Scalable Neural Synthetic Control with Individual Counterfactuals under Hidden Confounding

Maha Ouali[1,2] **Badih Ghattas**[1] **Emmanuel Flachaire**[1]
Philippe Charpentier [2] **Laurent Bozzi**[2]
[1] Aix Marseille Université [2] EDF

Abstract

As causal inference scales to high-frequency, individual-level observational data, traditional estimators face a dual challenge: computational intractability and bias from hidden confounding. In settings with staggered treatment adoption, prevalent in energy grid management and clinical monitoring, classical Synthetic Control (SC) methods suffer from prohibitive computational overhead when applied to individual units, while neural "black-box" regressors often diverge under non-stationary dynamics. We propose **B-Twin** (Balanced-Twin), a scalable neural framework that bridges representation learning with the structural rigor of synthetic control. B-Twin learns latent representations of individual trajectories that act as proxies for hidden confounders and uses a neural weight regressor to construct synthetic controls as weighted combinations of observed units. This replaces costly per-unit optimization with efficient neural inference while preserving interpretable counterfactual construction. Our experiments on synthetic and real-world datasets demonstrate that B-Twin effectively recovers treatment effects where outcome-regression baselines fail, while scaling efficiently to large populations.

1 Introduction

In the era of AI scaling, causal inference is shifting from small-scale aggregate studies to high-frequency, individual-level observational data. A primary challenge in these settings, such as evaluating Demand-Response (DR) programs in modern energy grids, is the presence of *hidden confounding*, where unobserved factors influence both treatment participation and outcome trajectories.

While classical econometric methods like Synthetic Control (SC) (Abadie and Gardeazabal, 2003; Abadie et al., 2010; Klößner et al., 2018; Abadie and L'hour, 2021) are theoretically robust to latent factors via factor-model structures, they face a **polynomial scaling bottleneck**. These methods typically require solving a constrained quadratic program for each treated unit ($O(N_c^3)$ where N_c is the size of the donor pool). For a population of N units, the resulting $O(N^4)$ total complexity renders individual-level SC intractable for modern datasets with thousands of sensors. Conversely, neural estimators like DragonNet (Shi et al., 2019) or TARNet (Shalit et al., 2017) scale efficiently ($O(N)$) but often function as "black-boxes," lacking the explicit balancing weights and interpretability that make SC-based counterfactuals trustworthy for policy-making. To bridge this gap, we propose B-Twin, a neural framework that identifies the Average Treatment Effect on the Treated (ATT) through the construction of individual counterfactuals. B-Twin reformulates synthetic control as a joint representation learning and balancing task with three key contributions:

- **Latent Representation Scaling:** We learn embeddings Z that serve as proxies for unobserved confounders, filtering the temporal noise that causes traditional SC and SyncTwin (Qian et al., 2021) to diverge.
- **Reduced Matching Complexity:** Unlike classical SC which solves a separate optimization for every treated unit, B-Twin utilizes a shared neural matching block. This reduces the total population complexity to $O(N^2)$, which can be further amortized via mini-batching.
- **Interpretable Neural Balancing:** We preserve the structural form of SC by producing explicit, non-negative matching weights, which ensures the counterfactual remains a transparent, inspectable combination of donor units.

Benchmarks on healthcare (MIMIC-III) and energy (Sowee) datasets demonstrate that B-Twin remains stable where classical methods fail. Our results show that B-Twin provides higher-fidelity counterfactuals than pure neural baselines under non-stationary shifts by anchoring predictions to observed control trajectories.

2 Methodology: Scaling Synthetic Control via B-Twin

We address the estimation of the *Average Treatment Effect on the Treated* (ATT) on time series by constructing individual counterfactual outcomes for treated units.

