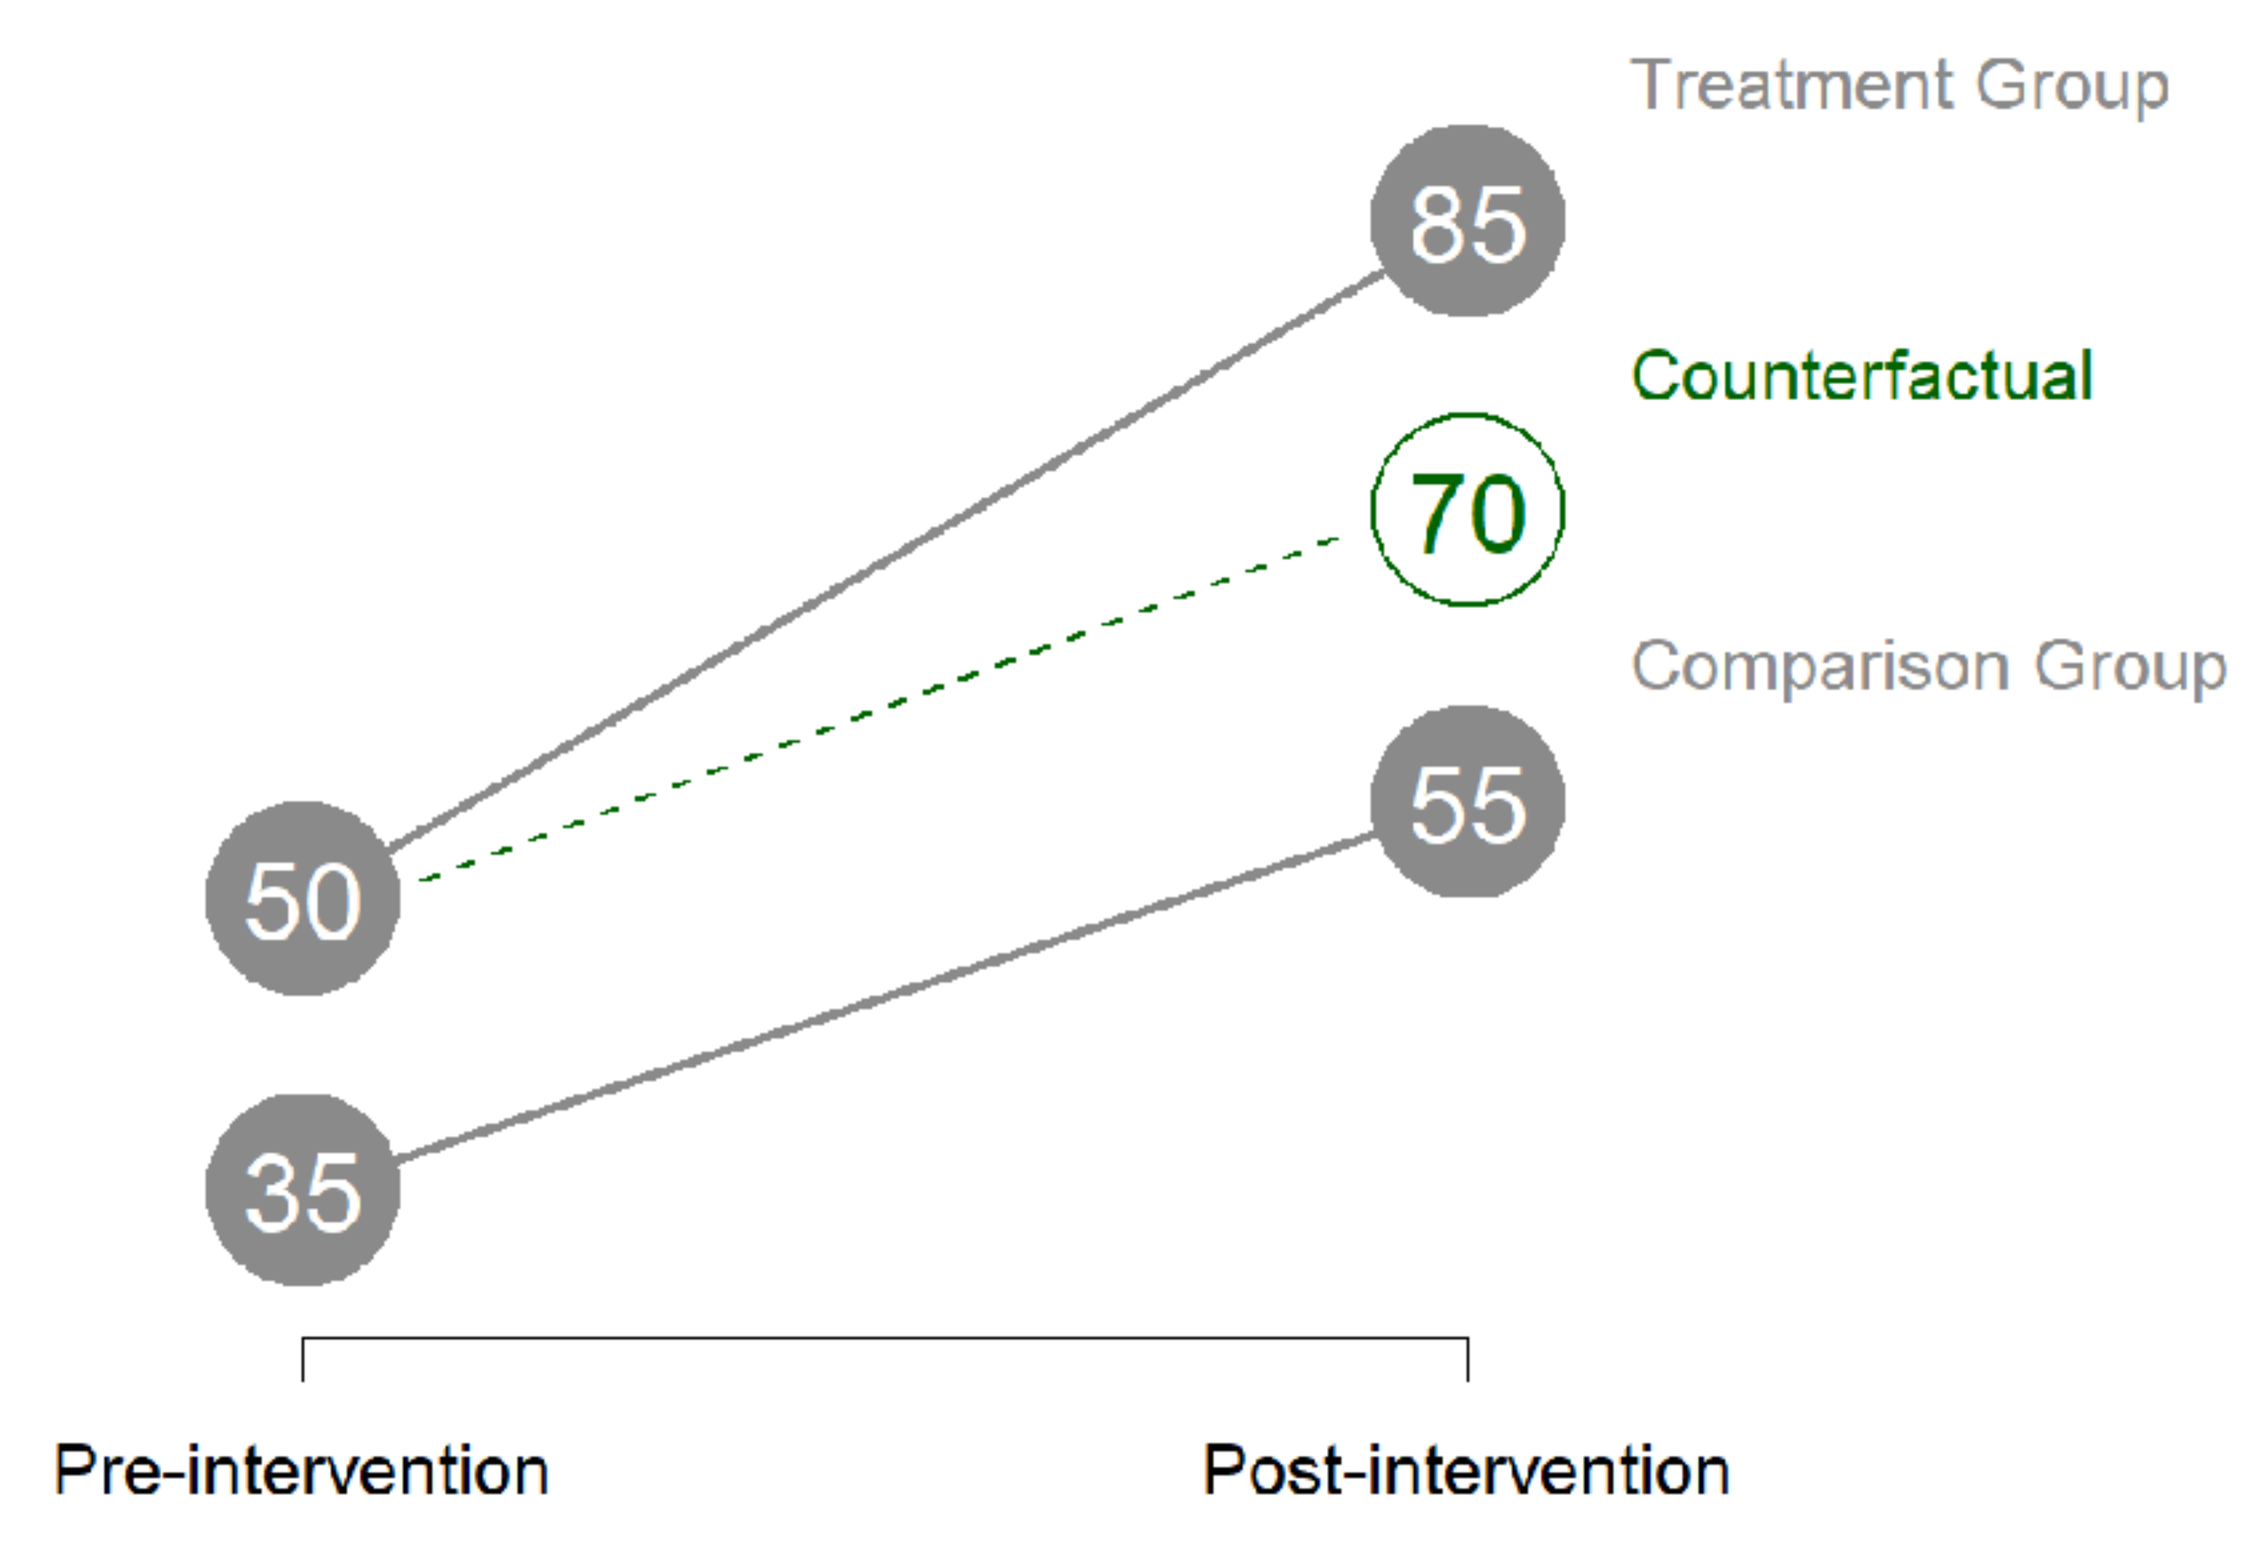

The Coefficients

```
##
## Call:
## lm(formula = Y ~ T + P + T * P)
##
## Coefficients:
##              Estimate Std. Error   t value Pr(>|t|)
## (Intercept) 3.500e+01  1.954e-14 1.791e+15   <2e-16 ***
## T           1.500e+01  2.764e-14 5.427e+14   <2e-16 ***
## P           2.000e+01  2.764e-14 7.236e+14   <2e-16 ***
## T:P         1.500e+01  3.909e-14 3.837e+14   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 2.764e-14 on 4 degrees of freedom
## Multiple R-squared:      1,  Adjusted R-squared:      1
## F-statistic: 1.151e+30 on 3 and 4 DF,  p-value: < 2.2e-16
```

