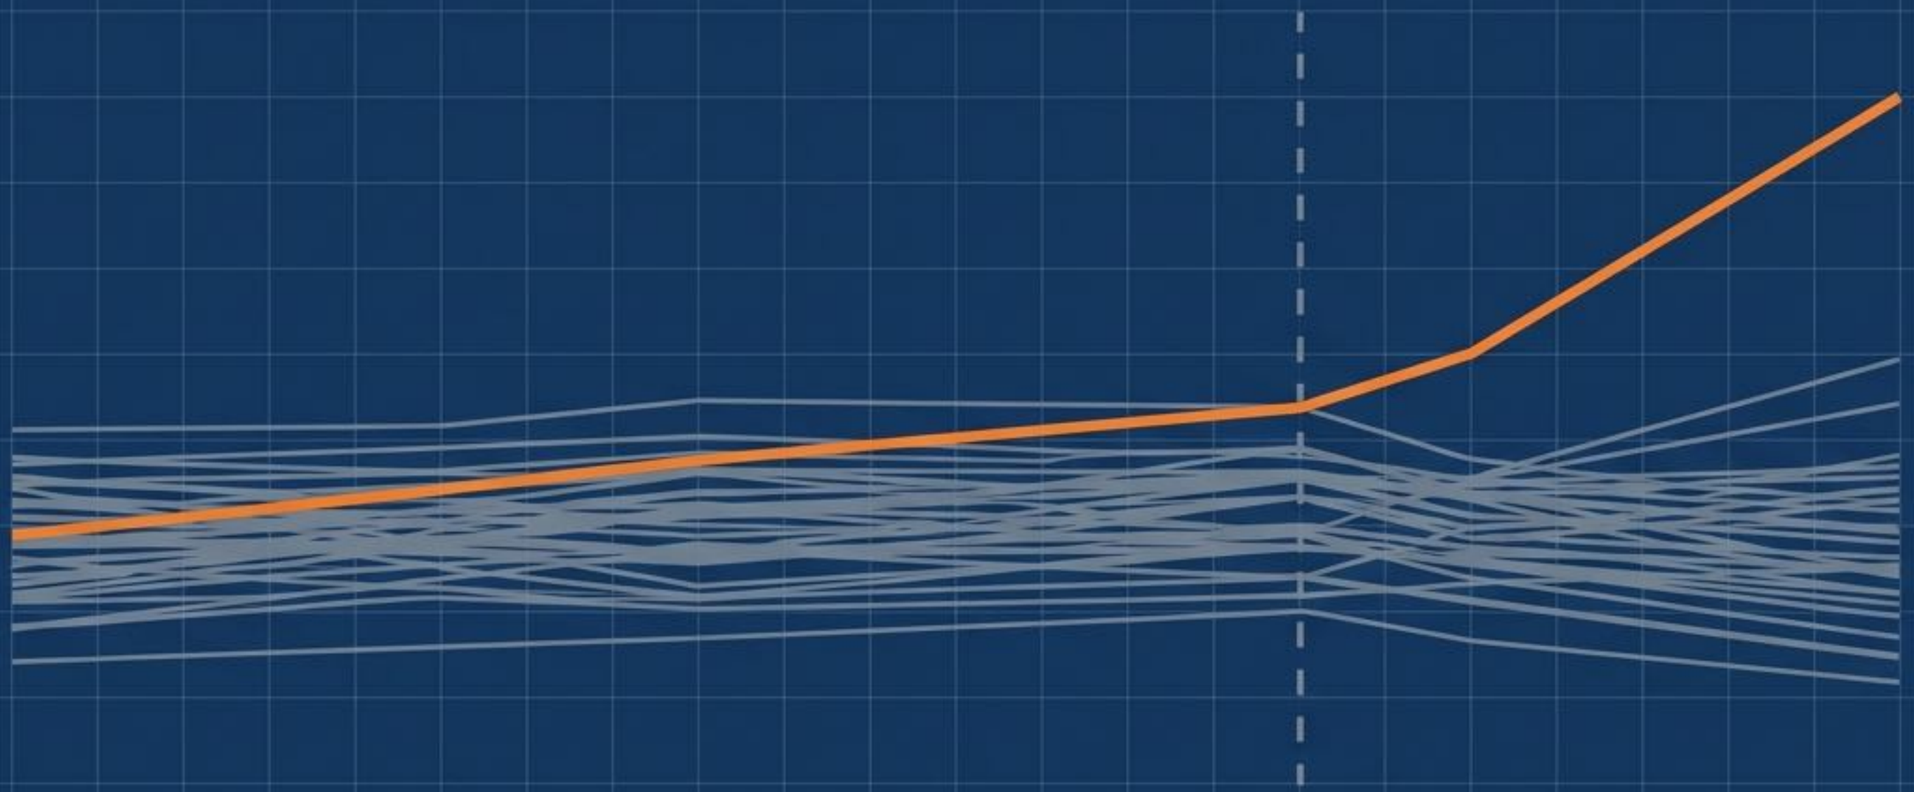

From DiD to SDID: A Ladder of Synthetic Control Estimators

Methodological deep-dive and the estimated cost of Brexit on UK GDP.

A study guide for panel data causal inference.

