Evaluating the Economic Impact ofNatural Disasters

A Causal-Inference Case Study of the Aceh Tsunami in Python

Study Guide Scope

- **Focus**: Evaluating disaster impact when the treatment is strictly observational.
- **Methods Covered**: **Dynamic Difference-in-Differences, Night-Lights Dose-Response, Synthetic Control, Spatial-HAC Standard Errors.**
- **Data Note**: Built on pedagogically calibrated synthetic data replicating the findings of **Heger & Neumayer (2019)**.

The Natural Experiment & Study Design

The Core Question

A decade later, was Aceh richer or poorer than it would have been without the tsunami?

The Natural Experiment

The tsunami destroyed capital and killed 130,000 people, but triggered USD 7.7 billion in reconstruction aid. We only observe the post-tsunami reality.

The Research Pipeline

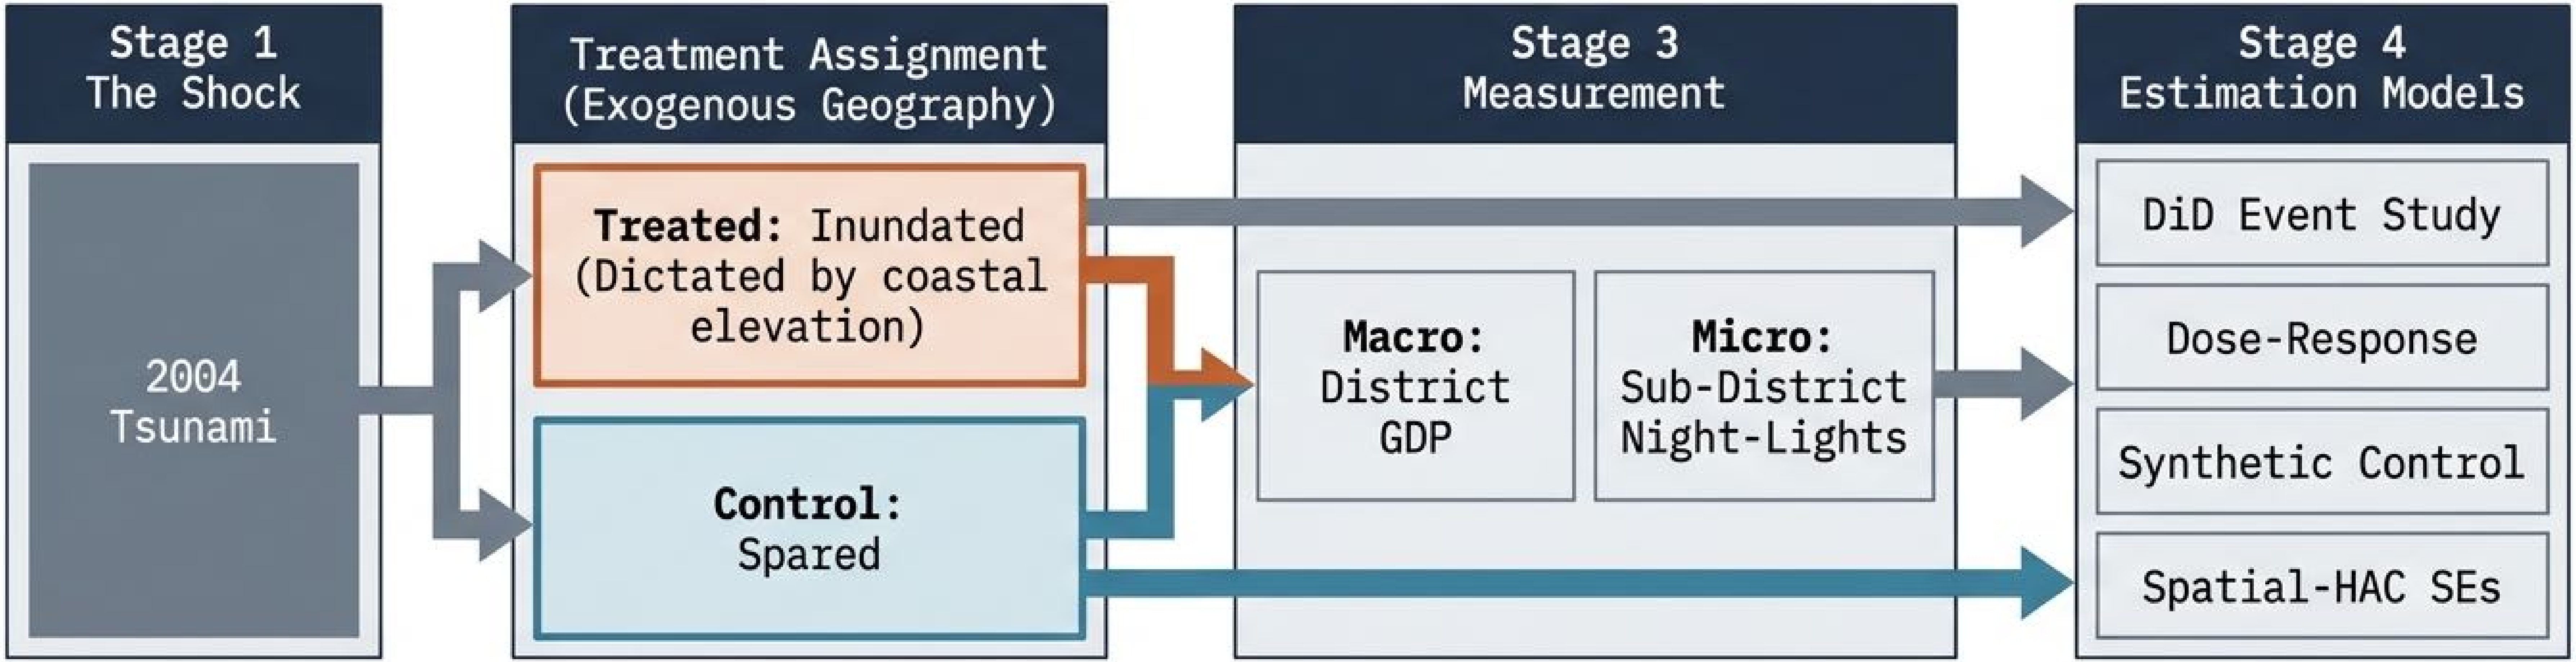

Identification Strategy: Treatment (inundation) was dictated strictly by coastal geography, entirely exogenous to a district's prior economic trajectory.

Core Causal Inference Lexicon

Concept	Formal Definition & Notation	The Aceh Example	Intuitive Analogy
Difference-in-Differences (DiD)	Change in Treated minus Change in Control.	Flooded growth (+0.0103) minus Control growth (-0.0022) = +0.0125.	Timing two runners on parallel tracks; credit coaching only with the extra burst.
Parallel Trends	Identifying assumption. Pre-treatment trends must match.	Pre-tsunami (2003–04) DiD coefficient is +0.0172 ($p=0.28$, insignificant).	Two boats drifting on the same current; only an engine (treatment) makes one pull ahead.
ATT	$E[Y(1) - Y(0) \mid D=1]$. Average Treatment on the Treated.	The +18.3% gap is the effect on flooded districts, not any random district.	The bonus speed measured on the specific car that received the engine tuning.