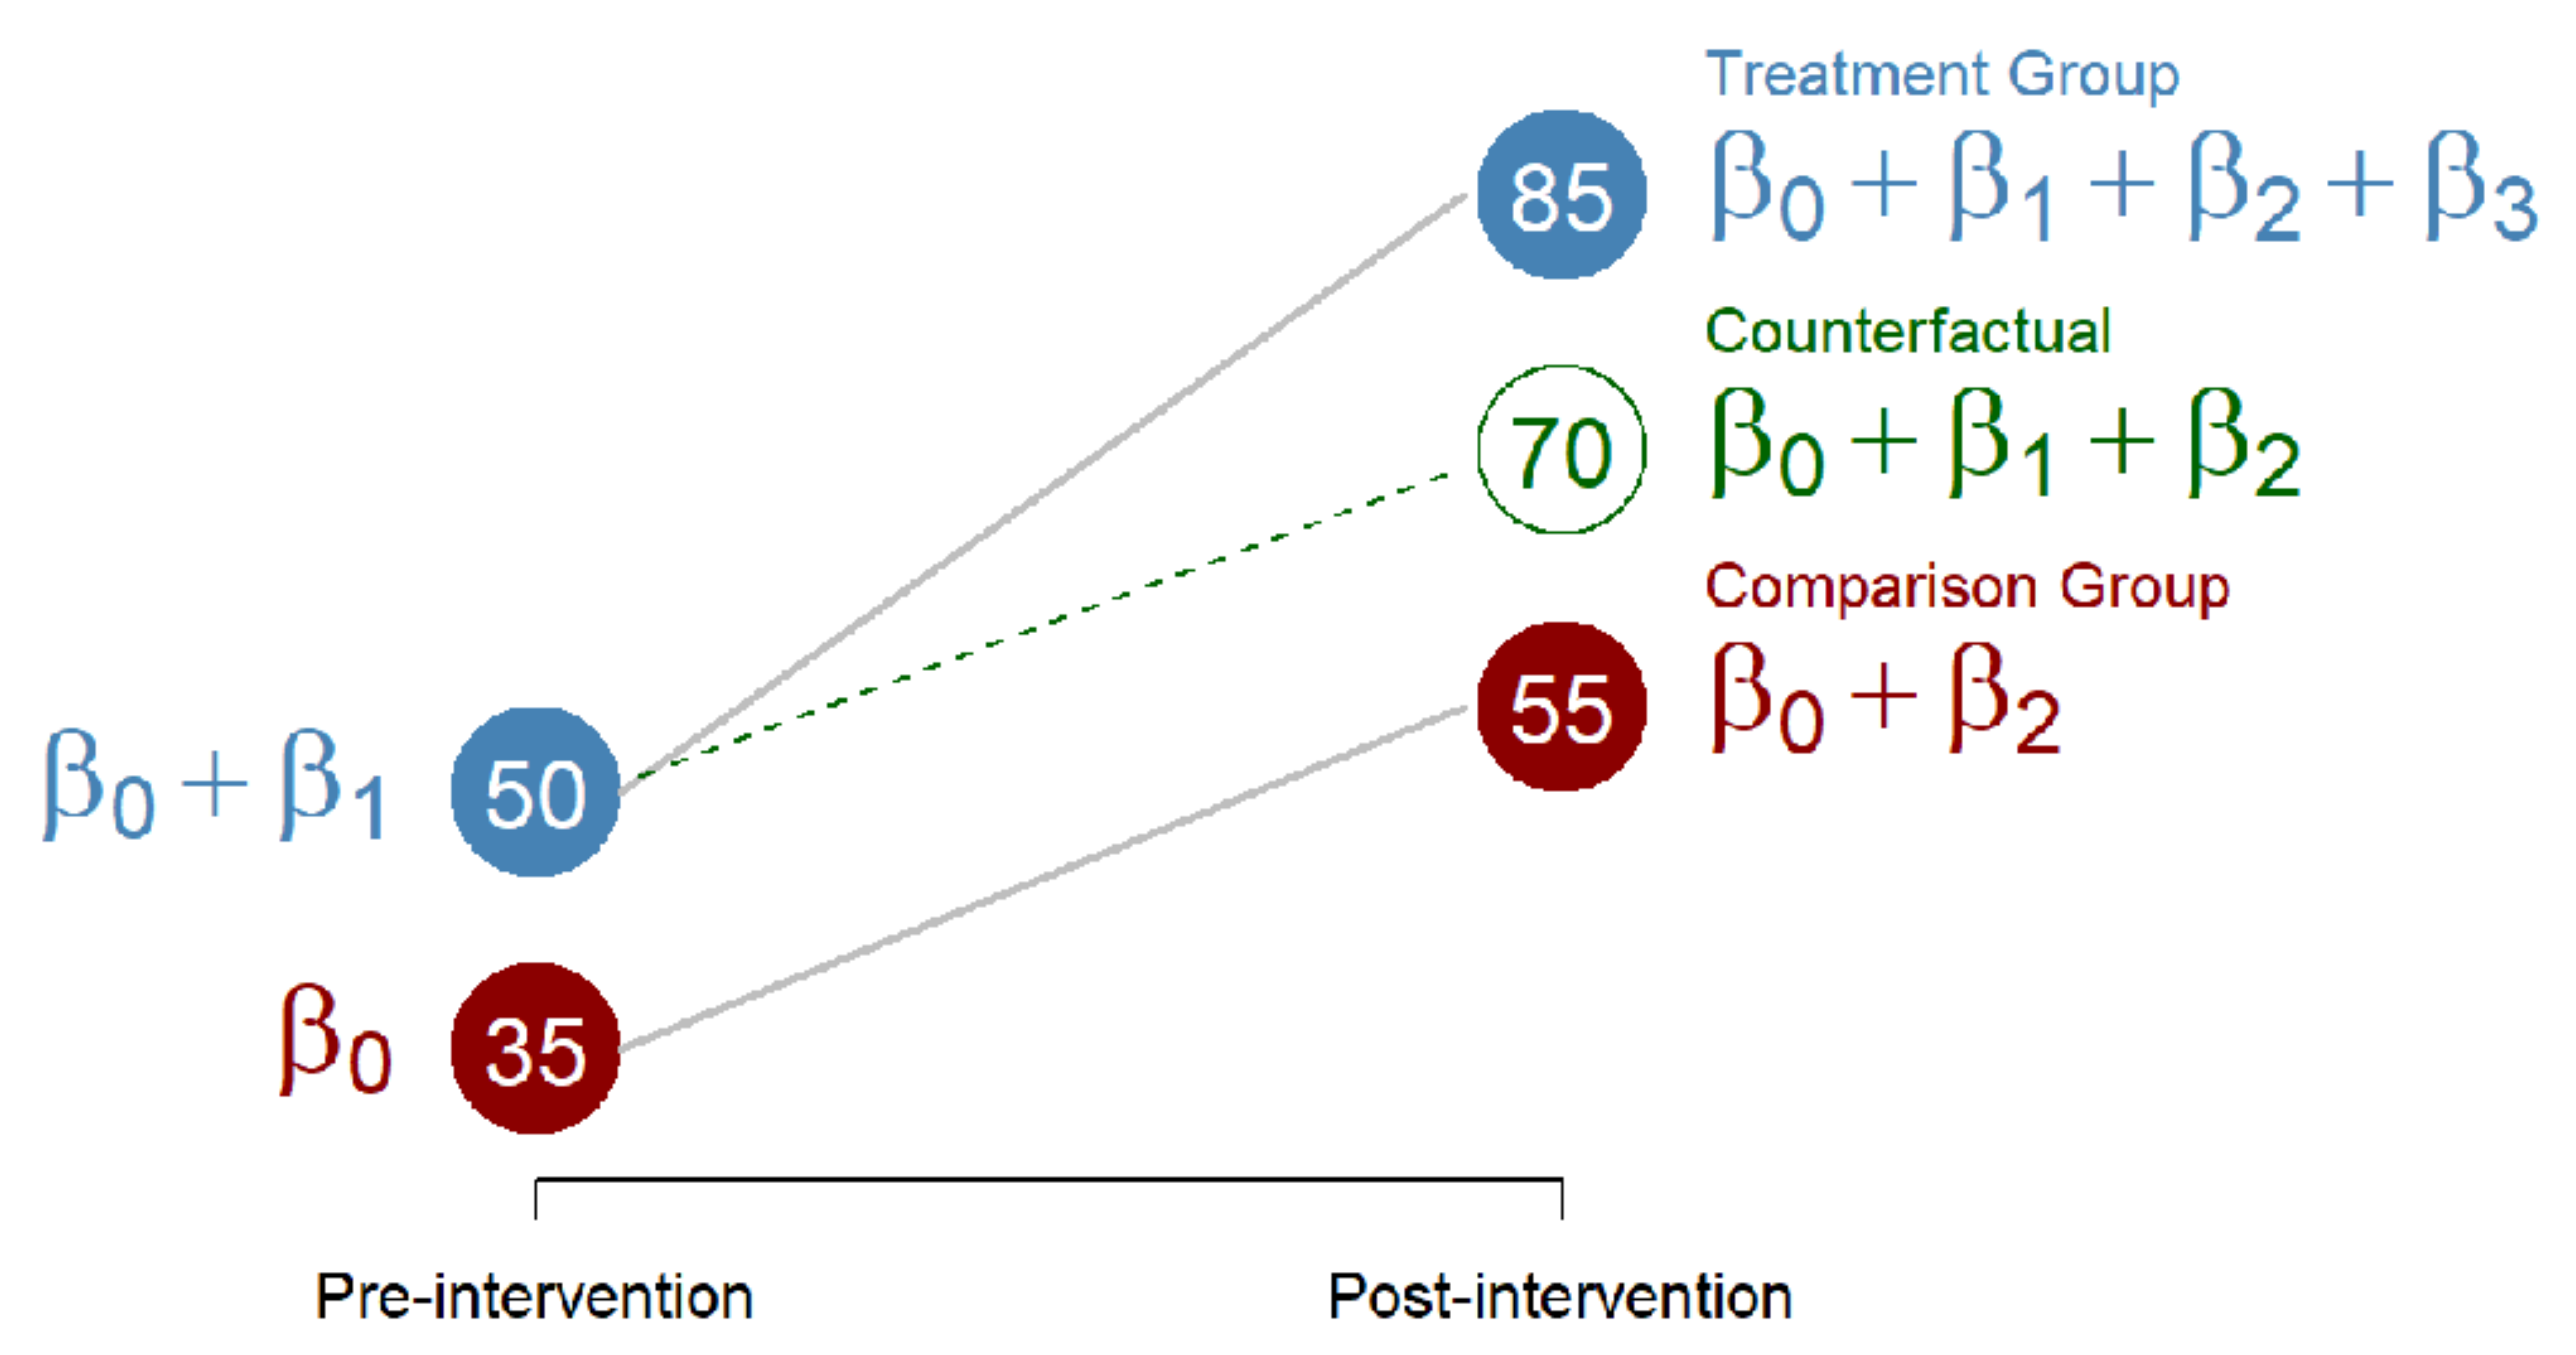

Causal Impact

https://google.github.io/CausalImpact/CausalImpact.html

Slides from: Inferring the effect of an event using CausalImpact by Kay Brodersen, Big Data Spain 2016 Conference https://youtu.be/GTgZfCltMm8

Key Idea

- The motivation to use Causal Impact methodology is that the Difference in differences in limited in the following ways:
  - DD is traditionally based on a static regression model that assumes independent and identically distributed data despite the fact that the design has a temporal component.
  - Most DD analyses only consider two time points: before and after the intervention. In practice, we also have to consider the manner in which an effect evolves over time, especially its onset and decay structure.
- The idea is to use the trend in the control group to forecast the trend in the treated group which would be the trend if the treatment had not happened.
- Then the actual causal estimate would be the difference in the actual trend vs. the counter-factual trend of the treated group that we predicted.
- Causal Impact uses Bayesian structural time-series models to explain the temporal evolution of an observed outcome.