The Core Problem & Case Study Setup

The Timeline

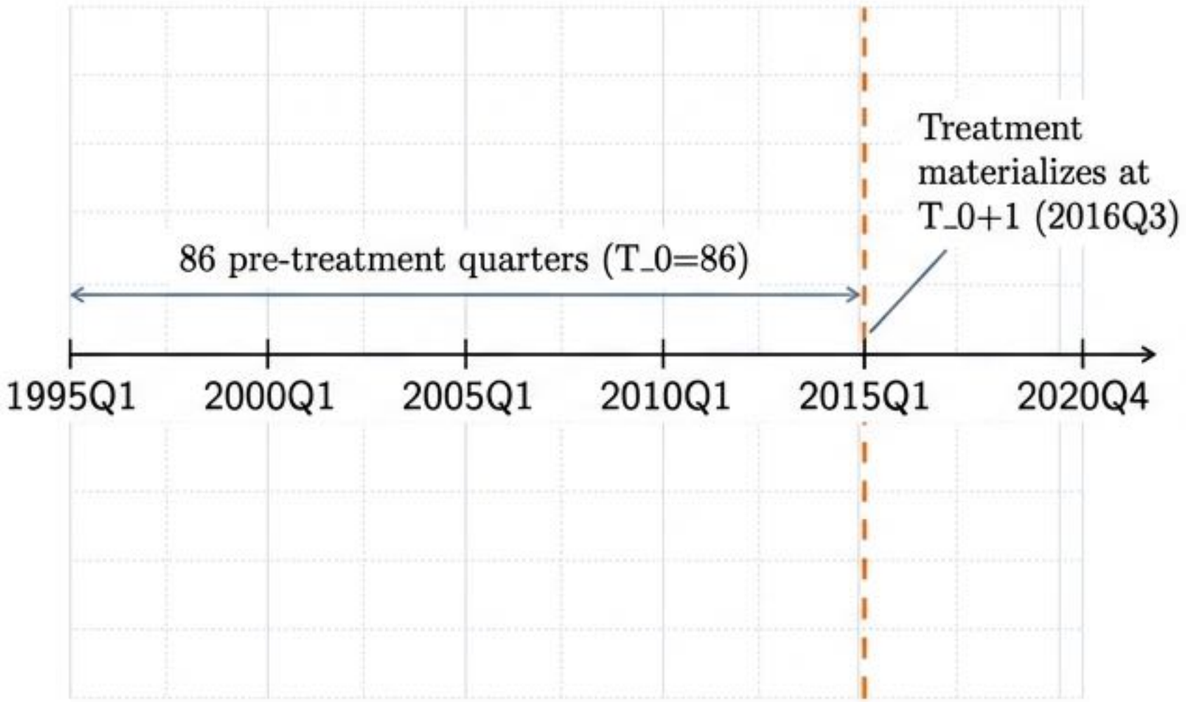

The Donor Pool

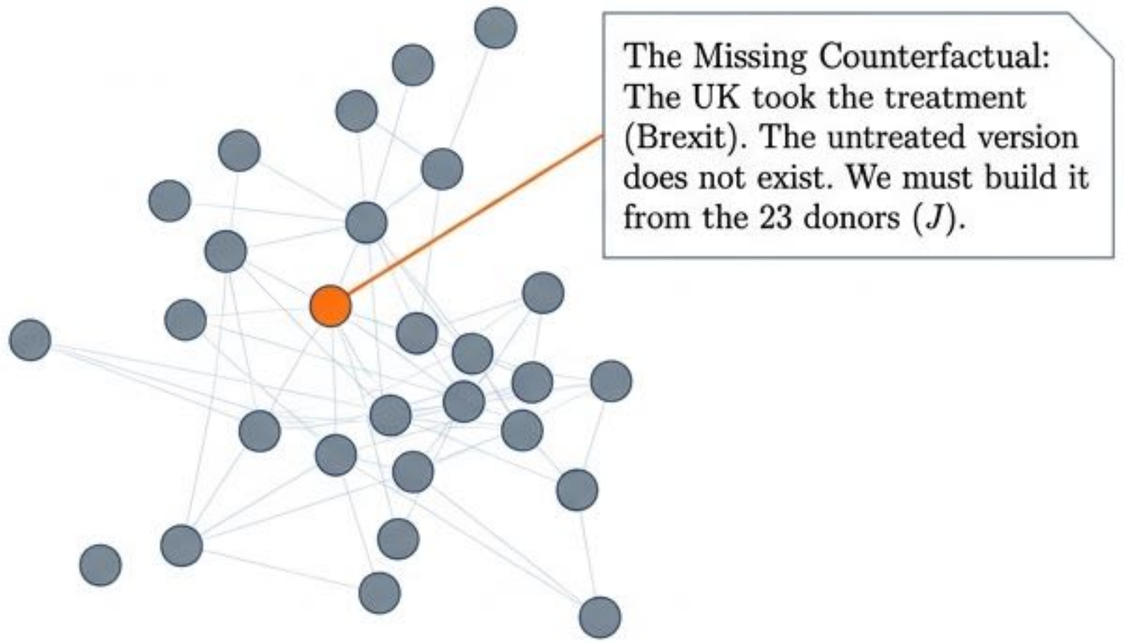

$$\hat{\tau}_t = y_{1,t} - \hat{y}^0_{1,t}$$

Actual UK GDP minus the Estimated Counterfactual.
Outcome is log real GDP indexed to 1995.

The Unified Regression Framework

$$(\hat{\tau}, \hat{\alpha}, \hat{\beta}) = \arg\min \sum_{j=1}^{J+1} \sum_{t=1}^{T_0} (y_{j,t} - \alpha_j - \beta_t - w_{j,t}\tau)^2 \hat{\omega}_j \hat{\lambda}_t$$

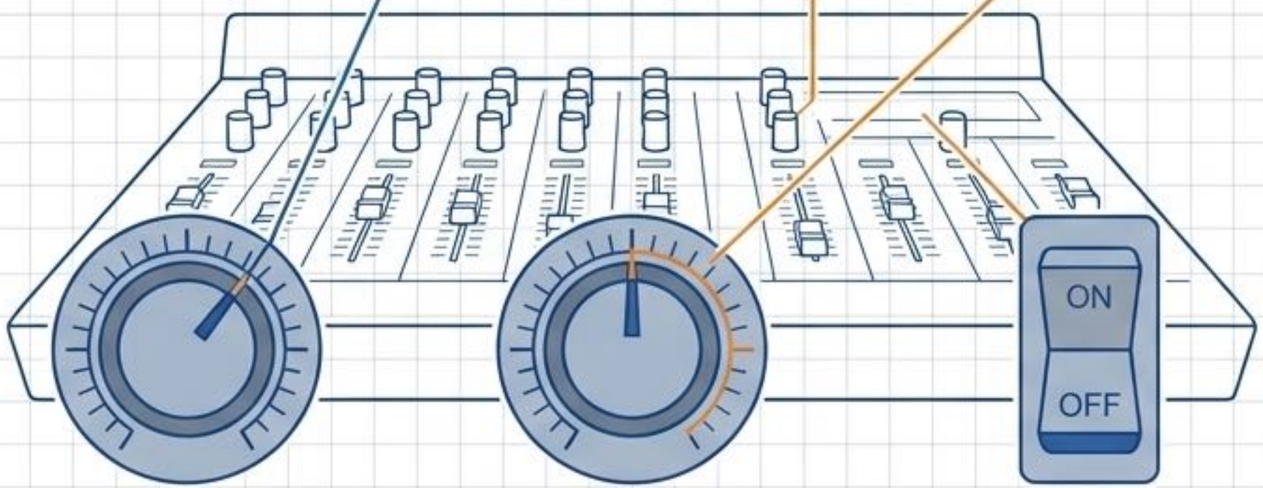

ω Dial: Unit Weights (Which countries matter?)

λ Dial: Time Weights (Which pre-treatment quarters matter?)

α Switch: Unit Fixed Effect (Is there a constant level offset?)

The Evolutionary Ladder: Moving from Difference-in-Differences (DiD) to Synthetic Difference-in-Differences (SDID) is simply a matter of changing the constraints on ω, λ, and α.