Counterfactual	Unobserved potential outcome $Y(0)$ for treated units.	Estimated path of the no-tsunami “Synthetic Aceh”.	The parallel-universe Aceh where the wave never arrived.

NotebookLM

Data & Measurement Strategy

Macro: District-Level Analysis

- **Metric:** Annual growth rate of real district GDP (excluding volatile oil/gas).
- **Treatment:** Binary On/Off Flooded Dummy.
- **Sample:** N = 125 districts. 10 flooded vs 115 controls.
- **Purpose:** Establishes aggregate economic trajectory.

Micro: Sub-District-Level Analysis

- **Metric:** DMSP-OLS Satellite Night-Lights (Log Luminosity).
- **Treatment:** Continuous Dose (Share of Population/Area Flooded).
- **Sample:** N = 276 Aceh sub-districts.
- **Purpose:** Captures exact localized intensity where GDP is too coarse.

Night-Lights Transformation Equation

$$NL_{ct} = \log\left(\sum(DN_{nct} + 0.001)\right)$$

Note: Summing keeps the measure comparable to GDP totals. Log transformation tames the heavy right-skew of brightness. The `0.001` constant prevents breaking the log at zero.

Method 1: Difference-in-Differences

The Equation Mechanics

$$\Delta Y_{it} = \beta_1 D_i 1[t \in \text{pre}] + \beta_2 D_i 1[t = 2005] + \beta_3 D_i 1[t \in \text{recovery}] + \beta_4 D_i 1[t \in \text{post}] + \alpha_i + \gamma_t + \varepsilon_{it}$$

D_i = Flooded Dummy | α_i = District Fixed Effect | γ_t = Year Fixed Effect | **Baseline** = 2000-02 (omitted)

The Python Implementation

```python
# Library: pyfixest (Stata-flavored fast fixed-effects)
model = pf.feols('gdp_growth ~ D_pre + D_2005 + D_recov + D_post | district_id + year', data)
```

Insight: Everything to the right of the '|' is computationally absorbed as fixed effects. This prevents the generation of massive dummy-variable matrices, making estimation highly efficient.

Method 1: Difference-in-Differences (Event Study)