2.1 Problem Setup and Identification

Let $i \in \{\mathcal{T} \cup \mathcal{C}\}$ denote treated and control units. For a time horizon L, we observe outcomes $Y_i = \{Y_{i,t}\}_{t=1}^{L}$ and treatment $T_i \in \{0, 1\}$. While treatment can be staggered across units, we assume for simplicity that treatment starts at a fixed time t_0 for all individuals $i \in \mathcal{T}$, which allows controlled comparison with baseline methods. In principle, staggered adoption can be handled by aligning trajectories relative to treatment onset, although we do not evaluate this setting empirically. Crucially, we assume treatment is influenced by unobserved confounders $W_i \in \mathbb{R}^K$, violating standard ignorability (Rosenbaum and Rubin, 1983). Let Y_i^1 denote the potential outcome under treatment for unit i and Y_i^0 denote its potential outcome under control. Let (Y^1, Y^0, W, T) denote a generic unit following the same data generating process. To identify the causal effect, we rely on the following assumptions :

Assumption 1. *(SUTVA) The potential outcomes for one individual are unaffected by the treatment of others.*

Assumption 2. *(Positivity) Every individual has a non-zero probability of receiving treatment and control given the latent factor W:*

$$0 < P(T = 1 \mid W) < 1.$$

Assumption 3. *(Latent Ignorability) the potential outcomes are independent of treatment given hidden factors W:*

$$(Y^1, Y^0) \perp\!\!\!\perp T \mid W.$$

For each treated unit $i \in \mathcal{T}$, we estimate the counterfactual outcome under control $\hat{Y}_{i,t}^0$ as a weighted combination of units from the donor pool $\mathcal{C}$:

$$\hat{Y}_{i,t}^0 = \sum_{j \in \mathcal{C}} b_{ij} Y_{j,t}$$

where b_{ij} represents the weight (contribution) of control unit j to the synthetic version of treated unit i. We define the Individual Treatment Effect (ITE) at time t as:

$$ITE_{i,t} = Y_{i,t} - \hat{Y}_{i,t}^0 = Y_{i,t} - \sum_{j \in \mathcal{C}} b_{ij} Y_{j,t}$$

where $Y_{i,t}$ is the observed treated outcome and $\hat{Y}_{i,t}^0$ is the synthetic counterfactual.

To ensure these weights b_{ij} produce unbiased estimates under confounding, they must satisfy a balancing condition (derived in Appendix A) for each unit $i \in \mathcal{T}$:

$$\int \left(\frac{g(w)}{\mathbb{P}^1} - \sum_{j \in \mathcal{C}} b_{ij} \frac{1 - g(w)}{\mathbb{P}^0} \right) \mathbb{E}(Y^0(w)) P(w) dw = 0 \tag{1}$$

where $g(w) = \mathbb{P}(T = 1 | W = w)$ is the propensity score, $\mathbb{P}^1 = \mathbb{P}(T = 1)$, and $\mathbb{P}^0 = \mathbb{P}(T = 0)$. Under hidden confounding, equation (1) reveals that weights must account for differences in the distribution of W between groups. Since we cannot balance on W directly, we propose to use the propensity score $g(W)$ as a proxy (Rosenbaum and Rubin, 1983). This motivates the following proxy balancing condition for each unit $i \in \mathcal{T}$:

$$\mathbb{E}[g(W_i) \mid T = 1] = \mathbb{E}\left[\sum_{j \in \mathcal{C}} b_{ij} g(W_j) \mid T = 0 \right] \tag{2}$$

Equation (2) will guide the selection of weights b_{ij} by encouraging balance in the distribution of propensity scores. In practice, we do not observe W, so we learn a latent representation Z and assume it captures the confounding information needed for $\hat{g}(Z)$ to serve as an approximate balancing score. To overcome the $O(N^4)$ complexity of solving this for every unit via traditional methods, we propose **B-Twin**, a two-stage neural architecture that enables us to learn these weights efficiently (see Figure 1).

2.2 Stage 1: Latent Representation Learning

We map pre-treatment trajectories $Y_{pre} = \{Y_t\}_{t=1}^{t_0}$into a latent space Z using a Variational Autoencoder (VAE) (Kingma et al., 2013) to serve as a proxy for W. From the latent representation Z, we also estimate the propensity score $\hat{g}(Z)$ via a classifier head. The encoder f_ϕ and decoder f_θ and the propensity estimator g_ψ are trained jointly with the following loss function:

$$\mathcal{L}_{rep} = \mathcal{L}_{\mathrm{r}} + \beta D_{\mathrm{KL}} + \gamma \mathcal{L}_{\mathrm{p}} \tag{3}$$

$$\mathcal{L}_{\mathrm{r}} = \|Y_{\mathrm{pre}} - \hat{Y}_{\mathrm{pre}}\|^2,\ \mathcal{L}_{\mathrm{p}} = \mathrm{BCE}(T, \hat{g}(Z))$$

