

diff - diff

**Most methods require
a control group.**

Reality doesn't.

Heterogeneous Adoption Designs (HAD).

Now in diff-diff.


de Chaisemartin et al. (2026).

doi.org/10.48550/arXiv.2405.04465

diff-diff **v3.3.1**

Standard techniques subtract a control group.

What if there isn't one?



No control group, just different doses.

Sometimes **everyone** gets treated.

Universal Pricing Rollout

Every store moves to the new pricing; magnitude varies.

Industry-Wide Regulation

Every firm is subject to the new rule; exposure varies.

Company-Wide Training

Every employee receives the program; hours invested differ.

National Policy Floors

Every state implements the federal floor; intensity differs.

The HAD estimator.

For Heterogeneous Adoption Designs.

Recovers treatment effects when there's
no untreated comparison.

Anchor on the least-treated.

Treats them as a quasi-control group.

Recover the Weighted Average Slope.


Local-linear fit at the dose-support boundary.

Auto-detect the design path.

Three identification strategies, one API.

Returns WAS on continuous-at-zero designs; WAS_d_lower otherwise.

Use the least-treated as a quasi-control.



Local-linear at the dose boundary.


Weighted Average Slope.

(Design 1', $d_{\text{lower}} = 0$)

$$WAS = \frac{E[\Delta Y] - \lim_{d \downarrow 0} E[\Delta Y | D \leq d]}{E[D]}$$

Treatment effect per unit of dose,
weighted by the dose itself.

The HAD estimator reads the dose distribution.



One API, three identification strategies.

The Code.

Same sklearn-like API as every diff-diff estimator.

```
from diff_diff import HeterogeneousAdoptionDiD as HAD

result = HAD(design='continuous_at_zero').fit(

    data,

    outcome_col='outcome',

    dose_col='dose',

    time_col='period',

    unit_col='unit',

)

print(result.att)          # WAS estimate

print(result.conf_int)     # Bias-corrected 95% CI

print(result.design)       # 'continuous_at_zero'
```

Bias correction, three design paths,

event study - one fit() call.

Production-ready.

Bias-Corrected CIs

Calonico-Cattaneo-Farrell
ported in-house

Auto Design Detection

Three identification paths,
one API

Dynamic Event Study

Per-horizon estimates
with pointwise CIs

Survey Support

pweights, strata, PSU, FPC
via Binder TSL

Sup-t Bands

Simultaneous CIs across
event-study horizons

Pre-Test Diagnostics

QUG, Stute, Yatchew-HR,
joint workflow

Validated against R.

End-to-end match against DIDHAD v2.0.0 on continuous-at-zero designs.

R Parity vs DIDHAD

End-to-end match against the DIDHAD R package (v2.0.0).

Paper-Equation Parity

Theorem 1 / Equation 7 and Theorem 3 / Equation 11 implemented as written.

Bias-Corrected CI Bit-Identity

Uniform-weights matches the in-house nprobust port at atol=1e-14.

Monte Carlo Oracle Consistency

Recovers the known-tau DGP; coverage near nominal level.

Now in diff-diff.

The HAD estimator.

```
$ pip install --upgrade diff-diff
```

github.com/igerber/diff-diff

diff - diff

Difference-in-Differences for Python

diff-diff v3.3.1