Event study Plot

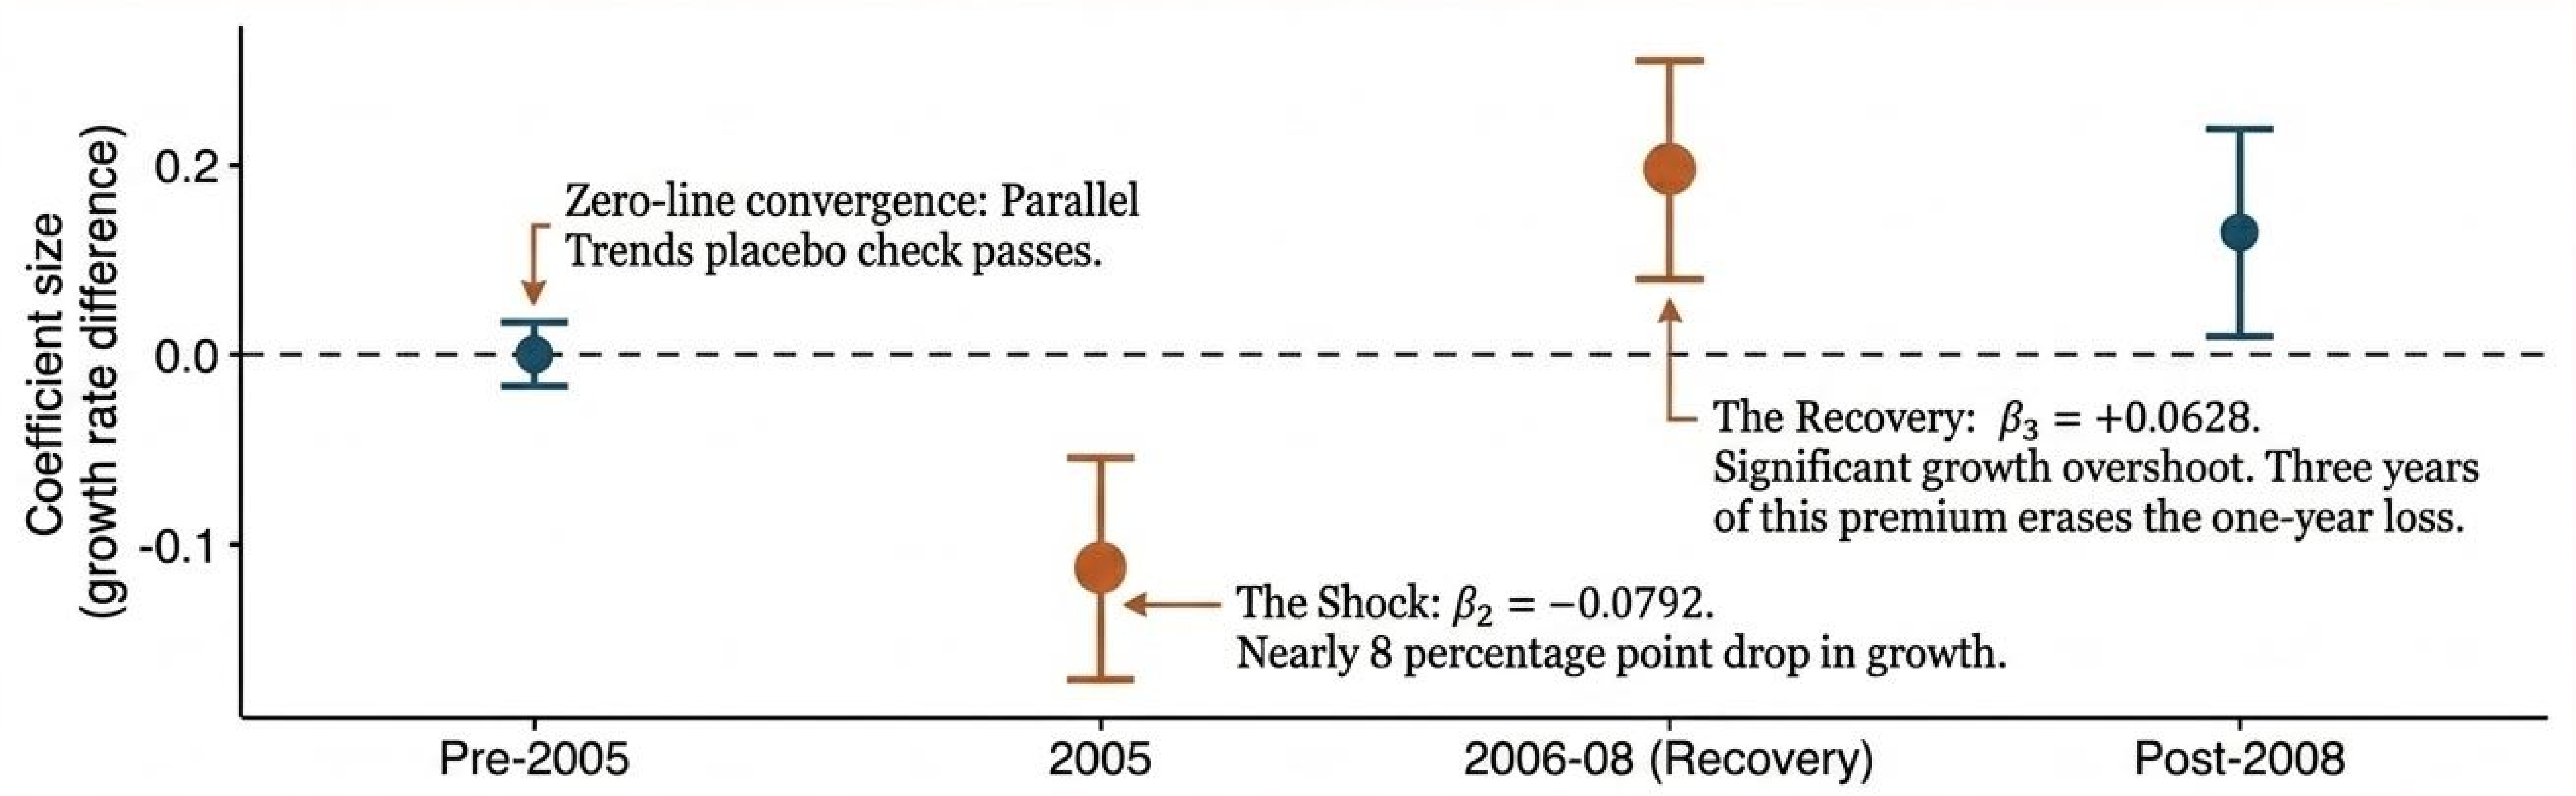

The Narrative Result

Sustainable recovery beyond the counterfactual trend.