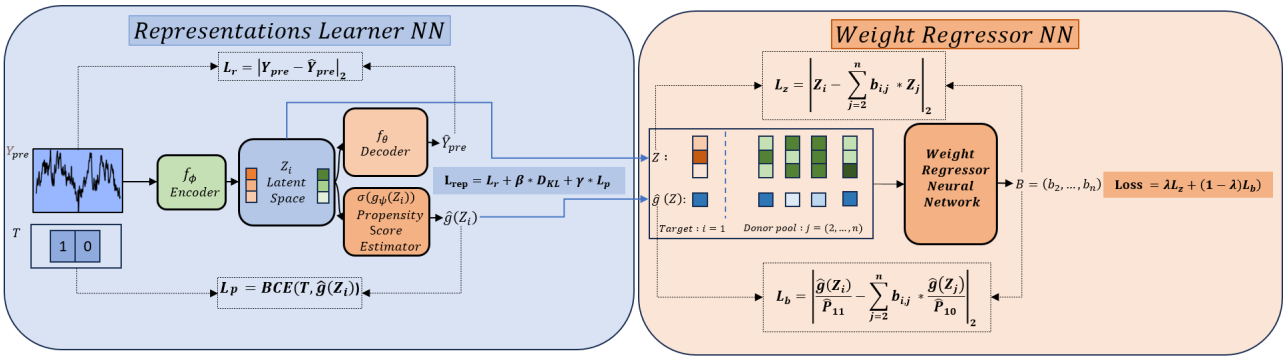

Figure 1: Overview of B-Twin architecture

$$D_{\mathrm{KL}} = \mathrm{KL}\left[q(Z \mid Y_{\mathrm{pre}}) \parallel p(Z)\right]$$

This ensures Z captures both the temporal dynamics of the outcome and the features governing treatment assignment.

2.3 Stage 2: Neural Weights Regression

Given learned representations and estimated propensity scores, a weight regressor module predicts balancing weights B. Unlike classical SC, this module is trained to minimize a **Balancing Loss** derived from Equation (2):

$$\mathcal{L}_{\mathrm{match}} = \lambda \mathcal{L}_z + (1-\lambda)\mathcal{L}_b,$$

where

$$\mathcal{L}_z = \left\| Z_i - \sum_{j \in \mathcal{C}} b_{ij} Z_j \right\|^2, \; \mathcal{L}_b = \left\| \frac{\hat{g}(Z_i)}{\hat{P}_1^1} - \sum_{j \in \mathcal{C}} b_{ij} \frac{\hat{g}(Z_j)}{\hat{P}_0^1} \right\|^2$$

with

$$\hat{P}_1^1 = \frac{1}{|\mathcal{T}|} \sum_{i \in \mathcal{T}} \hat{g}(Z_i), \; \hat{P}_0^1 = \frac{1}{|\mathcal{C}|} \sum_{i \in \mathcal{C}} \hat{g}(Z_i)$$

Scaling Advantage: Unlike traditional SC requiring $O(N_c^3)$ optimization per treated unit, B-Twin uses a shared neural architecture. By processing batches of treated units against the donor pool in a single pass, we reduce the complexity per-unit to $O(N_c)$. This results in a total population complexity of $O(N^2)$, a significant improvement over the polynomial bottleneck of classical solvers.

2.4 The Importance-Weighted ATT Estimator

While B-Twin identifies individual effects $ITE_{i,t}$, our target is the Average Treatment Effect on the Treated (ATT), defined as the empirical mean of these individual estimates: $ATT = \mathbb{E}[Y^1 - Y^0 | T = 1]$. By applying the importance-weighting identity, we reformulate the ATT as an expectation over the marginal distribution of latent confounders:

$$ATT = \mathbb{E}_W[\frac{g(W)}{\mathbb{P}^1}(Y^1 - Y^0)].$$

This identity allows us to leverage the entire population ($\mathcal{T} \cup \mathcal{C}$) for a more stable Monte Carlo approximation (See Appendix C for the full empirical formulation). The final results are provided using this estimator.