A Simple Example

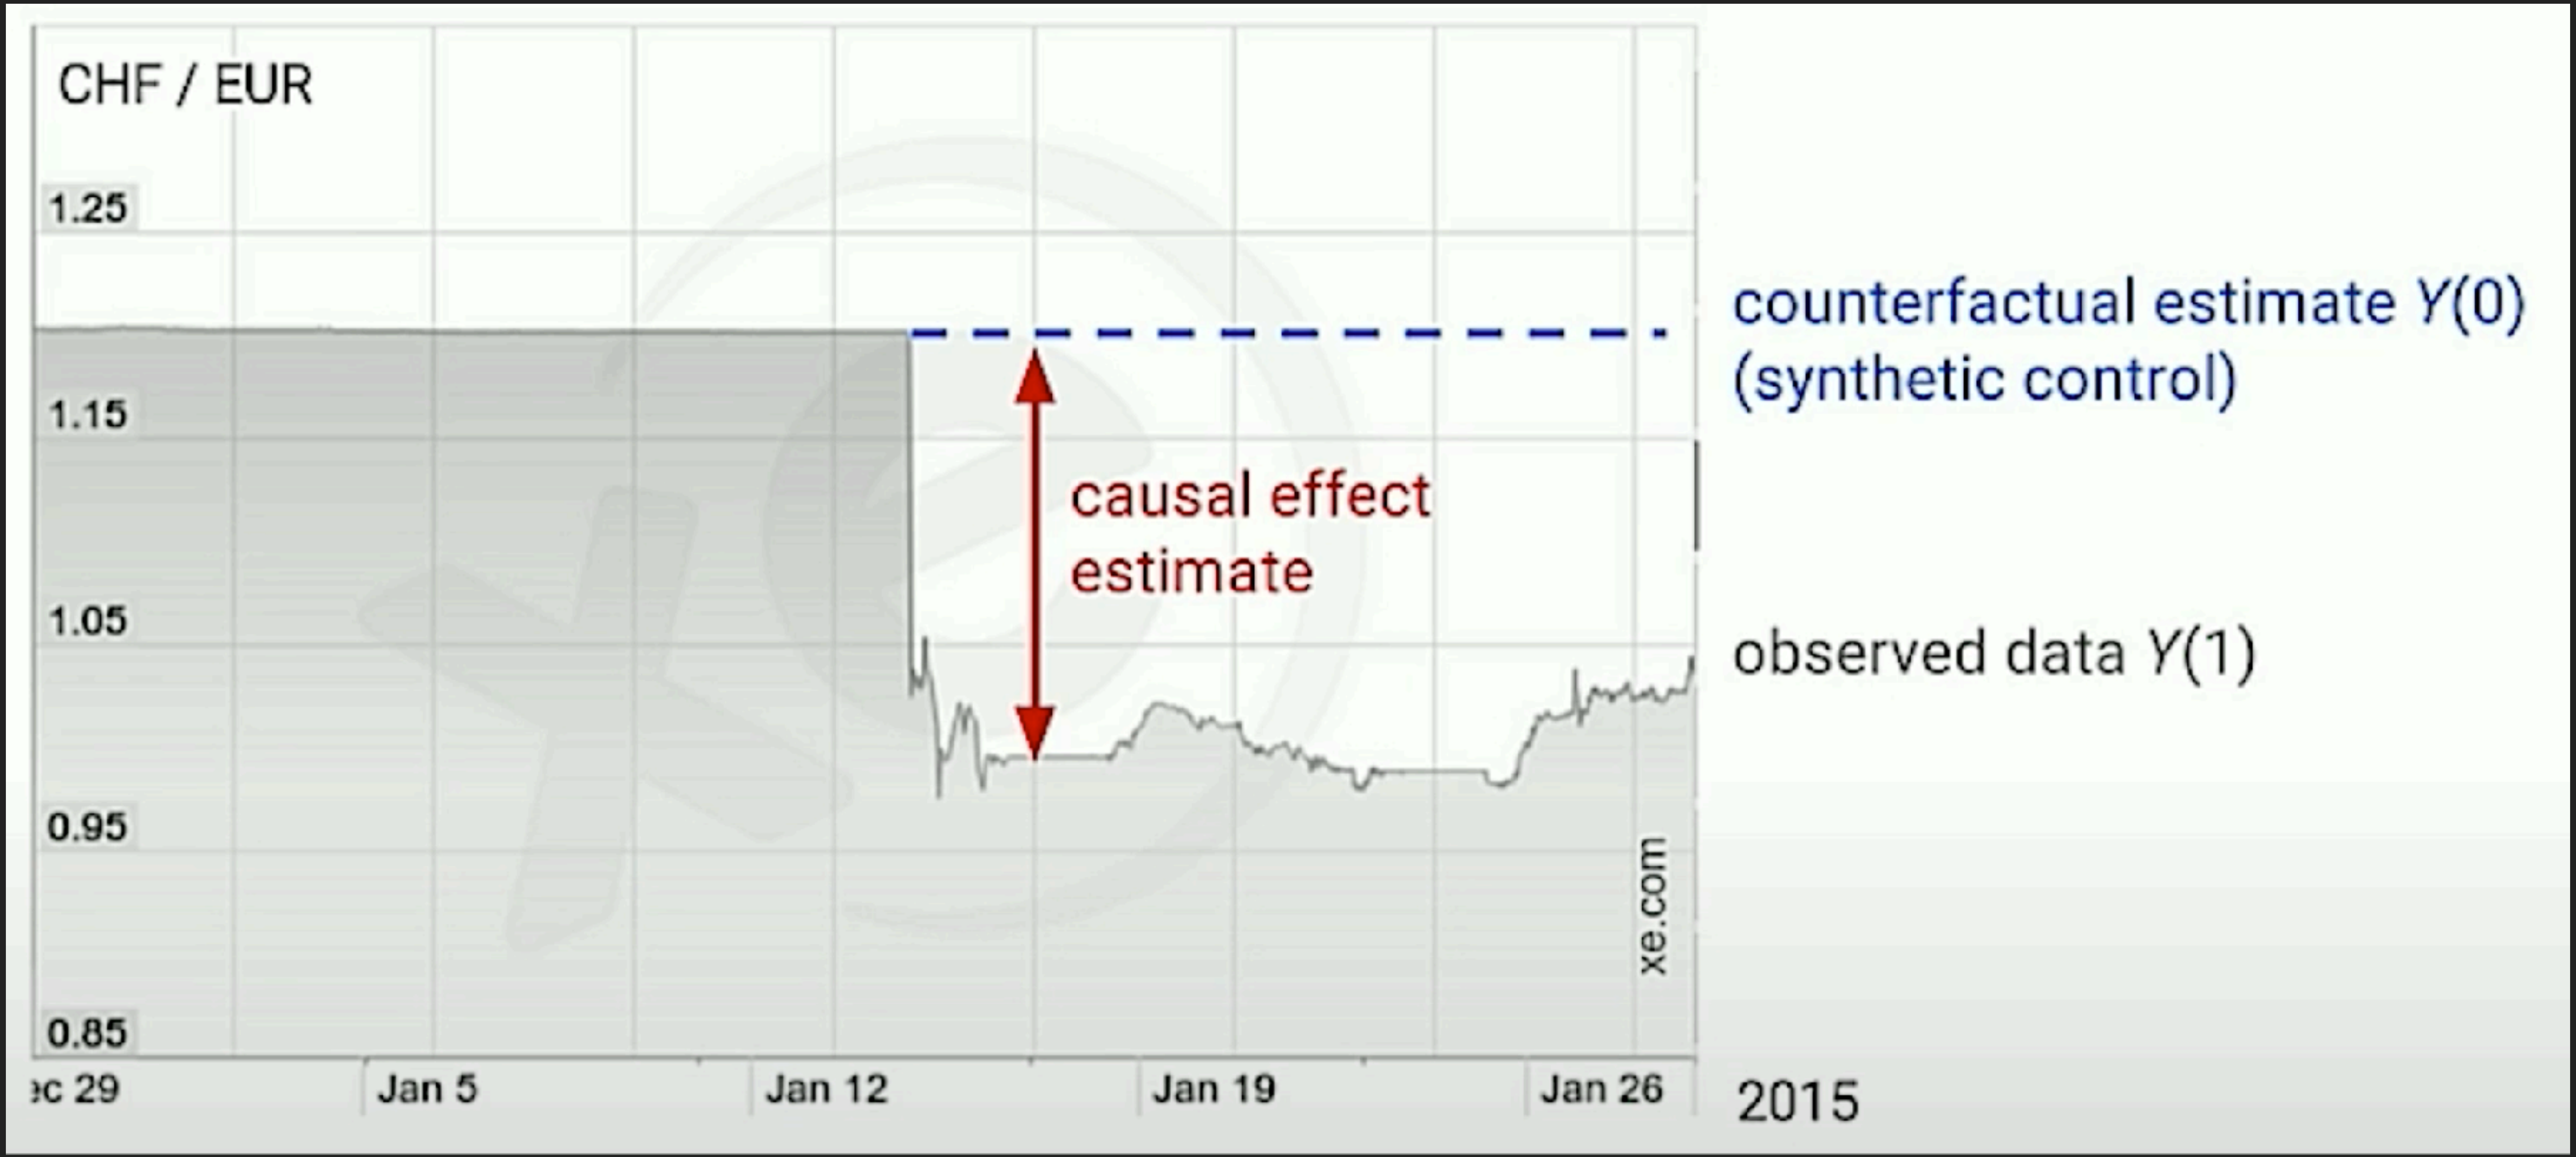

A Harder Example

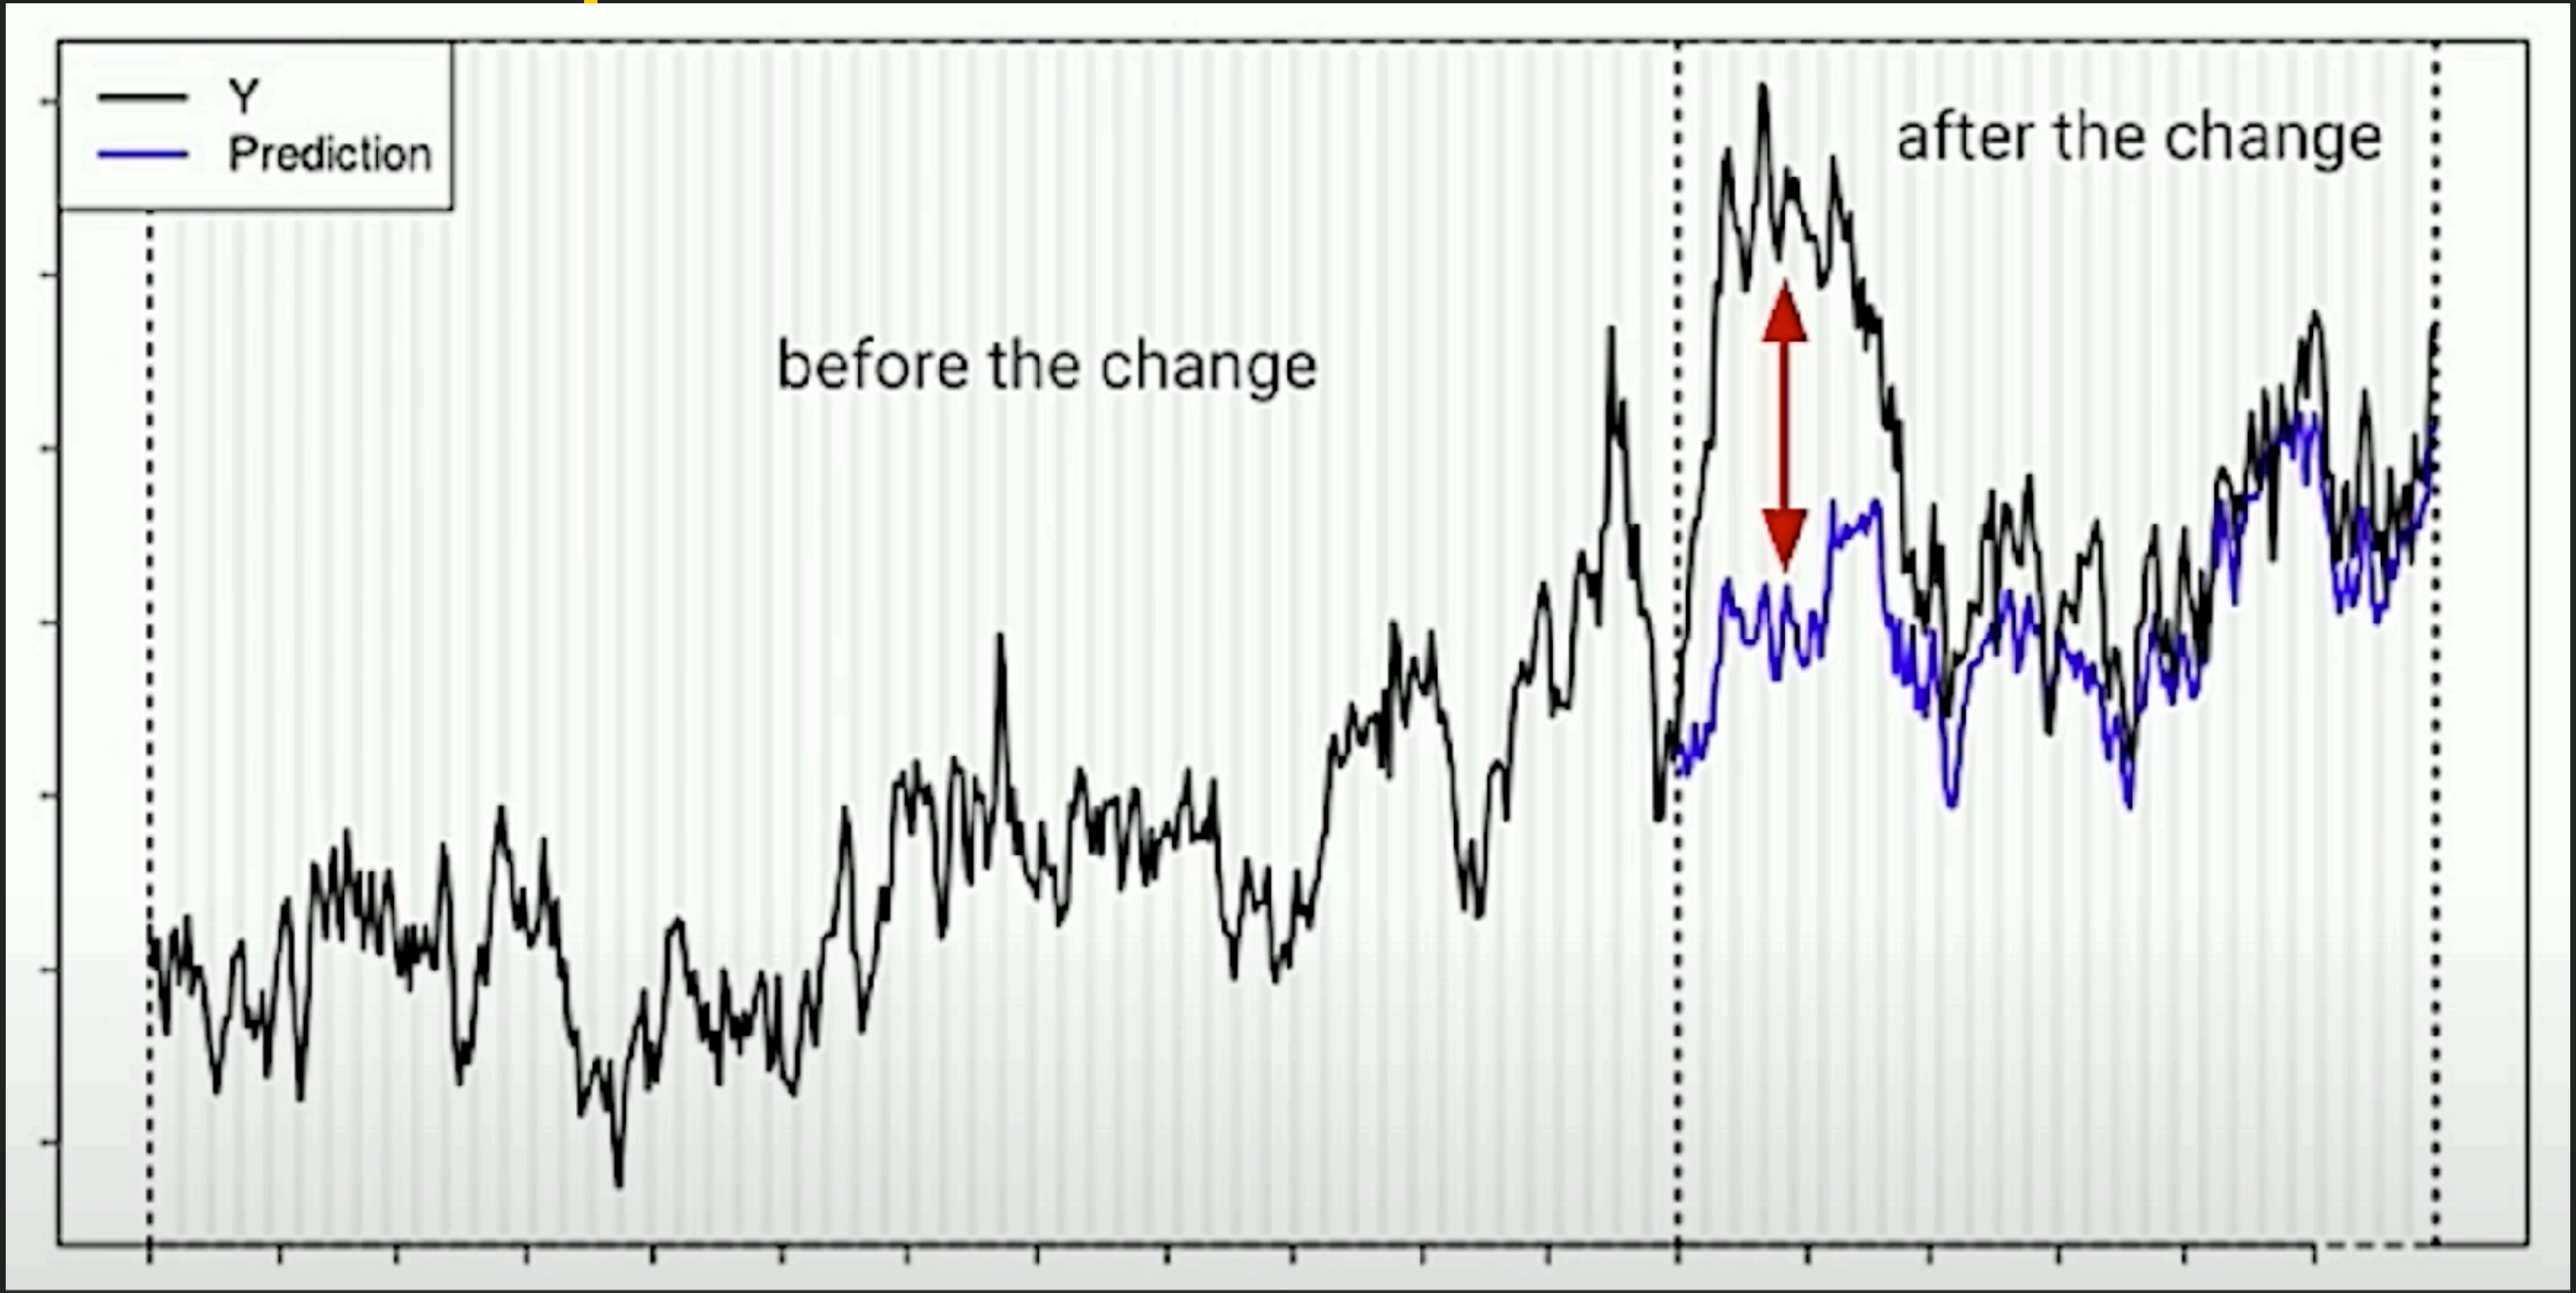

Inference

