

I love this paper but it's barely been noticed.¹

Andrew Gelman

2 Jan 2023

What happens when you drop a rock into a pond and it produces no ripples?

My 2004 article, Treatment Effects in Before-After Data, has only 23 citations and this goes down to 16 after removing duplicates and citations from me. But it's one of my favorite papers. What happened?

It is standard practice to fit regressions using an indicator variable for treatment or control; the coefficient represents the causal effect, which can be elaborated using interactions. My article from 2004 argues that this default class of models is fundamentally flawed in considering treatment and control conditions symmetrically. To the extent that a treatment “does something” and the control “leaves you alone,” we should expect before-after correlation to be higher for controls than in the treatment group. But this is not implied by the usual models.

My article presents three empirical examples from political science and policy analysis demonstrating this point and proposes some statistical models. Unfortunately, these models are complicated and can be noisy to fit with small datasets. It would help to have robust tools for fitting them, along with evidence from theory or simulation of improved statistical properties. I still hope to do such work in the future, in which case perhaps this work will have the influence I hope it deserves.

References

Gelman, A. (2004), Treatment effects in before-after data. In *Applied Bayesian Modeling and Causal Inference from an Incomplete Data Perspective*, ed. A. Gelman and X. L. Meng, 195-202. London: Wiley.

¹ For a collaborative article, “What are your most underappreciated works?”, in *Econ Journal Watch*.