3 Experiments

We evaluate B-Twin against neural baselines (SyncTwin(Qian et al., 2021), TARNet(Shalit et al., 2017), CFRNet(Shalit et al., 2017), DragonNet(Shi et al., 2019), BCAUSS(Tesei et al., 2023)) across simulated and real-world datasets focusing on hidden confounding and non-stationarity. Due to $O(N^4)$ complexity, classical Synthetic Control (SC) is limited to group-level aggregates in simulations and used only for individual timing benchmarks in real data.

3.1 Simulated Study: Robustness and Non-Stationarity

We design a DGP where a latent factor W_i influences both treatment T_i and outcomes Y_i. At $t = t_0$, a temporal shift breaks the pre-treatment outcome relationship to challenge extrapolation-based learners. Full DGP specifications are in Appendix B.

Results and Scaling Limits: We fix the population to $N = 10,000$ to test scaling limits. As shown in Fig. 2, B-Twin consistently recovers the ground-truth ATT across all settings. Notably, in cases (c) and (d) (high noise and heterogeneity), regression-based base-

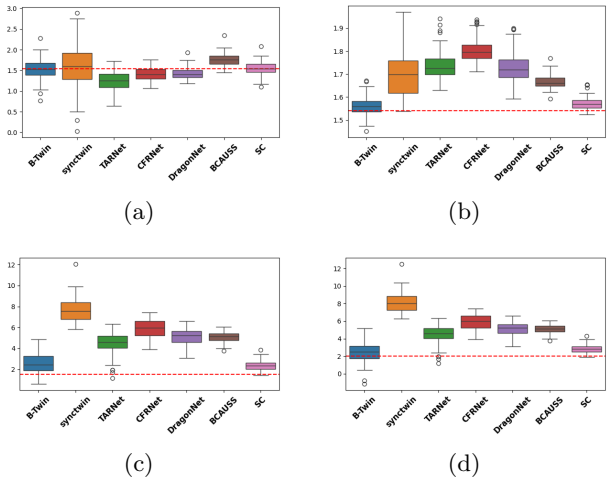

Figure 2: ATT estimation across four settings , 100 Monte Carlo simulations. (a) Randomized experiment; (b) low noise and low confounding; (c) high noise and strong confounding; (d) heterogeneous treatment effects. The red dashed line denotes the ground-truth ATT.

lines (TARNet, CFRNet, DragonNet, BCAUSS) exhibit significant bias. This degradation stems from their reliance on outcome extrapolation: when the DGP shifts, the regressor's learned mapping becomes obsolete.

In contrast, B-Twin's stability demonstrates that representation-based matching is more robust to non-stationary scaling than direct regression. By constructing counterfactuals as weighted combinations of actual control units, B-Twin anchors its predictions to observed data points that share the same latent trajectory, rather than attempting to predict shifted values through a fixed functional head.

3.2 Real-World Scaling: MIMIC-III and Sowee Energy datasets

We evaluate the methods on $N = 12481$ ICU trajectories (MIMIC-III) (Johnson et al., 2016) and $N = 7,924$ electricity load curves (Sowee). Since ground truth is unknown, we generate *placebo datasets* while setting true effect to zero (see Appendix D for details).

Scaling Performance: As shown in Figures 3 and 4, classical SC is computationally impractical for individual-level estimation at this scale (10^4x slower than B-Twin). B-Twin's $O(N^2)$ architecture overcomes the $O(N^4)$ per-unit optimization of classical SC. On the Sowee dataset, classical SC requires 700s per individual unit (projecting to 750+ hours for the population), while B-Twin completes the entire population inference in 6 minutes on a single GPU. This $10^4\times$ speedup enables individual-level synthetic control at

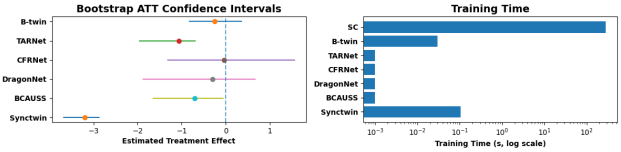

Figure 3: Mimic III Placebo Results: Bootstrap ATT estimation (left panel) and training time (right panel, log-scale).

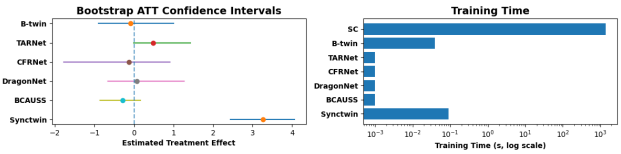

Figure 4: Sowee Placebo Results: Bootstrap ATT estimation (left panel) and training time (right panel, log-scale).

scales where traditional solvers are computationally prohibitive.