Per-Capita Check

Did people just leave? Re-running with GDP per capita yields a recovery coefficient of +0.0827. The economic rebound is real, not just an arithmetic artifact of a shrunken population denominator.

Method 2: Dose-Response Using Night-Lights

The Conceptual Metaphor

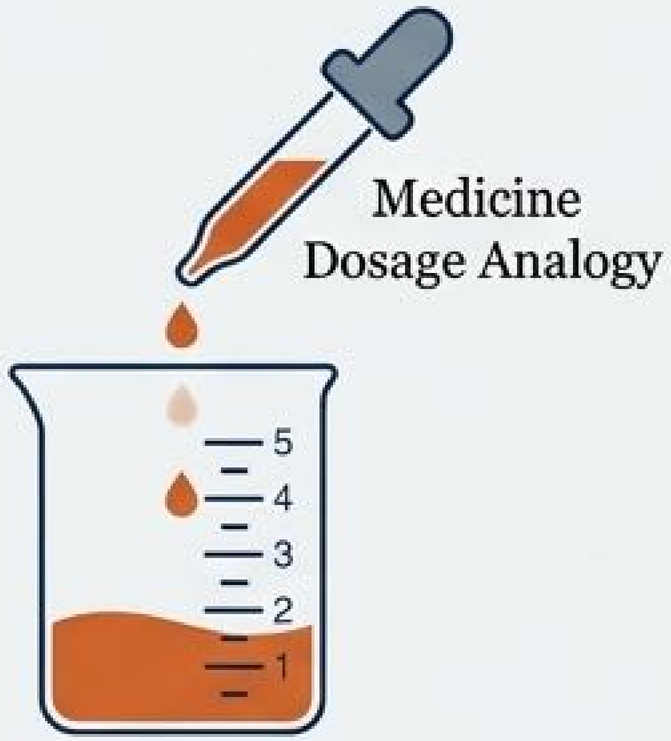

Treating exposure not as an on/off switch, but as a dial. A sip does little; the full dose moves the needle. If only the largest doses show an effect, the drug works, but average effects hide where the action is.

Context Note: In 2004, flooded coastal sub-districts were 2.5x brighter than non-flooded ones.

The Regression Mechanics

The Methodological Pivot

Instead of a binary `D_i` dummy, substitute continuous measures of destruction:

- `share_pop_flooded`
- `share_area_flooded`
- `flood_intensity_quintile`

Interact these continuous metrics with the recovery period to test if locations hit harder rebounded proportionally harder.

Interpreting the Dose-Response

The Continuous Effect

+0.0160

Coefficient on population-share flooded. Each additional unit of population flooded yields specific extra luminosity growth during reconstruction.