Visual Glossary: Constraints & Analogies

The Simplex ($\omega \geq 0, \sum = 1$):

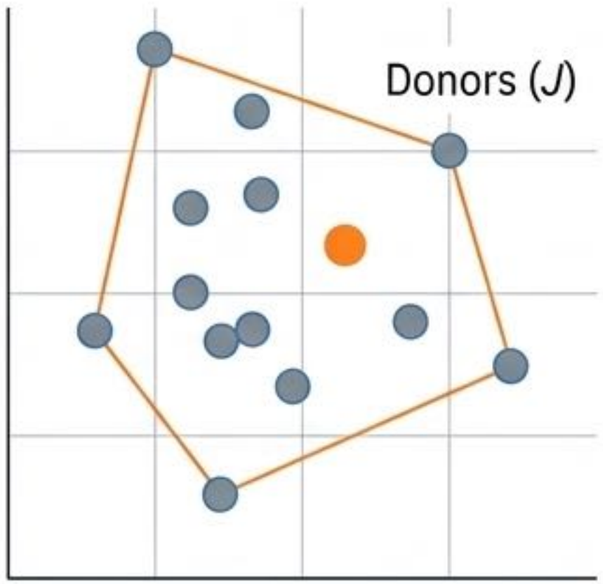

The synthetic blend is trapped inside the rubber band of donors.

No negative weights allowed.

The Mixing Desk:

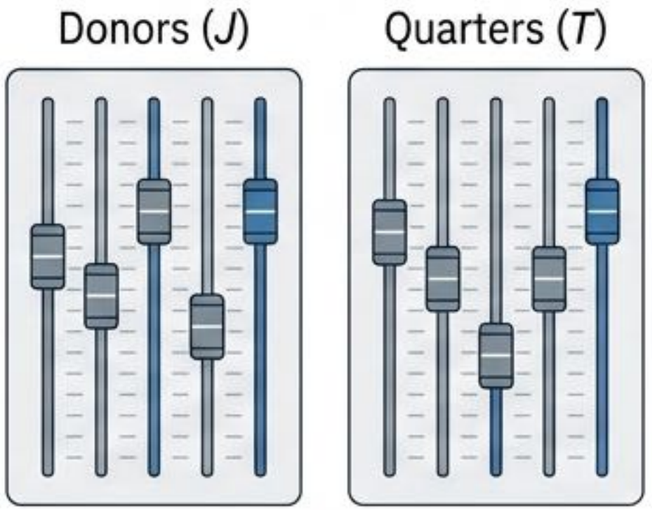

Unit vs. Time Weights:

ω sets how much each donor counts.

λ sets how much each pre-treatment quarter counts.

The Bathroom Scale:

Bias Adjustment (α_j):

A constant offset for parallel but level-mismatched trends. It calibrates the scale without changing the shape of the growth path.

Rungs 1 & 2: DiD vs. Synthetic Control (SC)

$$(\hat{\tau}, \hat{\alpha}, \hat{\beta}) = \arg\min_{\tau} \sum_{t} \sum_{t} (y_{j,t} - \alpha_j - \beta_t - w_{j,t}\tau)^2 \hat{\omega}_j \hat{\lambda}_t$$

Rung 1: DiD

HUD State: omega = flat | λ = flat | α = ON

1/J	1/J	1/J	1/J	1/J	1/J	1/J	1/J
1/J	1/J	1/J	1/J	1/J	1/J	1/J	1/J
1/J	1/J	1/J	1/J	1/J	1/J	1/J	1/J
1/J	1/J	1/J	1/J	1/J	1/J		

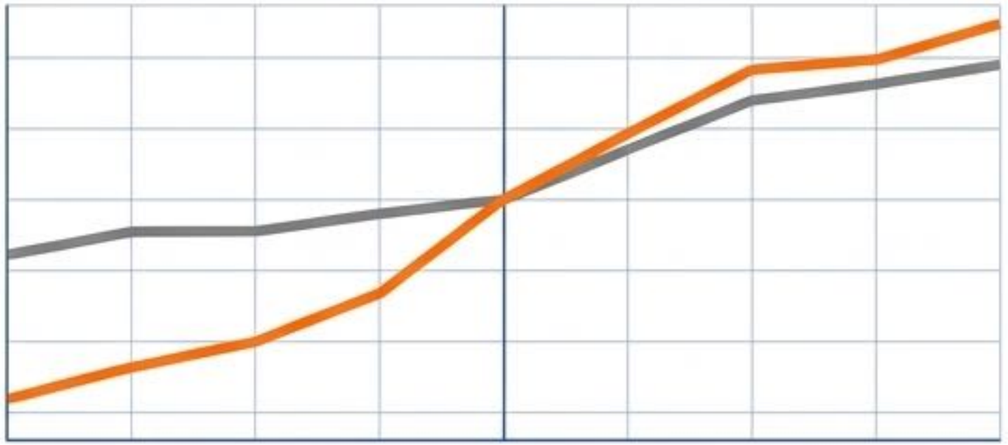

Result: 4.98% gap. Pre-trends fail visibly.

Rung 2: SC

HUD State: omega = optimized | λ = flat | α = OFF

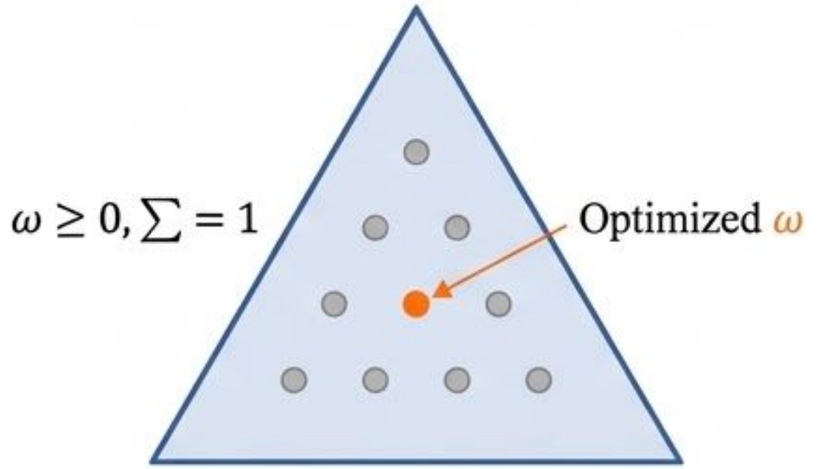

1) Restrict to simplex.
2) Minimize $\sum (y_{UK} - \sum \omega y_{Donors})^2$.
3) Multiply donor outcomes by ω.

Constraint Note: SC switches \alpha OFF, forcing it to match both shape and level.

Result: 3.06% GDP shortfall (2018Q4). Pre-treatment error drops.