Reliability: In placebo datasets, B-Twin achieves the most concentrated confidence intervals around zero. Baselines like CFRNet and DragonNet exhibit wider intervals and spurious shifts, indicating sensitivity to the induced selection bias.

4 Conclusion

We proposed B-Twin, a neural framework designed to scale causal inference to high-frequency, individual-level observational data. By bridging representation learning with the structural rigor of synthetic control, B-Twin enables ATT estimation through the construction of individual counterfactuals in the presence of hidden confounding. Across synthetic and real-world benchmarks, B-Twin demonstrated superior stability under non-stationary shifts by anchoring counterfactuals to observed control trajectories in a balanced latent space. By decoupling representation learning from matching and enabling efficient batch-wise inference, B-Twin reduces complexity from $O(N^4)$ to $O(N^2)$, scaling to large datasets where traditional synthetic control is impractical while maintaining interpretability through explicit balancing weights.

References

Alberto Abadie and Javier Gardeazabal. The economic costs of conflict: A case study of the basque country. *American Economic Review*, 93(1):113–132, 2003.

Alberto Abadie and Jérémy L'hour. A penalized synthetic control estimator for disaggregated data.

Journal of the American Statistical Association, 116 (536):1817–1834, 2021.

Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American statistical Association*, 105(490):493–505, 2010.

Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3 (1):1–9, 2016.

Diederik P Kingma, Max Welling, et al. Auto-encoding variational bayes, 2013.

Stefan Klößner, Ashok Kaul, Gregor Pfeifer, and Manuel Schieler. Comparative politics and the synthetic control method revisited: A note on abadie et al.(2015). *Swiss Journal of Economics and Statistics*, 154:1–11, 2018.

Zhaozhi Qian, Yao Zhang, Ioana Bica, Angela Wood, and Mihaela van der Schaar. Synctwin: Treatment effect estimation with longitudinal outcomes. *Advances in Neural Information Processing Systems*, 34:3178–3190, 2021.

Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.

Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*, pages 3076–3085. PMLR, 2017.

Claudia Shi, David Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. *Advances in Neural Information Processing Systems*, 32, 2019.

Gino Tesei, Stefanos Giampanis, Jingpu Shi, and Beau Norgeot. Learning end-to-end patient representations through self-supervised covariate balancing for causal treatment effect estimation. *Journal of Biomedical Informatics*, 140:104339, 2023.

Supplementary Materials

A Balancing condition

Let W denote a continuous hidden confounder distributed as $W \sim \mathcal{P}_W$, with density $P(w)$. We assume that the propensity score is described by a function g of W and that the treatment assignment follows a Bernoulli $T \mid W = w \sim \text{Bernoulli}(g(w))$. Conditional on $W = w$, the potential outcome under control $Y^0(w) = \{Y_t^0(w)\}_{t=1}^L$ follows $\mathcal{P}_{Y^0(w)}$ and the potential outcome under treatment $Y^1(w) = \{Y_t^1(w)\}_{t=1}^L$ follows $\mathcal{P}_{Y^1(w)}$.

We assume that individual units $(Y_i^0, Y_i^1, W_i, T_i)_{i \in \mathcal{T} \cup \mathcal{C}}$ are i.i.d. random variables from the joint distribution of (Y^0, Y^1, W, T). In particular:

$$Y_i^0 \mid W_i = w \sim Y^0(w), \quad Y_i^1 \mid W_i = w \sim Y^1(w).$$

Inspired by the synthetic control framework, for each treated unit i, we assign a set of constant positive weights $B_i = (b_{ij})_{j \in \mathcal{C}}$, forming a weight matrix $B \in \mathbb{R}^{|\mathcal{T}| \times |\mathcal{C}|}$.