The Quintile Breakdown

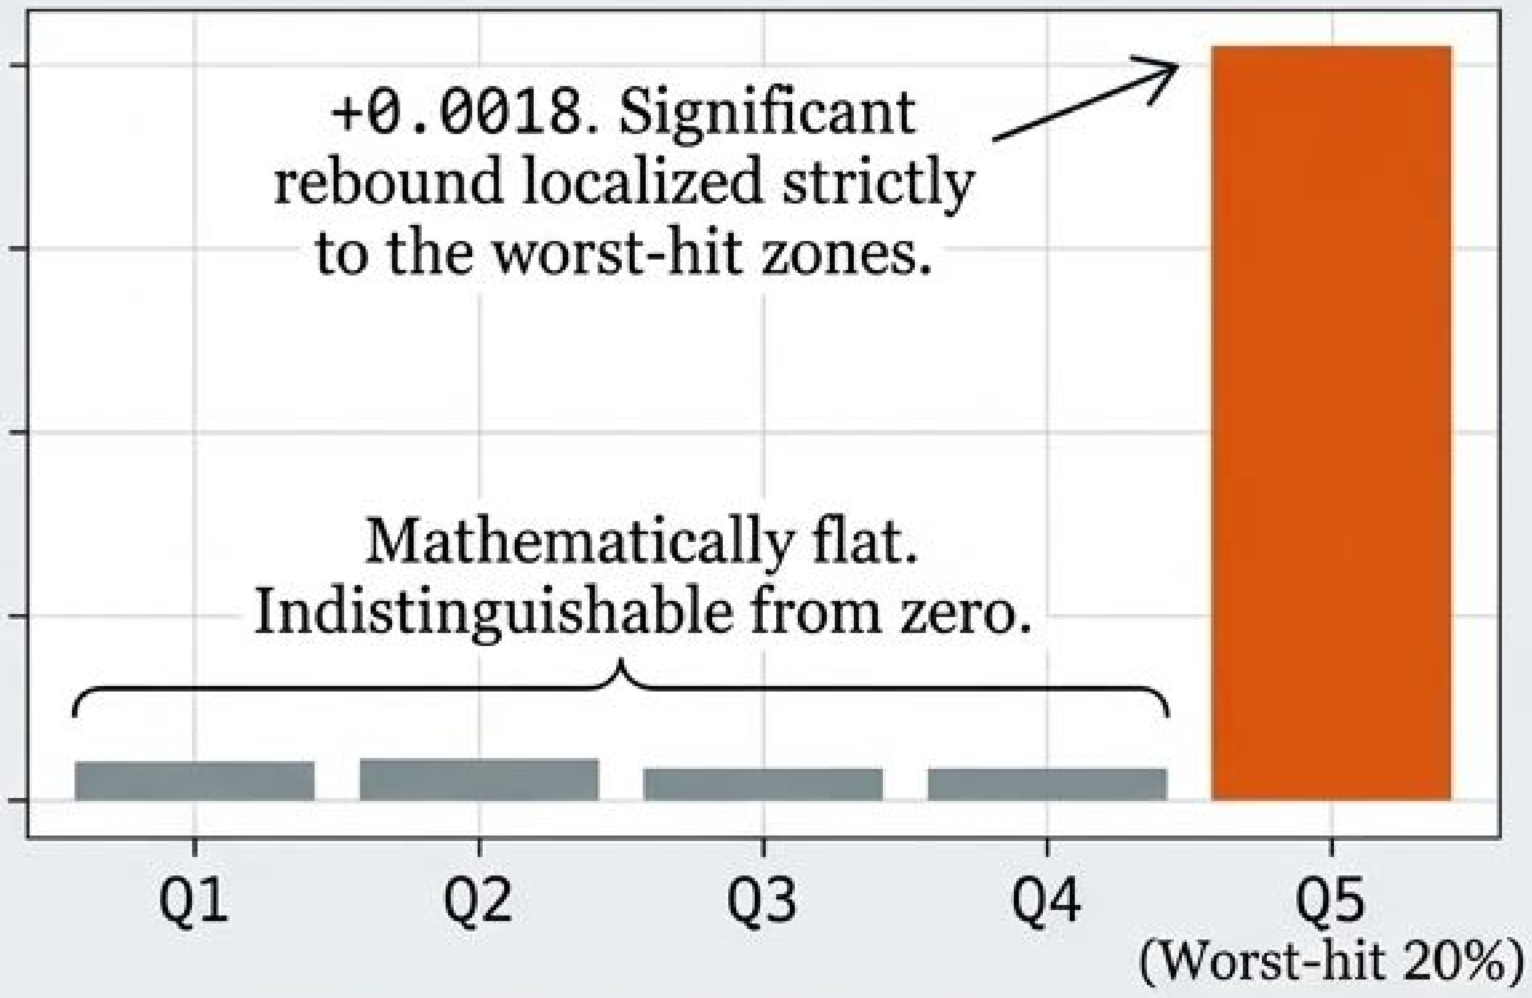

Takeaway: The 'average' effect hides the spatial reality. Reconstruction capital aggregated almost exclusively where damage was absolute.