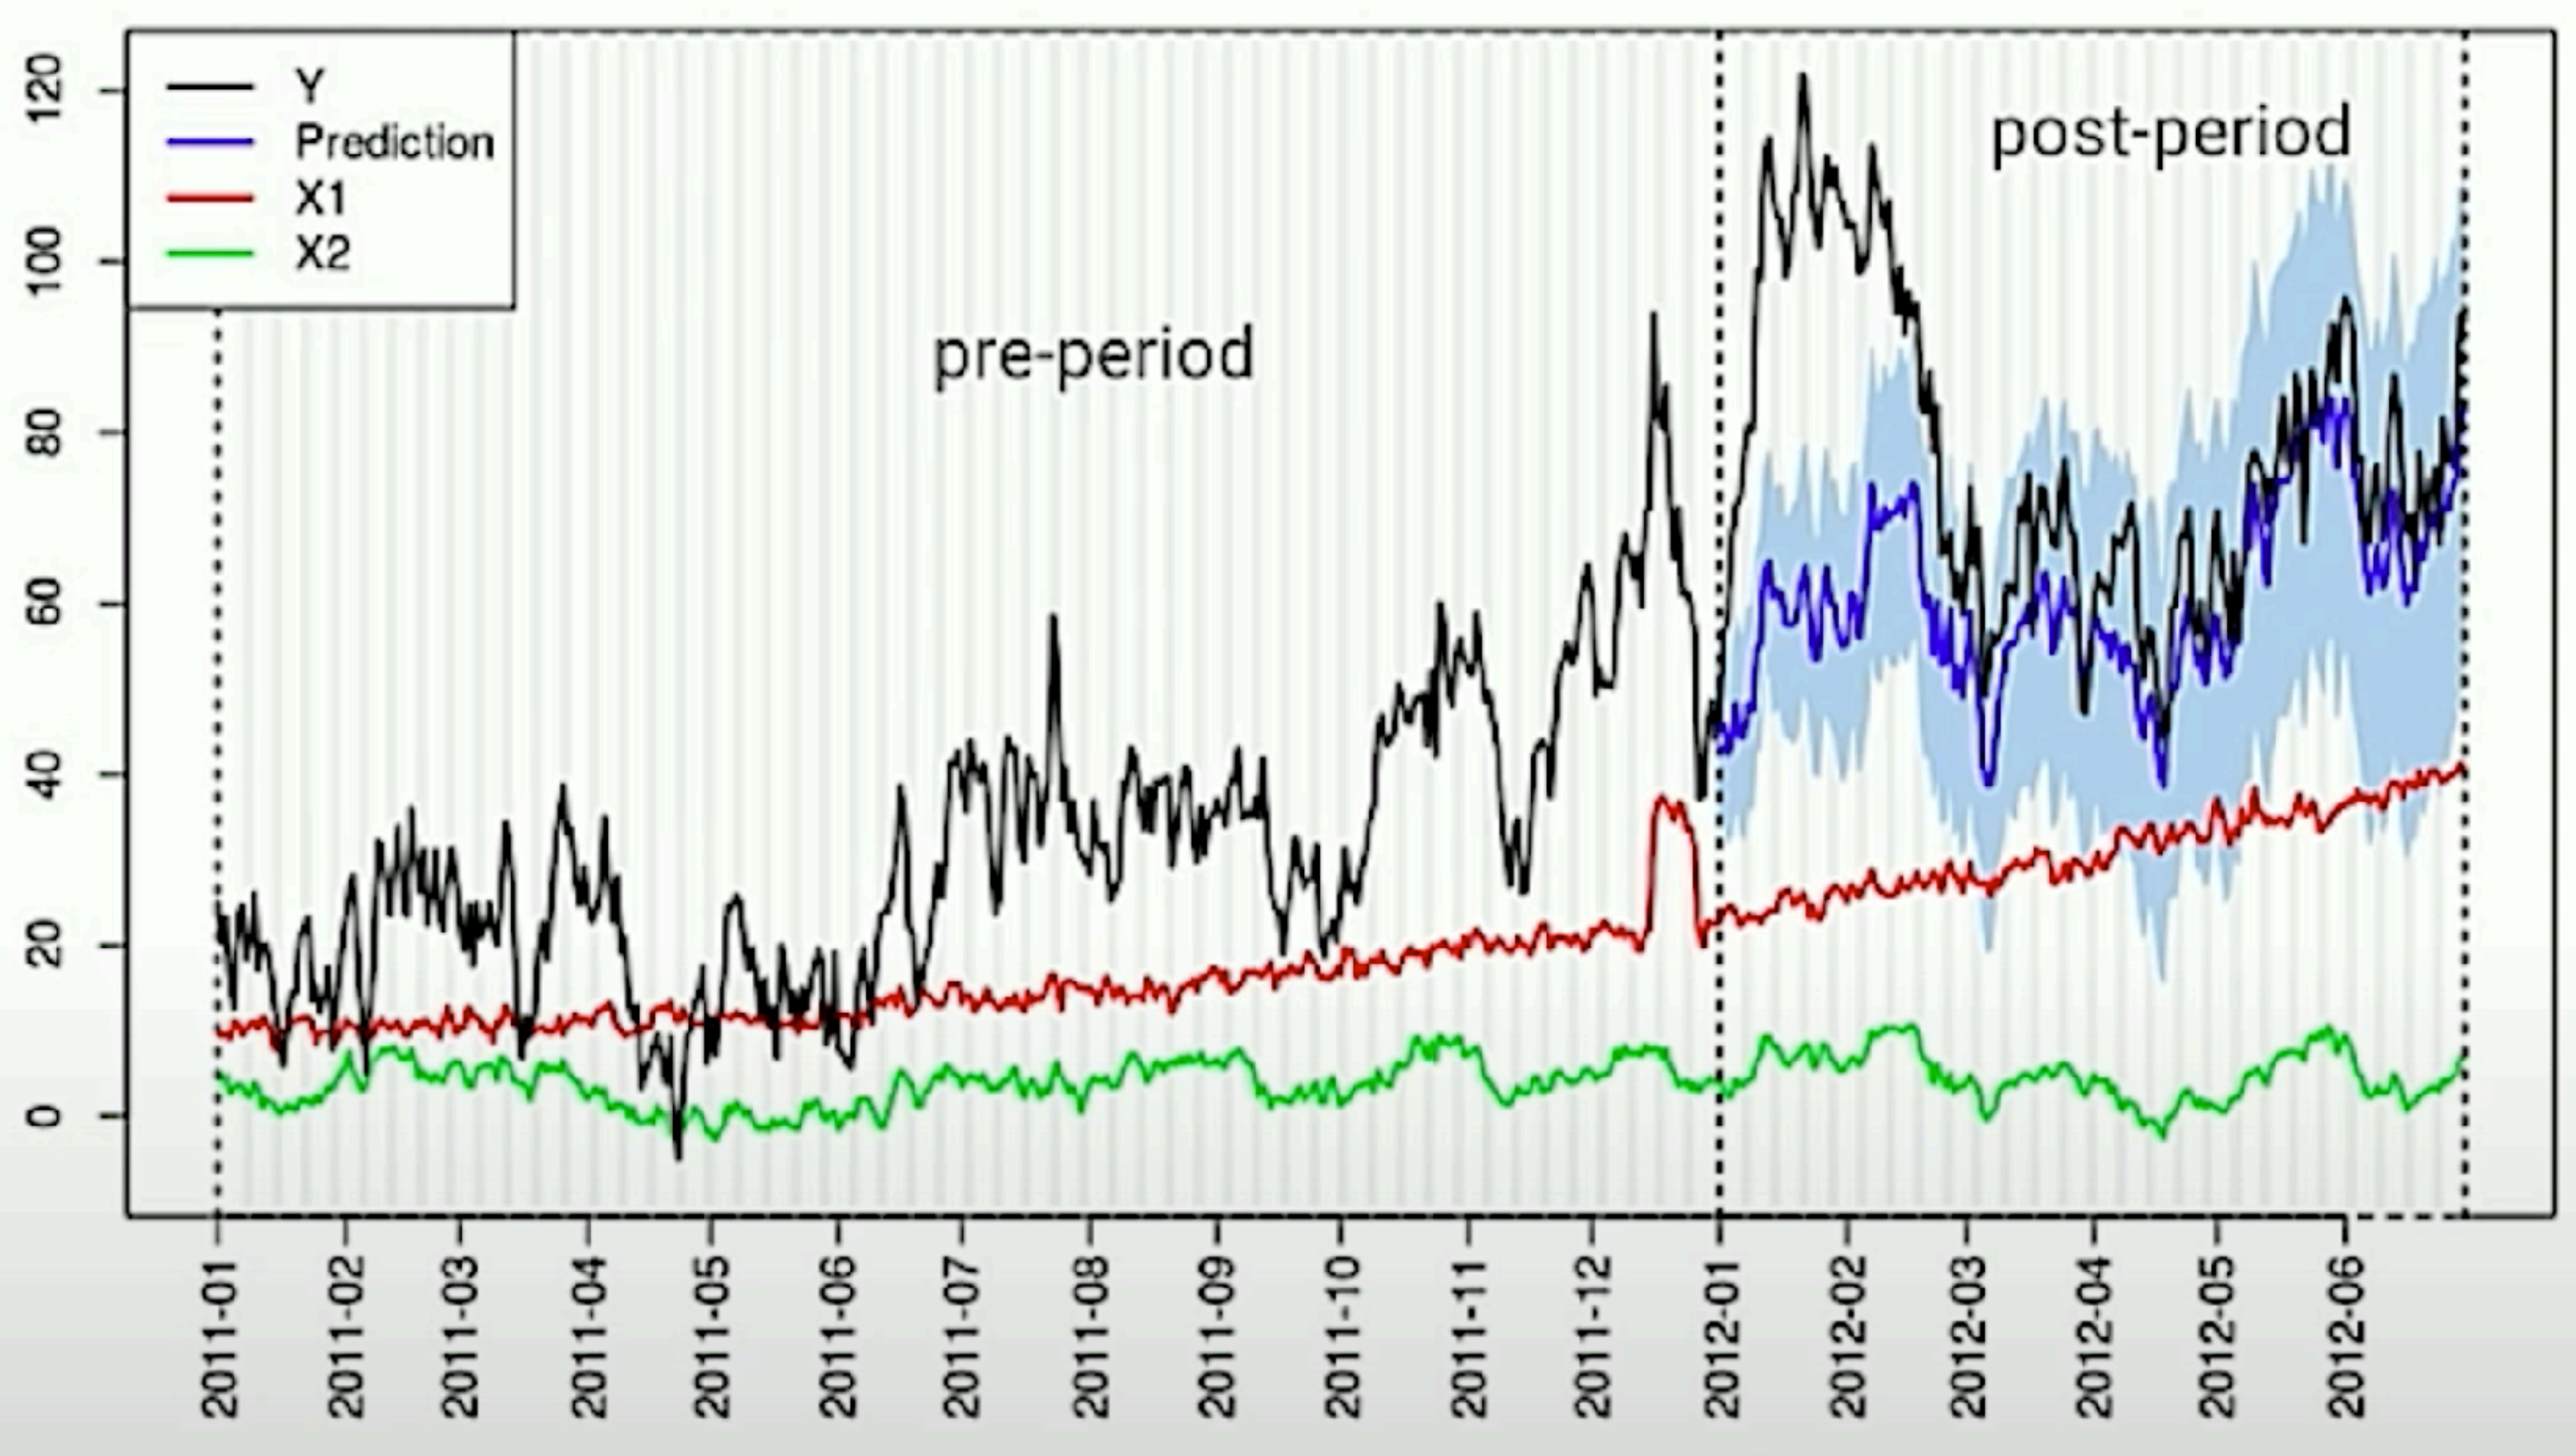

Y
Prediction
X1
X2
pre-period
post-period
0
20
40
60
80
100
120
2011-01
2011-02
2011-03
2011-04
2011-05
2011-06
2011-07
2011-08
2011-09
2011-10
2011-11
2011-12
2012-01
2012-02
2012-03
2012-04
2012-05
2012-06

