Case Study:

Answering the Age-Old Question of "Lift" Using Synthetic Control

The opportunity

Measuring true market impact with casual inference

The client, a consumer goods company, sought to evaluate the effectiveness of its $20 million state-wide marketing campaign launched in Florida. While sales increased post-launch, determining causality was elusive, as broader economic trends, competitor actions, and seasonal factors could explain the growth without the campaign's influence.

Attribution was hindered on multiple fronts: lack of a reliable control group due to each state's unique demographics, economies, and market dynamics; inability to run randomized experiments at scale in real-world settings; and reliance on simplistic before-after comparisons that failed to isolate the campaign's effect from confounding variables. The mandate became clear: deploy a rigorous method to quantify true lift, creating a counterfactual scenario that mirrored what sales would have been absent the intervention, enabling data-driven decisions on future investments.

At the time, no comparable "twin" state existed. Sales data showed upward trends, but without isolation, executives risked misallocating budgets or overlooking inefficiencies. The new approach was designed to synthesize controls from available data, turning observational challenges into precise insights and unlocking confident, evidence-based marketing strategies without experimental disruption.

"Our synthetic control method finally answered whether the campaign drove real growth or just rode the wave."

–Luke Matthews, Data Scientist,
Deimos-One

The solution

Building a counterfactual twin for accurate lift

To get there, the team reframed the goal from "track sales increases" to "isolate causal impact through simulation." That meant architecting a synthetic control model to construct a virtual baseline, weighting untreated states based on pre-campaign similarities, and comparing post-intervention divergence to measure true effect.

The solution was to engineer a "Synthetic Florida"—a composite counterfactual built from states without the campaign—so pre-launch trends matched the real Florida closely, allowing any post-launch gap to attribute directly to the marketing effort. We anchored the model in a variant of linear regression as the core algorithm, optimizing weights across historical data like past sales, population, income levels, and economic indicators to find the ideal blend. In parallel, we integrated predictive variables across datasets, ensuring the synthetic twin replicated Florida's behavior dynamically: for example, assigning 40% weight to California, 30% to Texas, 20% to Illinois, and 10% to Georgia based on empirical fit.

On the analytical side, we shifted from aggregate comparisons to granular matching, incorporating time-series alignment for monthly sales in millions—January's $21.7M real vs. $21.4M synthetic, February's $20.4M identical, up to May's near-perfect parity—validating the model's fidelity. State-specific confounders became precision levers: the algorithm adjusted for regional variations, enabling cleaner isolation within temporal pools. To visualize the output, we plotted diverging trends post-June launch, highlighting the gap as lift while iterative testing (covariate balance, placebo checks) pushed robustness back into executive reporting. Finally, we generalized the framework for scalability, routing non-campaign data to inform weights and yields—so the model learned faster, adapted to new variables, and justified its own refinements.

Note: *Prior to June, sales in Real Florida and the constructed Synthetic Florida tracked almost identically, demonstrating a strong baseline for comparison. When the campaign launched, the trajectories diverged: Real Florida's sales rose sharply, while Synthetic Florida's continued along its expected pre-campaign trend. This clear separation provided a measurable estimate of campaign lift. The upward movement of Real Florida against the steady control revealed what sales would have been without intervention, and the widening gap between the two confirmed the direct impact of the marketing campaign.*

The impact

Synthetic control reveals campaign lift: true incrementals not trends

19%

Monthly sales lift: gap analysis showed 19% average increase attributable to the campaign.

100%

Trends matched identically before June, validating the twin's accuracy.

$3.8M

Additional revenue per month post-launch.

The client focused on what matters—quantifying causal effects, faster budget decisions, and optimized future campaigns—and measured every step. By wiring dashboards to real vs. synthetic data flows, codifying pre/post divergences in reports, the team made allocations in days, not quarters.

That rigor cut ambiguity from ad hoc analyses to systematic proofs, routing insights to queues for scaling or refinement. Result: less wasteful spend, higher ROI confidence, lower misattribution risks, and empowered marketing staff.

This improved strategic agility and long-term profitability. Executive buy-in surged, with clearer visuals boosting cross-functional alignment. With robust tools and signals, teams innovated faster and achieved all-time high decision accuracy.

Lessons learned

Build

Model for casual outcomes

Rebuild the analysis end-to-end—linear regression variant as the spine, data integration for scale, covariates and time-series for precision—so every variable feeds weight optimization and divergence signals back into attribution. Standardize historical IDs, lock matching + placeholders, and carry metadata across datasets so the models learn from true baselines, not crude averages.

Optimize

Win on what matters: lift, incrementals, attribution

Optimize for causal effects per dollar, not raw metrics. Use state-smart weighting, covariate-optimized blending, and variant testing to lift accuracy and cut confounders. Tighter alignments, faster iterations, clearer visuals, and robust checks stack together—compounding gains in lift estimates, ROI clarity, and time-to-insight across campaigns.

Cut The Noise

Fit the past

Put Insist on a strong pre-period match (visual and statistical) so post-period divergence reads as impact, not noise.

Prove The Win

Stress test the model

Run placebo and sensitivity checks (donor pool, covariates, window length) to prove the lift is robust, not an artifact of a particular specification..

"Leadership proved the traditional methods wrong. We quantified true lift while improving decisions, accuracy, and confidence, because we treated attribution as a counterfactual, end-to-end system, not a stack of one-off comparisons."
–Jamin Thompson, CEO, Deimos-One

"For the first time, we could see exactly what our campaign delivered. Synthetic control gave us clear proof of lift, letting us redirect budgets with confidence and act faster than ever."
–VP of Marketing, Client