Methodological Deep Dive: The Optimization Trap

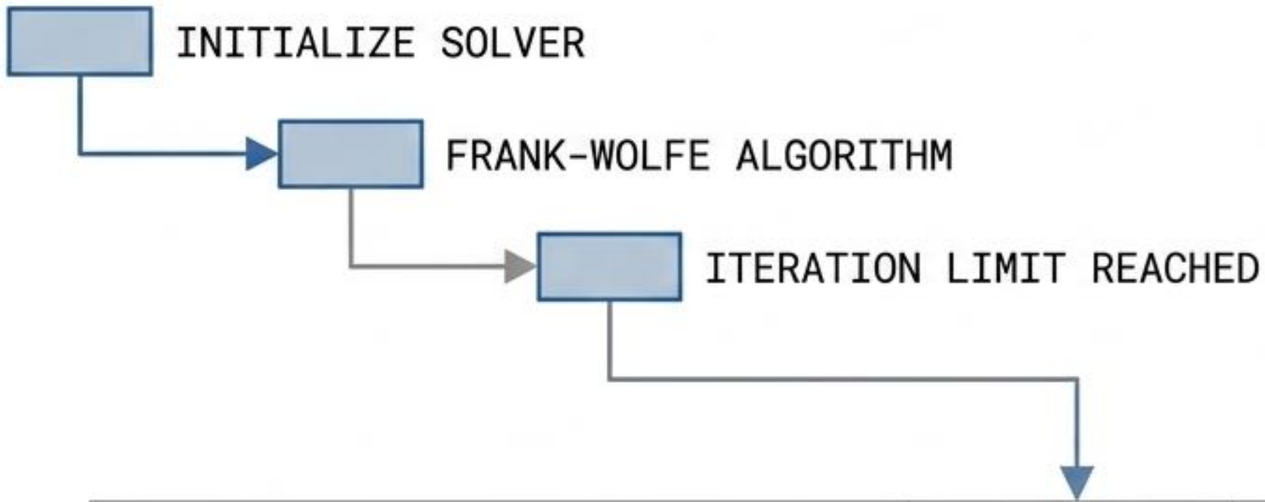

Solver	Iterations	Sum of Squares	Loss Estimate
Package Default (Frank-Wolfe)	10,000	0.00275	3.06%
Exact QP (True Minimum)	-	0.00268	3.04%

The Flat Objective:

With 23 donors and 86 quarters, the objective function is nearly degenerate (condition number ~1,000,000). The default synthdid algorithm halts on an iteration cap, not the true optimum.

Key Insight:

When the synthetic-control objective is flat, individual donor weights are not perfectly identified, but the aggregate treatment effect estimate remains highly stable (a 0.02 percentage point difference).

Rung 3: Demeaned Synthetic Control (DSC)

HUD
ω = optimized (demeaned) | λ = flat | α = ON

LADDER TRACKER

1. DiD
2. SC
3. **DSC**
4. SDID
5. MASC
6. ASCM

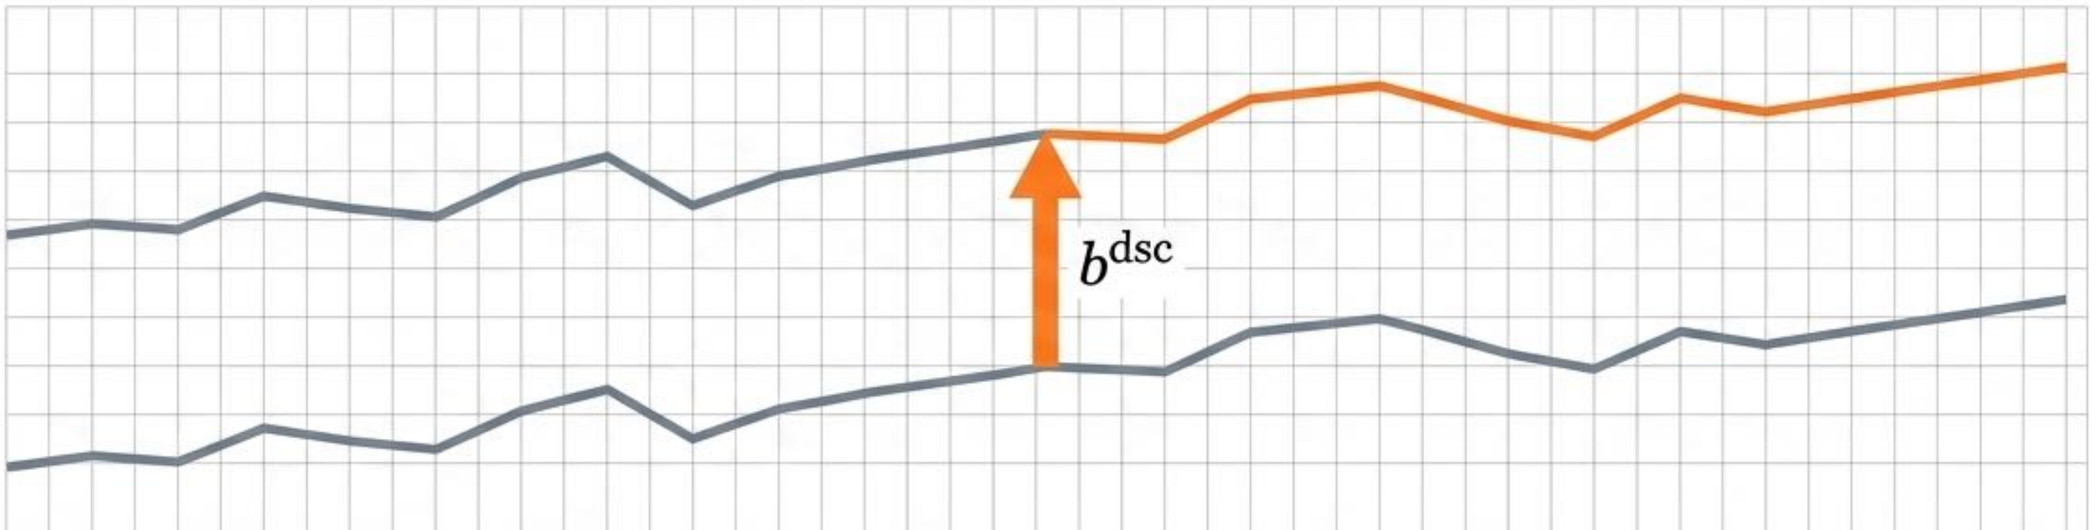

What SC gets wrong:

It lacks an intercept. It will reject a blend that tracks perfectly but sits at a different level.

Implementation:

Step 1: Remove each country's time-series mean.
Step 2: Optimize ω over the simplex on demeaned data.
Step 3: Add the average pre-treatment gap (b^{dsc}) back as a flat offset.

Result:

2.99% GDP shortfall.

(The small shift from 3.06% proves standard SC was already well level-balanced for the UK).