Method 3: Synthetic Control Mechanics

Building a Bespoke Counterfactual

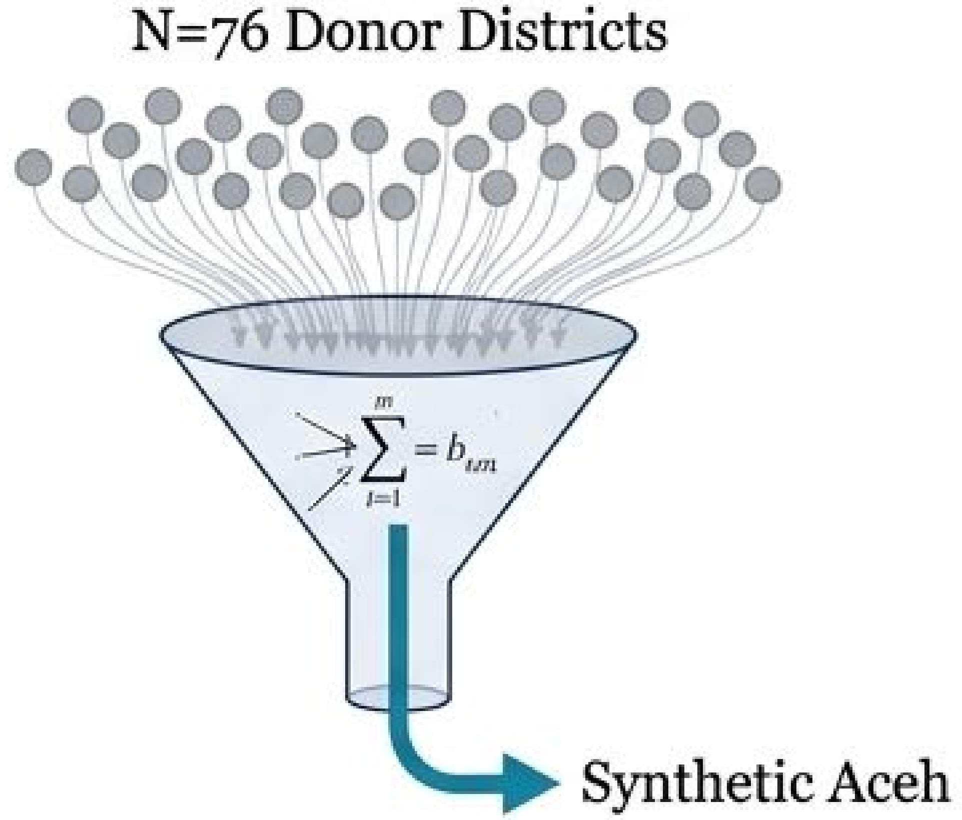

Goal: Build a bespoke counterfactual tracking pre-tsunami Aceh perfectly.
Implementation: Computed via `mlsynth` using the Abadie-Diamond-Hainmueller method.
Math: Pick non-negative weights summing to 1 to minimize pre-treatment RMSE.

Robustness of Weights

Top 6 donors carry 62% of the weight. The maximum single weight is only 0.13.
No single district dominates, preventing a fragile, single-lookalike counterfactual.