Estimating the Effect

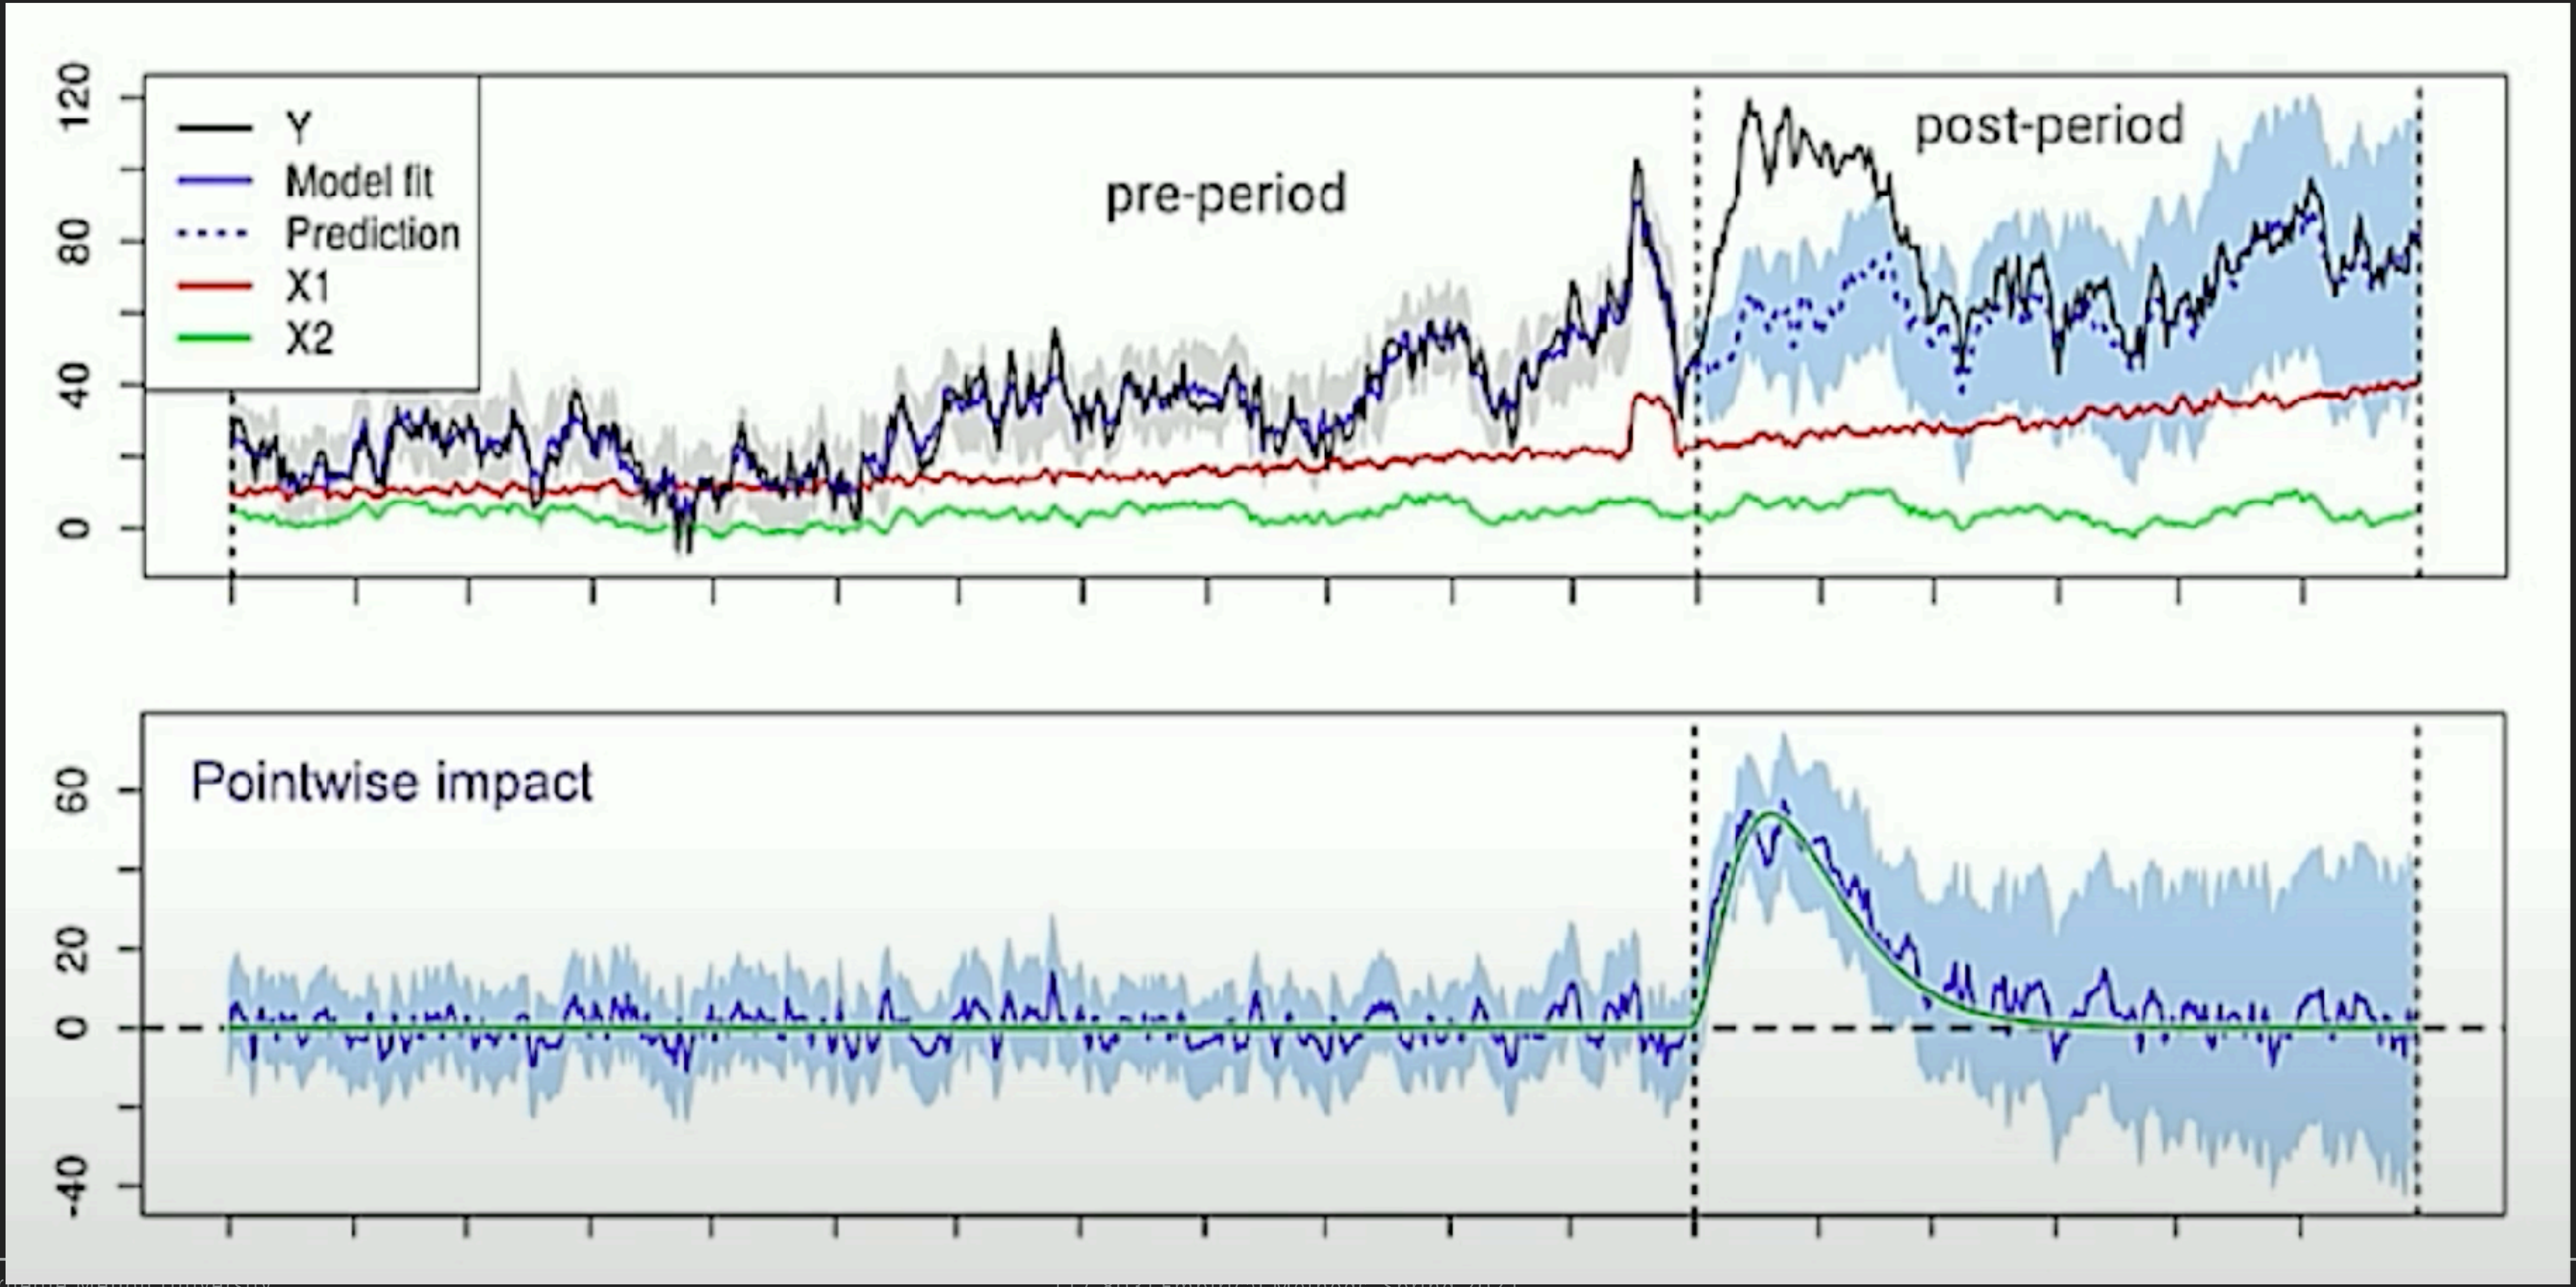

Causal Effect of Advertising on Clicks

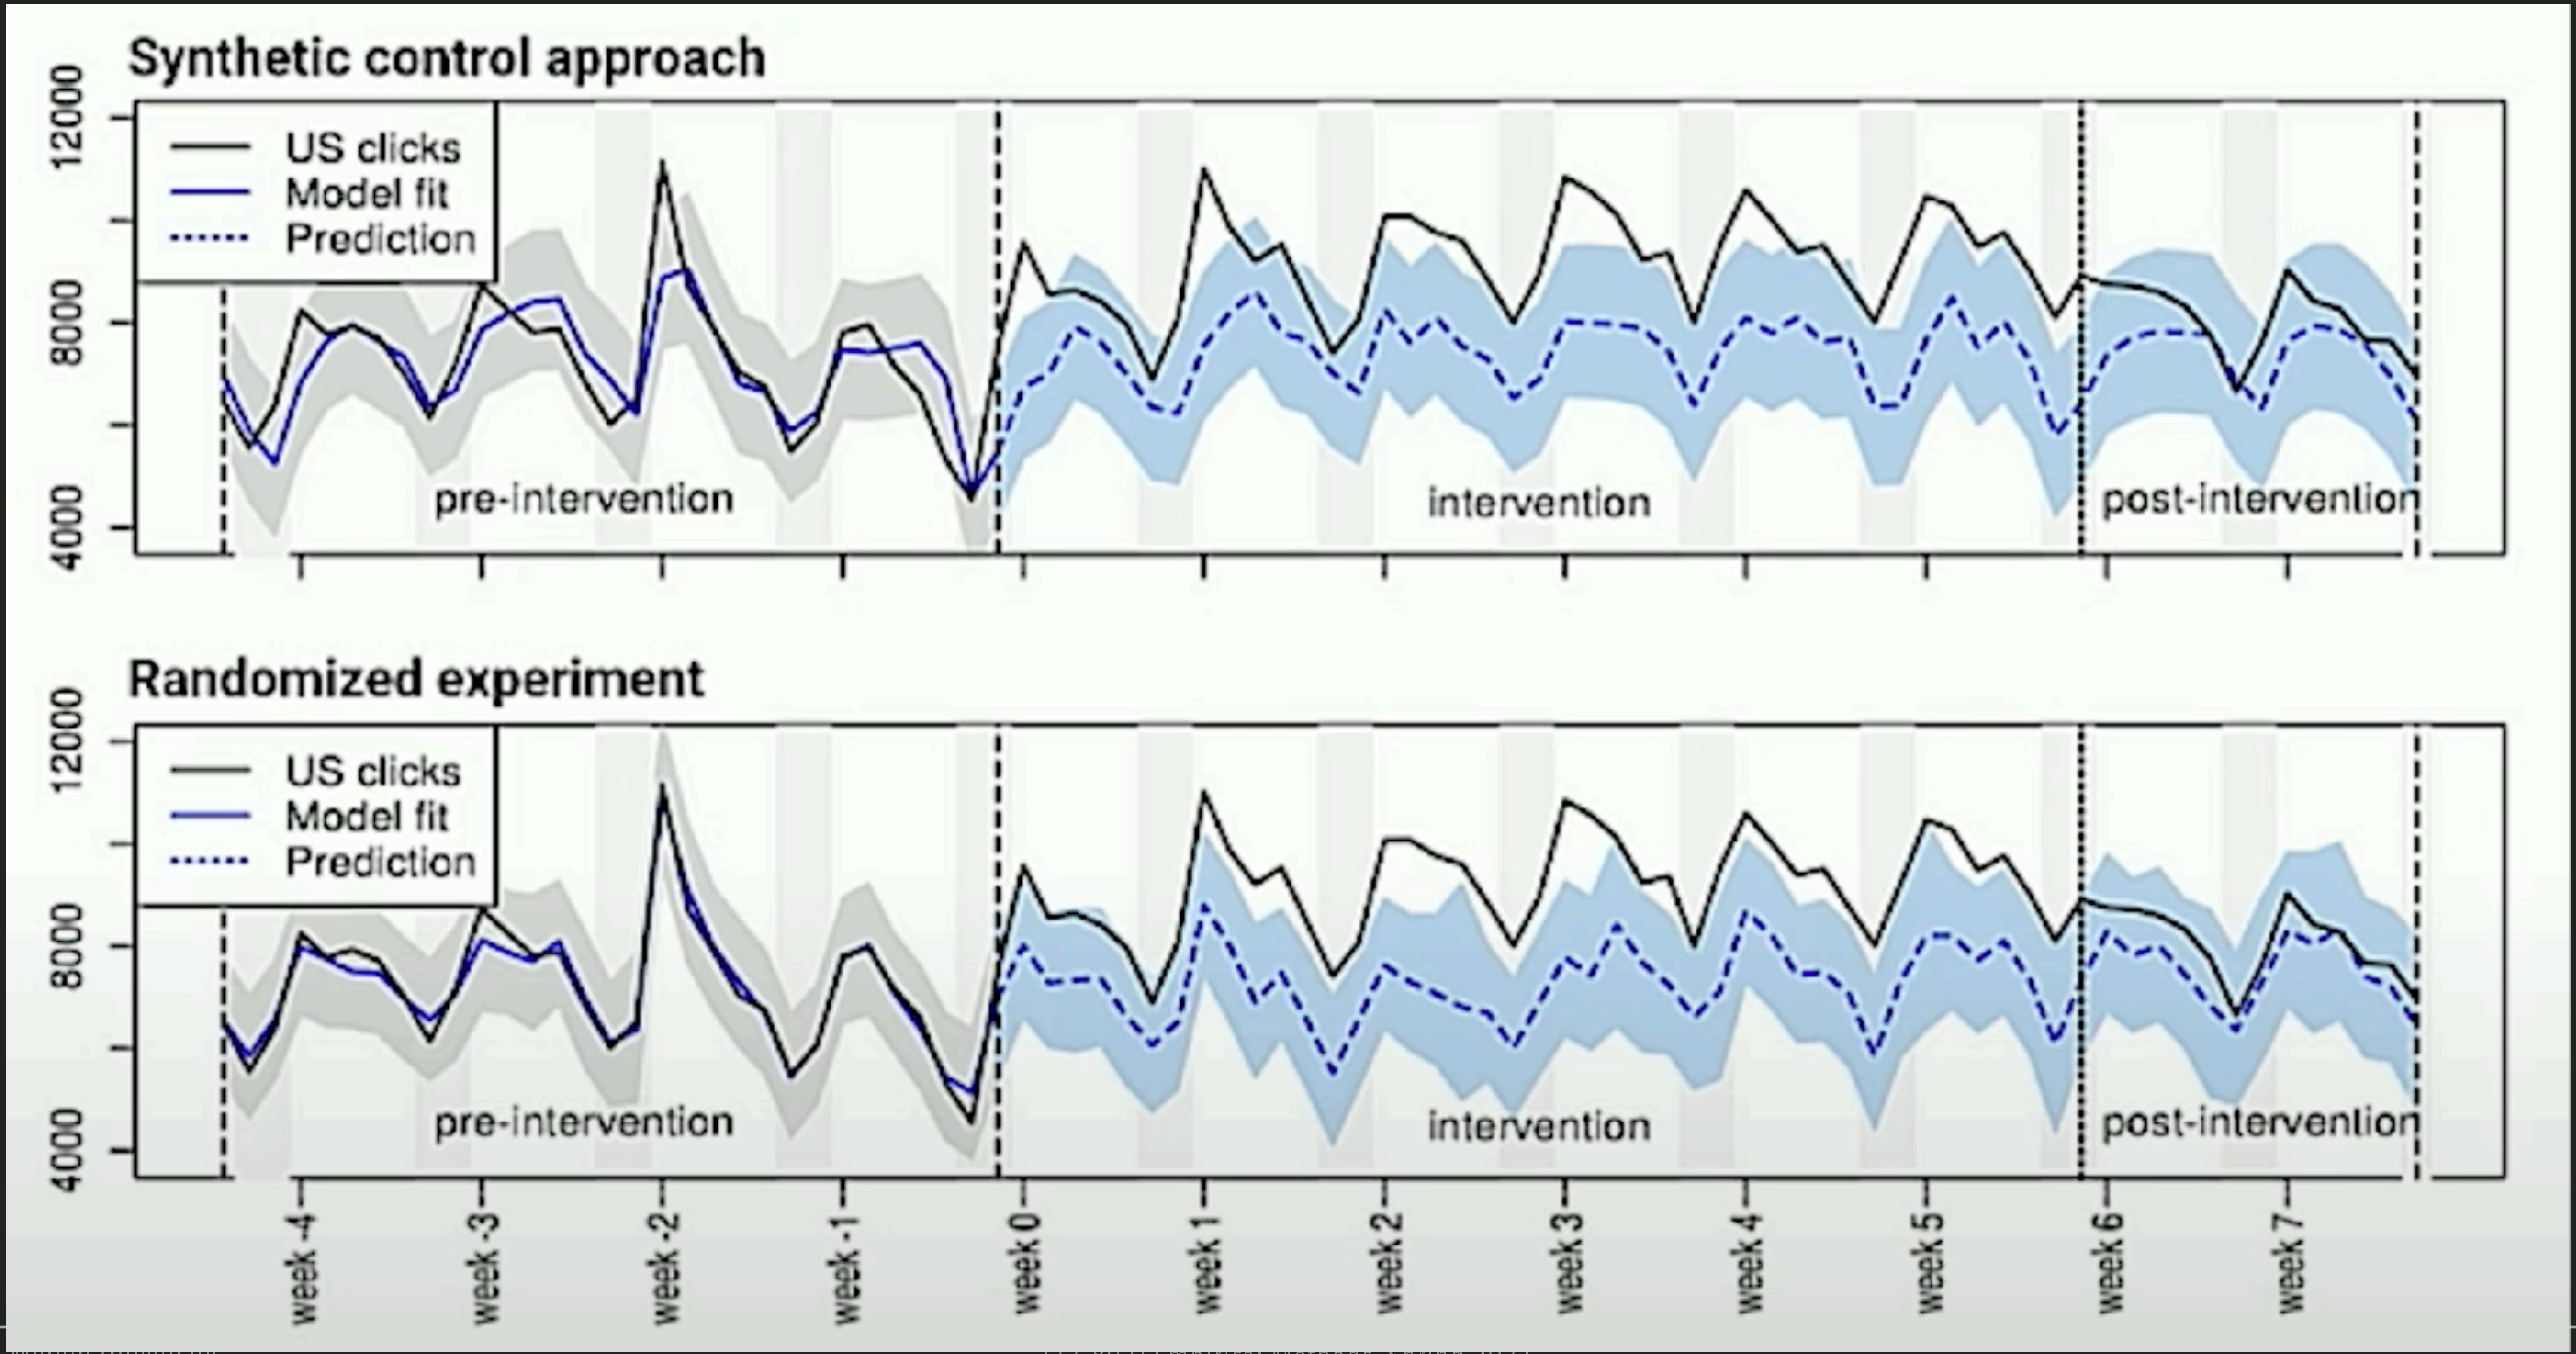

Credits

- Graphics: Dave DiCello photography (cover)
- Pearl, J., Glymour, M., & Jewell, N. P. (2016). Causal inference in statistics: A primer. John Wiley & Sons.
- Rajendran, P. Causal Inference using Difference in Differences, Causal Impact, and Synthetic Control. https://towardsdatascience.com/causal-inference-using-difference-in-differences-causal-impact-and-synthetic-control-f8639c408268
- Foundations of Program Evaluation III https://ds4ps.org/cpp-525-spr-2020/schedule/
- Gertler, P. J., Martinez, S., Premand, P., Rawlings, L. B., & Vermeersch, C. M. (2016). Impact evaluation in practice. The World Bank.
- Wing, C., Simon, K., & Bello-Gomez, R. A. (2018). Designing difference in difference studies: best practices for public health policy research. Annual review of public health, 39.
- See also:
  - https://www.rdocumentation.org/packages/did/versions/2.0.0
  - https://google.github.io/CausalImpact/CausalImpact.html