Rung 4: Synthetic Difference-in-Differences (SDID)

HUD	ω = optimized (reused) λ = optimized α = ON (weighted)

LADDER TRACKER

1. DiD
2. SC
3. DSC
4. SDID
5. MASC
6. ASCM

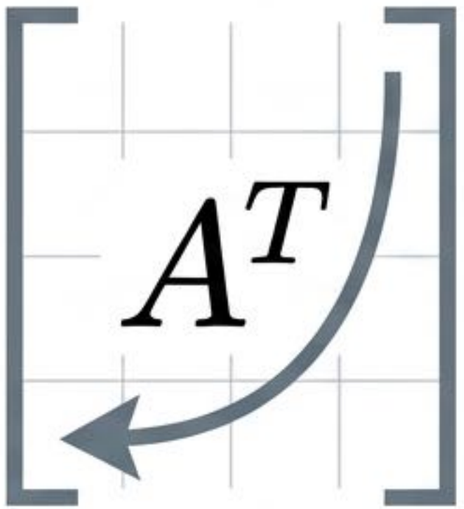

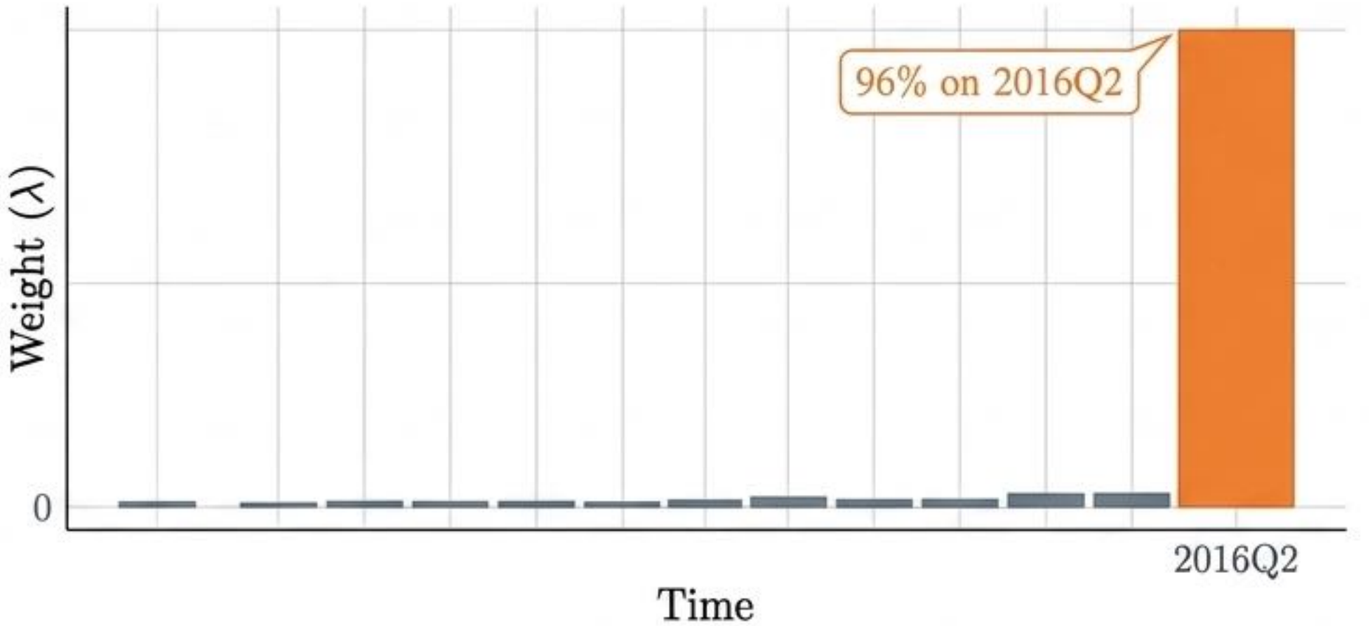

The Structural Insight:

What DSC gets wrong is treating the bias adjustment as a flat average. 1995 gets the same vote as 2016.

SDID lets the data choose λ. The time-weight problem is just the unit-weight problem run on the transpose.

Implementation:

Step 1: Remove each quarter's cross-sectional mean.

Step 2: Find the blend of pre-quarters that best predicts the treatment quarter.

Step 3: Use λ as a weighted average for the bias offset.

Result:

2.76%. Because GDP is a random walk, the best predictor of next quarter is simply this quarter.

96% of the time weight collapses onto the most recent pre-treatment quarter.