Synthetic Control: The Magnitude of Divergence

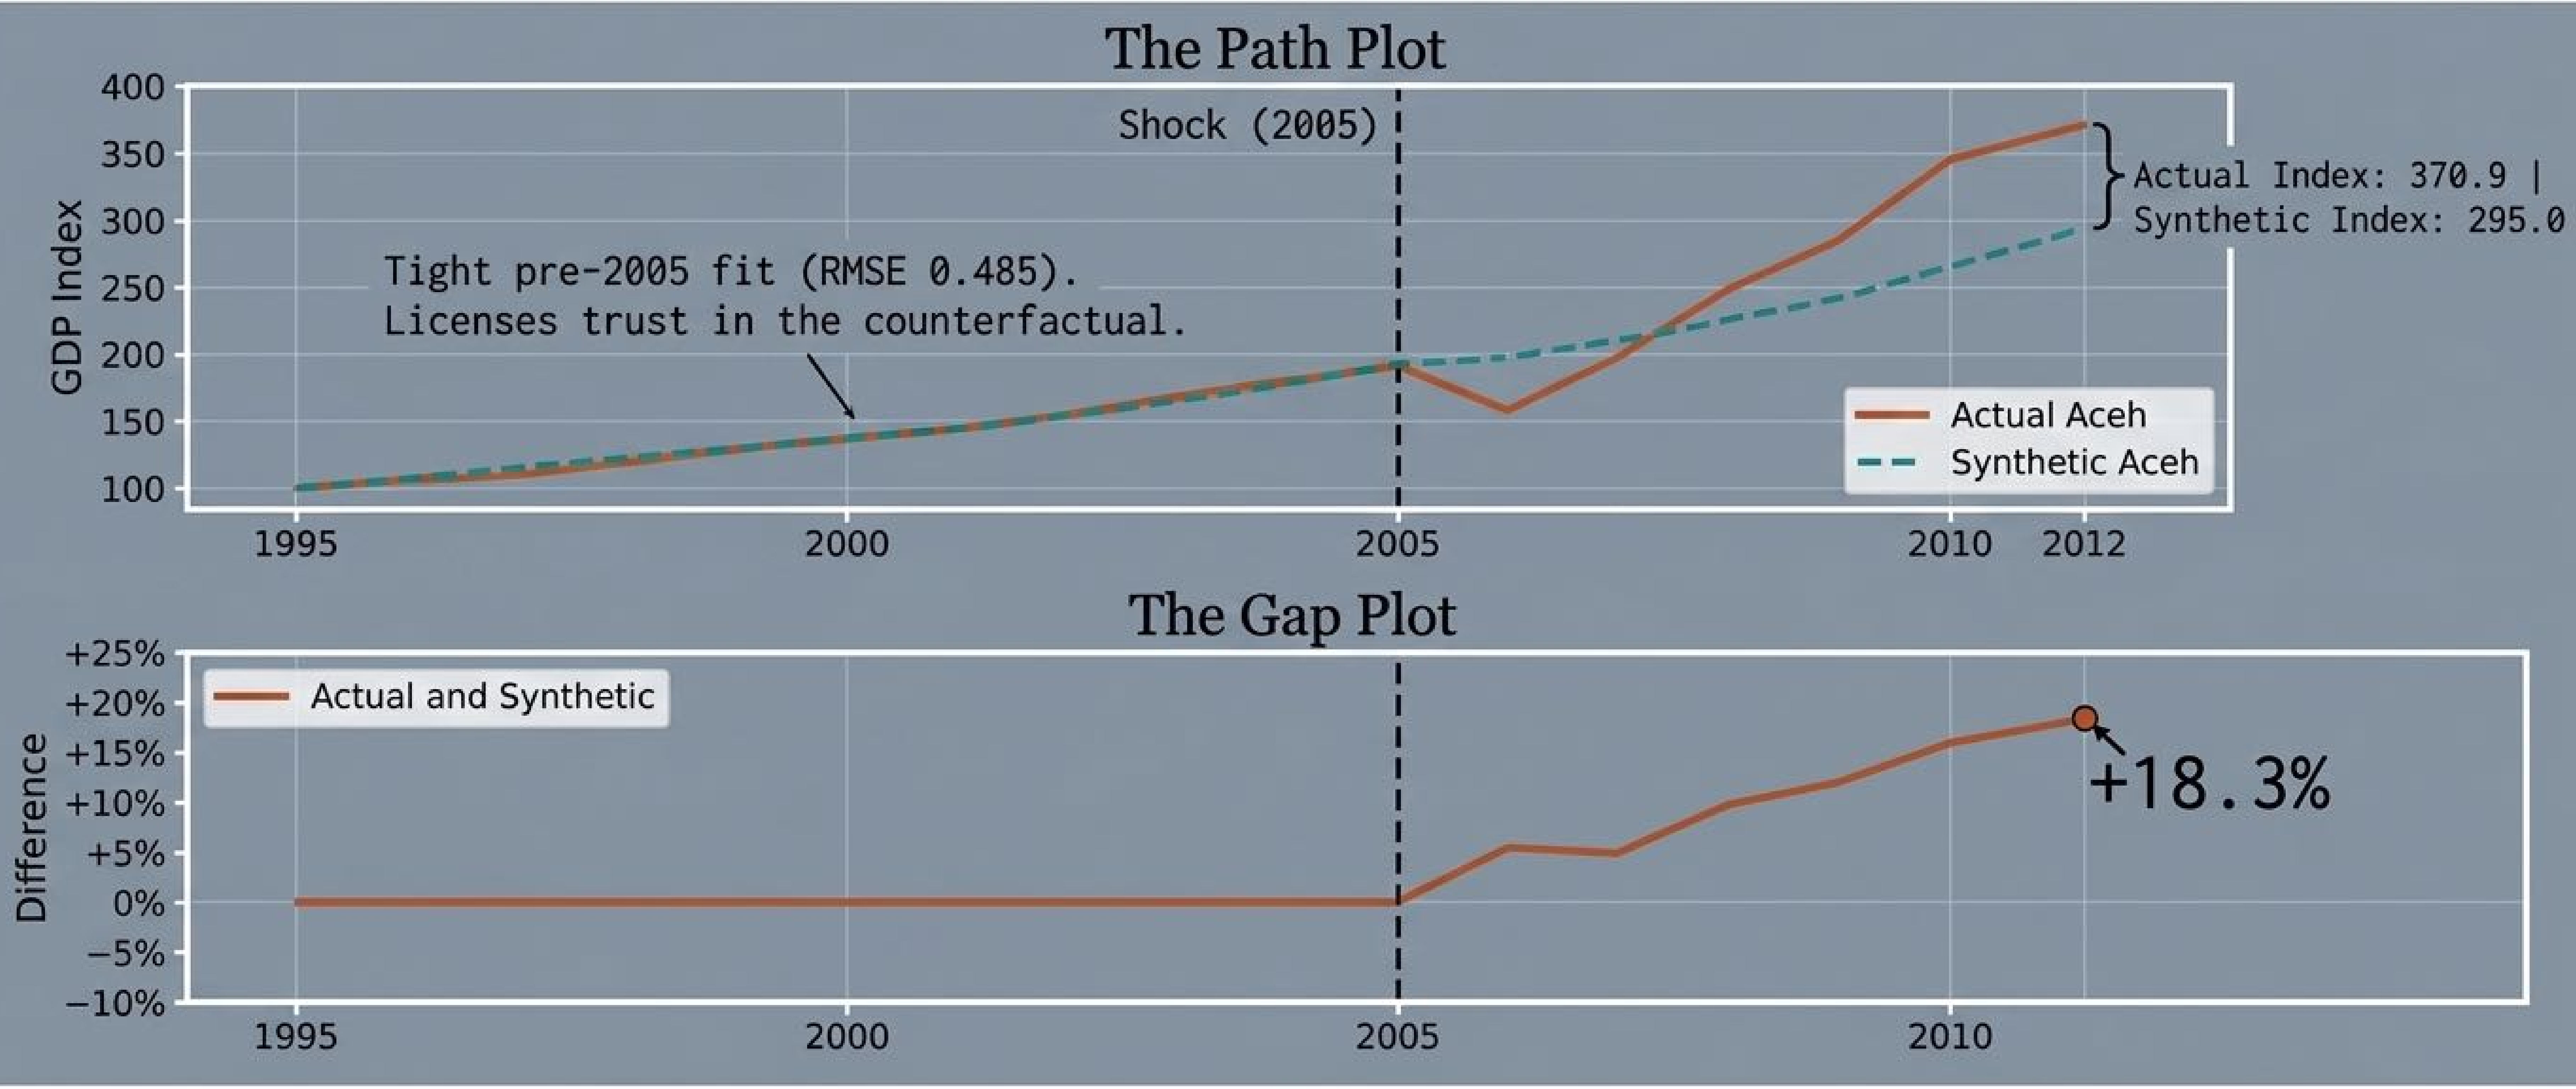

Takeaway: Quantifies the absolute magnitude of sustainable recovery beyond the counterfactual trend. Independent confirmation of the DiD results.