These weights will be used to estimate the counterfactual outcome under control for treated units. To ensure that the synthetic control provides an unbiased estimate of the untreated outcome for each treated unit, we require that for each $i \in \mathcal{T}$:

$$\mathbb{E}[Y_i^0 \mid T_i = 1] = \mathbb{E}\left[\sum_{j \in \mathcal{C}} b_{ij} Y_j^0 \mid T_j = 0\right]. \tag{4}$$

Under Latent Ignorability ($Y^0 \perp T|W$), this holds if and only if the following balancing condition is satisfied:

$$\int \left(\frac{g(w)}{\mathbb{P}(T=1)} - \sum_{j \in \mathcal{C}} b_{ij} \frac{1-g(w)}{\mathbb{P}(T=0)}\right) \mathbb{E}[Y^0(w)]\, P(w)\, dw = 0. \tag{5}$$

Proof. With $g(w) = \mathbb{P}(T = 1 \mid W = w)$, latent ignorability $\big(Y^0 \perp\!\!\!\perp T \big| W\big)$, the law of iterated expectation and Bayes rule give

$$\mathbb{E}(Y_i^0 \mid T = 1) = \int \frac{g(w)}{\mathbb{P}(T=1)}\, \mathbb{E}[Y_i^0 | W_i = w]\, P(w)\, dw. \tag{P1}$$

And since units are exchangeable conditional on W, we write $\mathbb{E}[Y_i^0 \mid W_i = w] = \mathbb{E}[Y^0(w)]$ which yields:

$$\mathbb{E}[Y_i^0 \mid T = 1] = \int \frac{g(w)}{\mathbb{P}(T=1)}\, \mathbb{E}[Y^0(w)]\, P(w)\, dw.$$

For the synthetic control:

$$\mathbb{E}\Big[\sum_{j \in \mathcal{C}} b_{ij} Y_j^0 \mid T = 0\Big] = \int \sum_{j \in \mathcal{C}} b_{ij} \frac{1-g(w)}{\mathbb{P}(T=0)}\, \mathbb{E}[Y^0(w)]\, P(w)\, dw. \tag{P2}$$

Equality of (P1) and (P2) is equivalent to (5). Conversely, if (5) holds, (P1) = (P2) follows immediately. □

B Simulation study : data generation process

In the simulation model, a hidden confounder influences both the treatment assignment and the outcome, with its effect on the outcome increasing gradually over time, in order to break the parallel trends assumption. This design makes the confounding effect difficult to detect in the pre-treatment period, posing a challenge for methods that assume parallel trends or rely on short-term observation. We also introduce non-stationarity between the pre- and post-intervention periods.

Untreated outcomes for each individual i at time t are generated using the following process:

$$Y_{i,t}^0 = 0.5\, q_t W_i^2 + \alpha t\, \mathbf{1}_{\{W_i > 0.5\}} + \epsilon_{i,t}, \quad \forall i,\ t \le t_0,$$

$$Y_{i,t}^0 = 0.5\, q_t W_i^2 + \alpha t\, \mathbf{1}_{\{W_i > 0.5\}} + \beta_1 \sin(kt) W_i + \beta_2 q_t^2 W_i + \epsilon_{i,t}, \quad \forall i,\ t > t_0.$$

where q_t is a time-varying factor following an AR(1) process: $q_t = \rho\, q_{t-1} + \epsilon_t$, and $W_i \sim \mathcal{U}[0,1]$ is an unobserved confounder. The noise terms follow $\epsilon_t \sim \mathcal{N}(0, \sigma)$.

Treatment begins at time $t = t_0$, with treatment assignment governed by the following logistic function:

$$p(T_i = 1 \mid W_i = w) = g(w) = \frac{1}{1 + \exp(-5(w - 0.5))} \tag{6}$$

introducing selection bias between treated and control units. The treated outcome is defined as:

$$Y_{it}^1 = Y_{it}^0 + \tau_i\, \mathbf{1}\{T_i = 1,\, t \ge t_0\}.$$

where τ_i is the individual treatment effect. The four settings in Figure 2 are defined as follows:

Setting	σ	α	τ_i	$g(w)$
a	5	0.05	1.54	0.5
b	1	0	1.54	logistic (6)
c	5	0.05	1.54	logistic (6)
d	5	0.05	$4\log(1 + W_i)$	logistic (6)

Table 1: Simulation settings: outcome noise (σ), confounding strength (α), treatment effect (τ_i), and treatment assignment function $g(w)$.