The Pivot: Extrapolation vs. Interpolation Bias

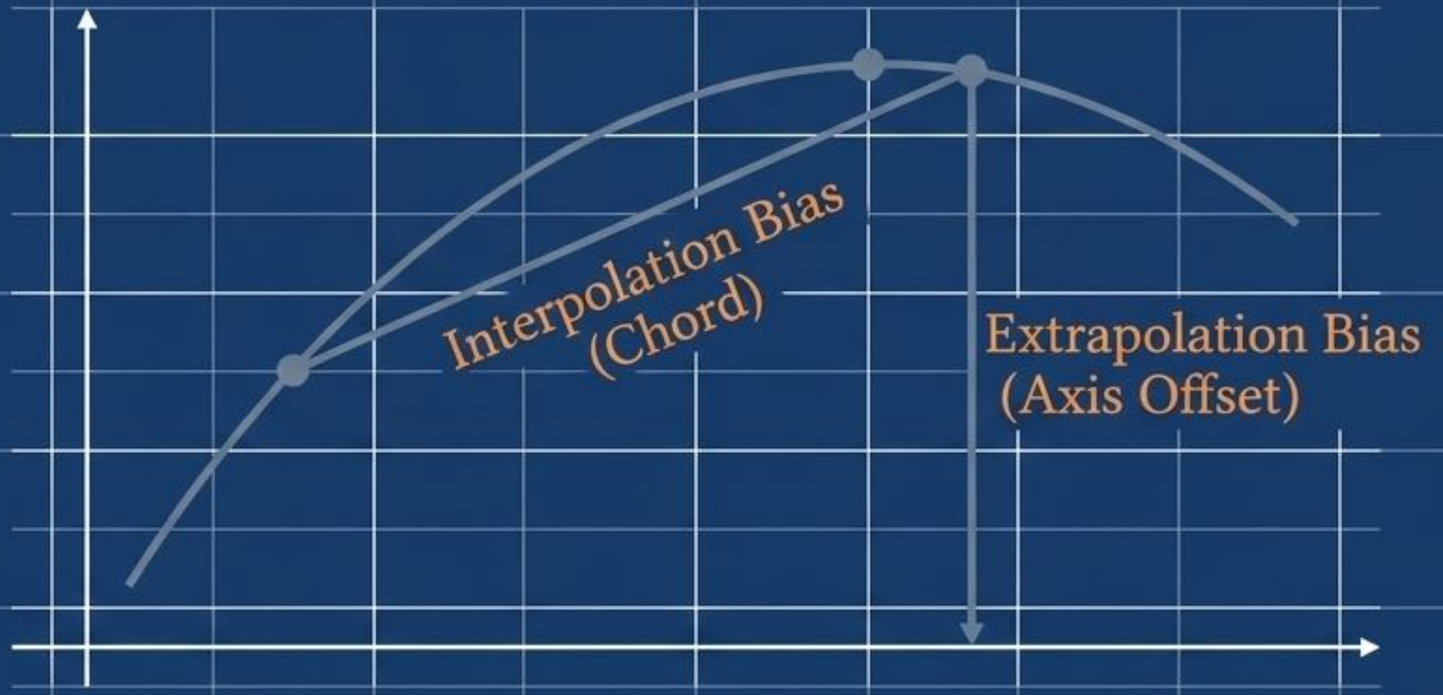

Extrapolation Bias:
Blend characteristics do not match the treated unit.
(Solved by unit weights, ω).

Interpolation Bias:
Characteristics match, but the outcome is a curved function.
Averaging outcomes $\neq$ the outcome of the average.
(Solved by time weights, λ).

The Dual Attack:
SDID is structurally powerful because unit weights attack extrapolation, while time weights attack interpolation simultaneously.

Rung 5: Matching and Synthetic Control (MASC)

HUD	ω = optimized (blended) λ = flat α = OFF

LADDER TRACKER

1. DiD
2. SC
3. DSC
4. SDID
5. MASC
6. ASCM

Nearest Neighbor Matching (kills interpolation)

Synthetic Control (kills extrapolation) (kills extrapolation)

The Trade-off:

If SC kills extrapolation and matching kills interpolation, buy a blend of both:

$$\hat{\omega} = \phi\hat{\omega}^{ma} + (1-\phi)\hat{\omega}^{sc}.$$

Implementation:

Use rolling-origin cross-validation (pre-treatment only) to choose m (neighbors) and ϕ (blend weight).

Result:

The data selects 15.8% matching, 84.2% SC (m=10).

Yields the lowest estimate on the ladder: 2.73%.

Trap Note:

Ensure cross-validation folds run properly from `6` to `T_0-1`. Defaulting to 5 folds forces ϕ=1 (pure matching).

Rung 6: Augmented Synthetic Control (ASCM)

HUD

Master equation summary:
ω = unconstrained (sum=1, sign free)
| λ = flat | α = OFF

Ladder Tracker

1. DiD
2. SC
3. DSC
4. SDID
5. MASC
6. ASCM

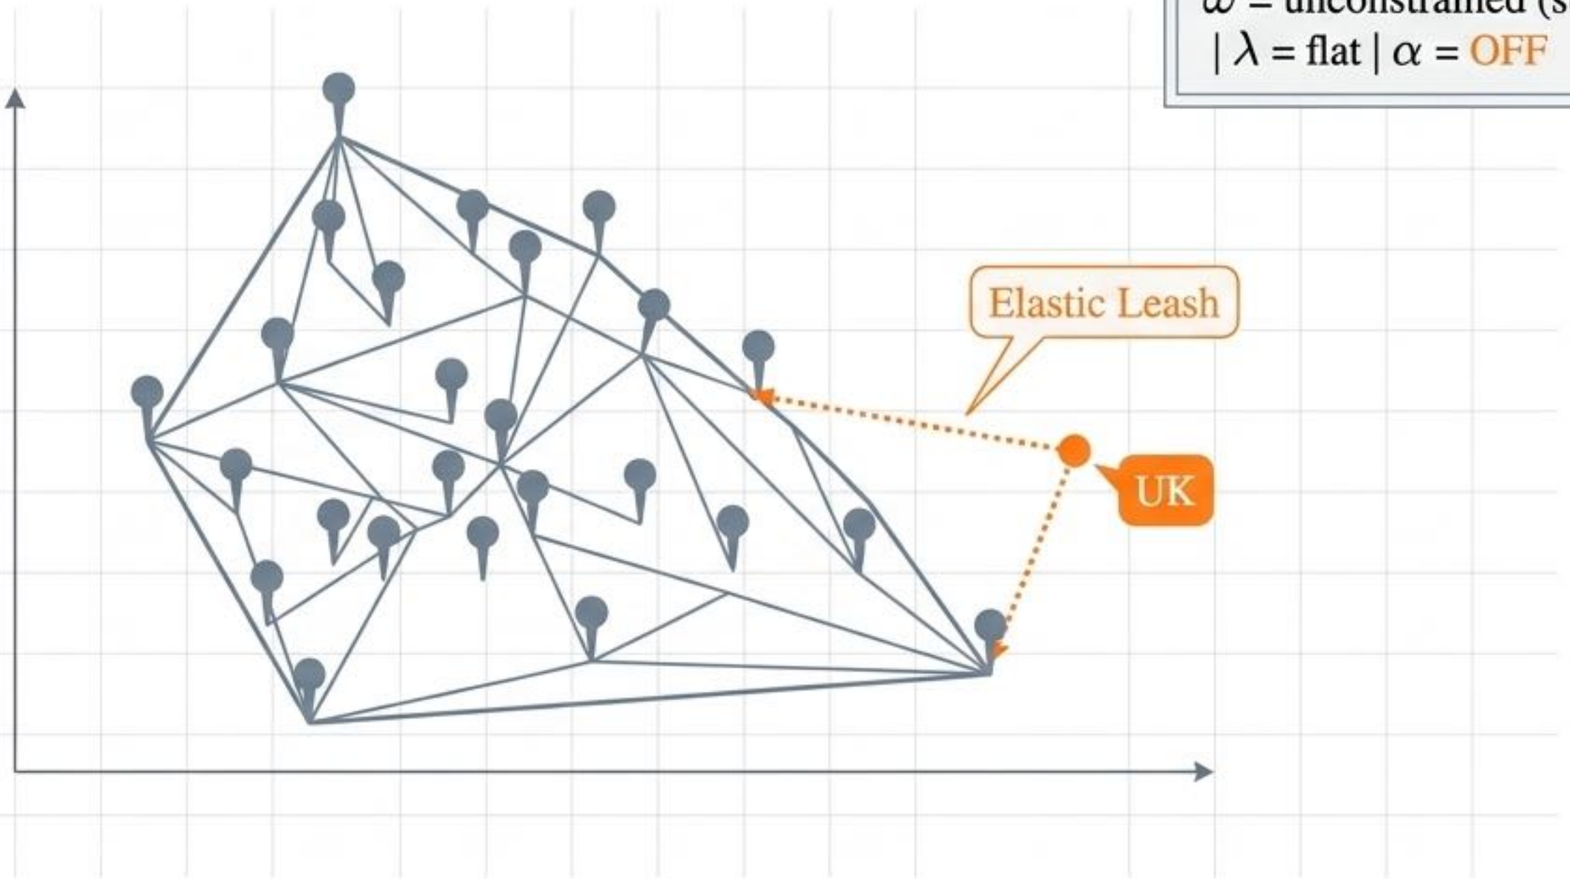

The Ridge Leash:

What if the UK is outside the convex hull? Drop the non-negativity constraint. Allow negative weights, but pay a quadratic penalty (ridge leash) for wandering too far from standard SC weights.