The Methodological Threat: Spatial Clustering

The Spatial Problem

Tobler's First Law: Near things are related. Shocks to tightly clustered districts (like a shared reconstruction boom) are not independent draws.

The Consequence of 'Naive' Errors

Default standard errors assume every observation is independent. If nearby districts share the same shock, naive errors over-count the amount of truly independent information. This makes standard errors artificially small and statistical confidence spuriously high.

The Diagnostic Check

`Moran's I: +0.065 (p = 0.003)`

Conclusion: Confirms highly significant spatial autocorrelation in the residuals. A statistical correction is mandatory.

The Fix: Conley Spatial-HAC Standard Errors

Method	Point Estimate	Standard Error	T-Statistic	Significance
Naive Errors	+0.0628	0.0146	> 4.00	1% (***)
Conley-HAC	+0.0628	0.0244	2.57	5% (**)

1.68x Inflation

Implementation Mechanics

Deploying Conley Spatial-HAC standard errors explicitly accounts for spatial correlation (districts clustered within ~100km) and serial correlation (a district correlated with itself over time).

The Lesson in Honest Inference

The point estimate (+0.0628) never moves. But the standard error inflates massively. The result remains statistically significant, but the spurious 1% certainty is eliminated. This is what honesty in causal inference looks like.

Robustness Checks: Placebos & Heterogeneity

The Spatial Placebo (Falsification Test)

- **Test:** Pretend the districts neighboring the flood were treated, omitting the actually flooded ones.
- **Result:** Coefficients are completely flat (approx. 0) and statistically insignificant.
- **Proof:** Ensures effects are specific to actual flood water destruction, not generalized regional economic spillovers.

Internal Heterogeneity (Urban vs Rural)

- **Rural Districts:** Took the brunt of the destructive 2005 shock (`-0.0883`) with a modest subsequent rebound.
- **Urban Districts (Cities):** Avoided the deep initial shock but captured massive reconstruction rebounds (`+0.1226`) due to capital concentration.
- **Caveat:** High statistical fragility here as there are only N=2 treated city units.

Synthesis: Methodological Triangulation

DiD Event Study

Proves temporal alignment and parallel pre-trends.

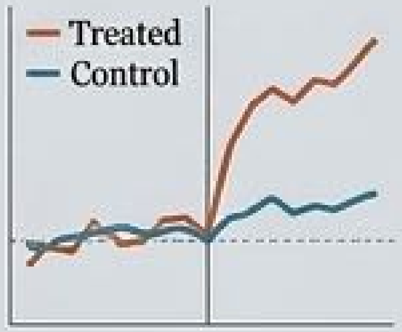

Addresses Blind Spot: Temporal Specificity

Dose-Response

Proves the effect scales precisely with physical exposure (worst-hit 20% saw the most light growth).

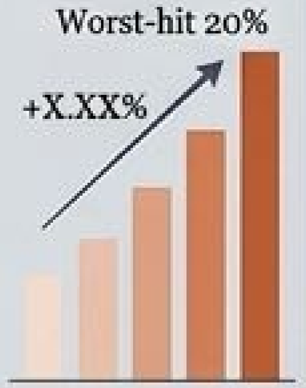

Addresses Blind Spot: Intensity & Scaling

Four methods converge on a single, robust story.

Synthetic Control

Removes reliance on general control trends; provides exact magnitude of of the shift (8.3%). (+18.3%).

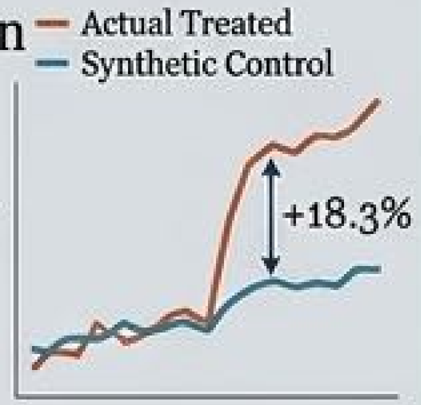

Addresses Blind Spot: Control Group Selection

Conley Spatial SEs

Validates statistical honesty. Ensures results survive despite severe geographic clustering.

Addresses Blind Spot: Spatial & Serial Correlation

Executive Summary & Academic Limitations

Numbers to Remember

Output Shock (2005):
-7.9%

Recovery Premium (06-08):
+6.3% / yr

Long-Term SC Gap (2012):
+18.3%

Spatial SE Inflation:
1.68x

The Lesson & Limits

The Economic Lesson

A localized catastrophe paired with massive, well-governed mega-reconstruction (funding >150% of damages) can shift long-run trajectories permanently higher.

Academic Limitations

- **Observational**: Parallel trends is a required assumption, not randomized proof.
- **Small N**: With only 10 treated districts, point estimates remain fragile.
- **Pedagogical Focus**: Relying on synthetically calibrated data engineered to teach causal Python methods.