

Efficient DiD

First Open-Source

Implementation

diff-diff v2.7

Chen, Sant'Anna & Xie (2025)

Semiparametric efficiency bound

for staggered DiD

The Problem


Standard staggered DiD uses one
comparison per target effect --
leaving valid information unused.



Single 2x2 comparison (CS default)

The Insight

Under PT-All, multiple valid comparisons exist. EDiD finds the optimal combination.



$$w^* = \frac{1'(\Omega^*)^{-1}}{1'(\Omega^*)^{-1}1}$$

GLS optimal weighting

The Assumption

Parallel Trends for All Groups

All cohorts share common outcome trends,
treated or not.

PT-Post

Trends hold between treated
cohort and comparison group

= CS post-treatment ATT

PT-All

Trends hold across all
cohorts, treated or not

= EfficientDiD

When does PT-All hold?

- Staggered policy rollouts across regions
- Administrative or geographic phasing
- Treatment timing unrelated to anticipated effects

The Payoff

Achieves the semiparametric efficiency bound

-- tightest possible confidence intervals.

**Callaway-
Sant'Anna**



EfficientDiD



"Often exceeding 40% gains in precision"

-- Chen, Sant'Anna & Xie (2025)

The Code

```
from diff_diff import EfficientDiD

edid = EfficientDiD(pt_assumption="all")

results = edid.fit(data, outcome="y",

                   unit="id", time="t",

                   first_treat="g",

                   aggregate="all")

results.print_summary()
```

sklearn-like API

Safety Net

Under PT-Post, post-treatment $ATT(g,t)$

matches Callaway-Sant'Anna exactly

ATT point estimates identical

Corollary 3.2, Chen, Sant'Anna & Xie (2025)

Get Started

```
pip install diff-diff
```

github.com/igerber/diff-diff

[arXiv:2506.17729](https://arxiv.org/abs/2506.17729)

diff-diff v2.7

Difference-in-Differences for Python