Implementation:

1) Calculate base SC weights.
2) Add ridge-predicted correction for pre-treatment imbalance via cross-validation.

Result:

8 donors receive negative weight (e.g., subtracting Switzerland). Estimate: 3.04%. Because the original SC fit was excellent, the ridge correction barely moves the needle.

Synthesis: The Estimator Matrix

Estimator	Unit Weights (ω)	Time Weights (λ)	Unit Effect (α)	Feasible Set Constraints	End-2018 Estimate
DiD	Fixed at 1/J	Fixed at 1/(T_0-1)	ON	Unconstrained	**4.98%**
SC	Optimized	Flat	OFF	Simplex	**3.06%**
DSC	Optimized (Demeaned)	Flat	ON	Simplex	**2.99%**
SDID	Reused from DSC	Optimized	ON (Weighted)	Simplex	**2.76%**
MASC	Blended (ϕ=0.84)	Flat	OFF	Simplex + Matching	**2.73%**
ASCM	SC + Ridge	Flat	OFF	Sum=1, Sign Free	**3.04%**

Core Academic Insight: Every methodological refinement on the ladder increases the estimated cost of Brexit compared to the published 2019 baseline (2.40%). The choices of ω, λ, and α govern this entirely.

Model Evaluation: In-Sample Placebo Tournaments

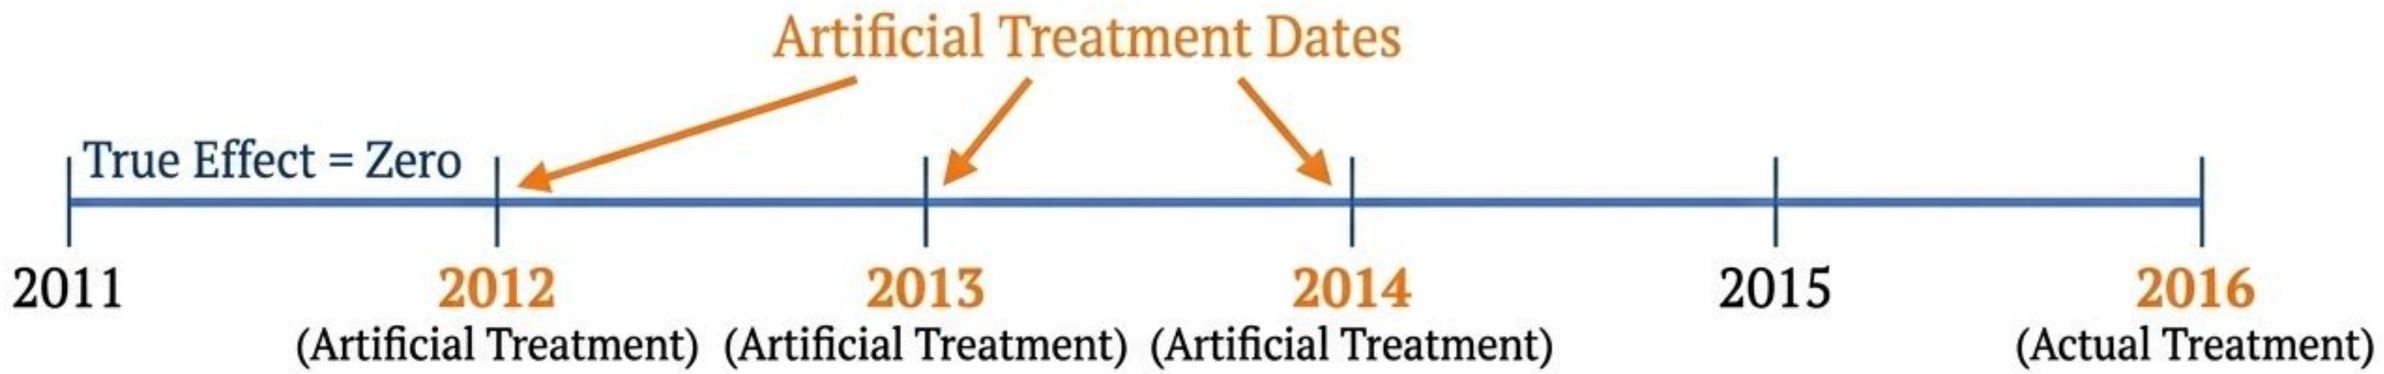

The Design: Pretend treatment happened earlier. Estimate the counterfactual. Since the true effect is zero, any gap found is pure false alarm (RMSE).

Method	RMSE at h=1	RMSE at h=4
SDID (i)	0.0067	0.0134
SDID (ii)	0.0066	0.0134
SDID (iii)	0.0066	0.0134
SC	0.0089	0.0150

The Horizon Trap:
The published paper claimed SDID variant (i) was vastly superior to variants (ii) and (iii). However, (i) was graded on a 1-quarter horizon (h=1), while (ii/iii) were graded on a 4-quarter horizon (h=4).

Corrected Result:
Graded on the same horizon, all SDID variants are indistinguishable. The SDID family as a whole wins the tournament.

Methodological Caveats: Covariates & Treatment Dates

Do Covariates Help?

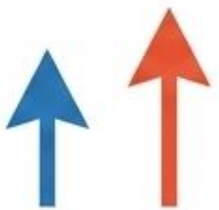

```
SC without covariates: 0.0089 →
SC with covariates: 0.0092

SDID without covariates: 0.0067 →
SDID with covariates: 0.0063
```

No. With 86 pre-treatment outcomes, the timeline already encodes the covariates. Adding predictors just injects estimation noise. The published 2.40% figure is a direct artifact of scoring covariates.