C Importance weighting estimator

Under the assumption of Latent Ignorability , the Average Treatment Effect on the Treated (ATT) is defined as:

$$\text{ATT} = \mathbb{E}[Y^1 - Y^0 \mid T = 1] \tag{7}$$

Using the law of iterated expectations and Bayes' Rule, we can rewrite this expectation over the latent distribution of confounders W :

$$\text{ATT} = \int \mathbb{E}[Y^1 - Y^0 \mid W = w, T = 1]\mathbb{P}(W = w \mid T = 1)dw \quad = \int \mathbb{E}[Y^1 - Y^0 \mid W = w] \frac{\mathbb{P}(T = 1 \mid W = w)\mathbb{P}(W = w)}{\mathbb{P}(T = 1)} dw \tag{8}$$

Defining the propensity score as $g(w) = \mathbb{P}(T = 1 \mid W = w)$ and $\mathbb{P}^1 = \mathbb{P}(T = 1)$, we obtain the importance-weighting identity:

$$\text{ATT} = \mathbb{E}_W \left[\frac{g(W)}{\mathbb{P}^1} (Y^1(W) - Y^0(W)) \right] \tag{9}$$

To compute Eq. 9, we replace the potential outcomes with observed outcomes Y_i and synthetic counterfactuals $\hat{Y}_i^0$ with $i \in \mathcal{T}$. The synthetic counterfactual is constructed using the neural weights $\hat{b}_{ij}$ which satisfy the balancing condition in the latent space:

$$Y^1(W_i) = Y_i, \quad \hat{Y}^0(W_i) = \sum_{j \in \mathcal{C}} \hat{b}_{ij} Y_j \tag{10}$$

Since this expectation is now taken over the marginal distribution of W, its Monte Carlo approximation must be calculated by averaging over all observed units (both treated and control). For control units $j \in \mathcal{C}$, we similarly construct a treated counterfactual $\hat{Y}_j^1$ as a weighted combination of treated units. The weights b_{ji} represent the contribution of treated unit i in the synthetic control of control unit j and are obtained by applying the same matching procedure with roles of treated and control units reversed. This symmetric augmentation does not change the target estimand, which remains the ATT. It is introduced only to reduce variance and mitigate errors in propensity estimation. The final estimator for the time horizon $\Delta t = L - t_0$ is:

$$\widehat{\text{ATT}} = \frac{1}{n\Delta t} \sum_{t=t_0}^{L} \left[\sum_{i \in \mathcal{T}} \left(\frac{\hat{g}(Z_i)}{\hat{P}_1^1} Y_{i,t} - \sum_{j \in \mathcal{C}} \hat{b}_{ij} \frac{\hat{g}(Z_j)}{\hat{P}_0^1} Y_{j,t} \right) + \sum_{j \in \mathcal{C}} \left(\sum_{i \in \mathcal{T}} \hat{b}_{ji} \frac{\hat{g}(Z_i)}{\hat{P}_1^1} Y_{i,t} - \frac{\hat{g}(Z_j)}{\hat{P}_0^1} Y_{j,t} \right) \right]$$

with

$$\hat{P}_1^1 = \frac{1}{|\mathcal{T}|} \sum_{i \in \mathcal{T}} \hat{g}(Z_i), \; \hat{P}_0^1 = \frac{1}{|\mathcal{C}|} \sum_{i \in \mathcal{C}} \hat{g}(Z_i)$$

D Semi-Synthetic Placebo Design and Scalability

To evaluate B-Twin's ability to recover treatment effects under selection bias and hidden confounding, we construct semi-synthetic placebo datasets using real-world trajectories where the ground-truth ATT is known to be zero. We use two distinct domains to test the framework's versatility:

- MIMIC-III (Healthcare): We analyze Mean Arterial Pressure (MAP) trajectories from $N = 12,481$ ICU patients. We select only patients who did not receive vasopressors to ensure a zero ground-truth effect. To ensure a uniform high-frequency comparison, each trajectory is interpolated and truncated to exactly 72 hourly observations. Artificial treatment is assigned at $t_0 = 60$ hours via a logistic function of patient age (an omitted confounder) to induce hidden confounding.
- Sowee (Energy): This dataset contains daily electricity consumption for $N = 7,924$ households. We select non-participants to ensure a zero ground-truth effect. Each time series consists of 301 daily observations spanning two winter periods. Treatment probabilities are assigned via a logistic function of house surface area, an omitted confounder correlated with baseline usage, to induce hidden confounding.