Treatment Date Sensitivity

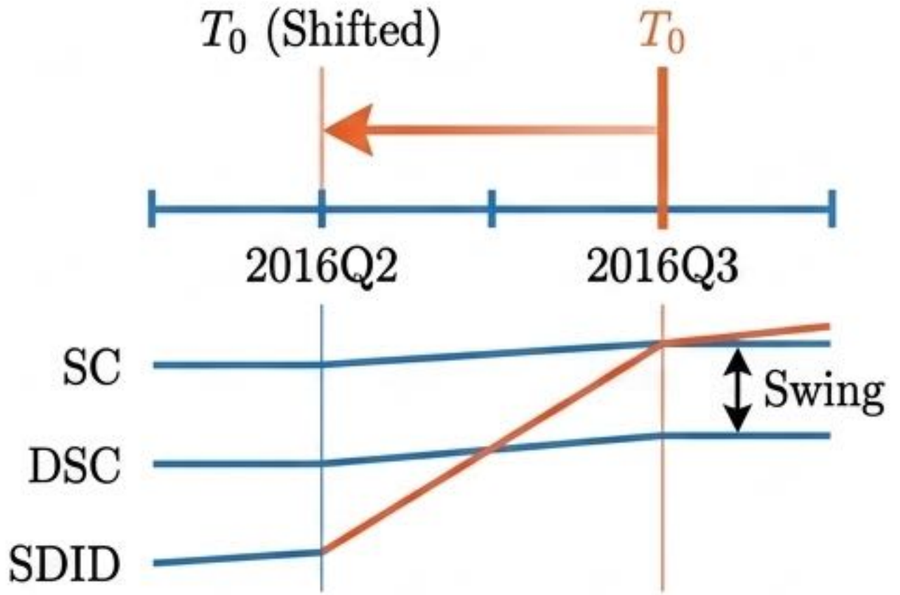

Moving the date barely affects SC/DSC. However, it swings SDID wildly (2.76% → 3.17%). Because SDID time weights are fitted against the treatment quarter, shifting the date changes the entire target of the fit.

Conclusion & Inference Interpretation

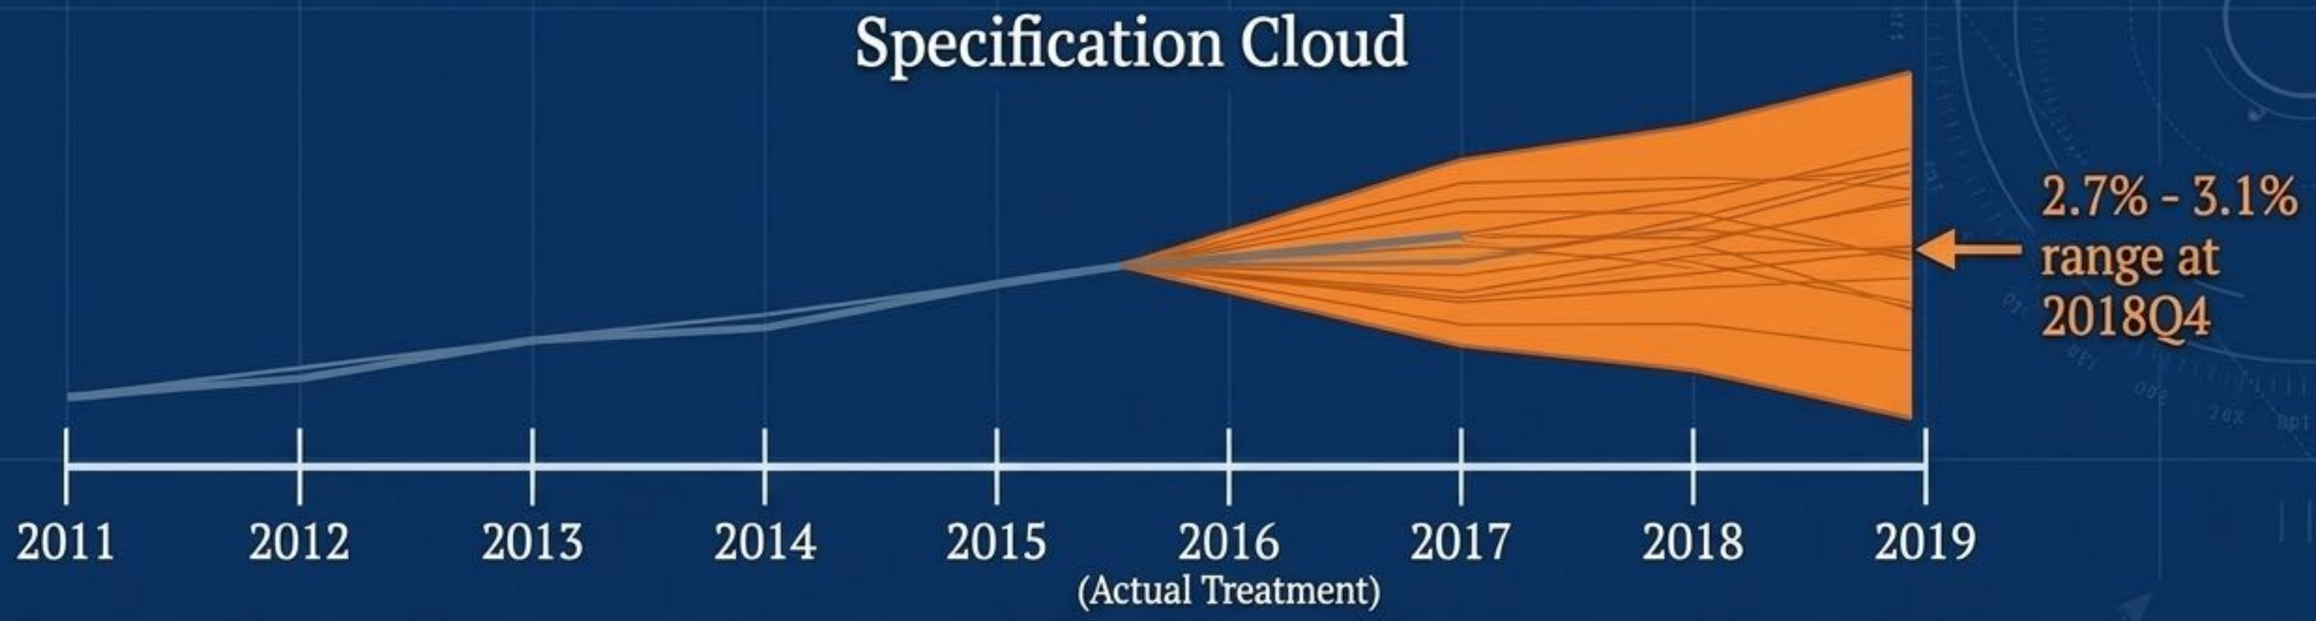

Empirical Conclusion:
Brexit cost the UK between **2.7%** and **3.1%** of GDP by end-2018 (**3.8% - 4.2%** by end-2019).

Inference Reality:
The point estimate is much better determined than the interval. Placebo-in-space p-values cannot reject below **0.042** (1 treated out of 24 donors). Synthdid standard errors yield a wide 95% CI (**0.9%** to **4.6%**).

Methodological Verdict:
Fit the ladder, not a rung. Report the specification cloud. Extrapolation and interpolation bias require separate fixes, but parameter costs must always be weighed.