

GAUSSIAN MIXTURE COUNTERFACTUAL GENERATOR

Anonymous authors

Paper under double-blind review

ABSTRACT

Generating synthetic control arms is a key challenge in clinical research, particularly in crossover trials where placebo data becomes unavailable after patients switch to active treatment. The absence of placebo data complicates estimating long-term efficacy and safety. To solve this, we propose a Gaussian mixture model that generates counterfactual data without needing control data for training. This method handles time-varying, continuous doses and estimates effects between treatment switchers and an extended placebo group, providing valuable insights for treatment effects, evidence generation, and decision-making.

1 INTRODUCTION

A crossover trial is a valuable design in clinical research, particularly for evaluating the long-term efficacy and safety of a new medication. Initially, a randomized clinical trial (RCT) is conducted in which participants are randomly assigned to either the active treatment group or the placebo group. After the initial phase, participants in the placebo group switch to the active treatment arm for additional weeks or months in an open-label crossover trial. This design allows researchers to collect more comprehensive data on the effects of the medication and provides ethical benefits by eventually offering active treatment to all participants (Elbourne et al., 2002; Nolan et al., 2016; Dwan et al., 2019).

The need for such extended trials arises from the desire to understand the sustained impact of a treatment beyond the initial trial period. By including a crossover phase, researchers can observe the effects of the medication on participants who were initially on placebo, thereby enhancing the robustness of the findings. This approach can reveal whether the benefits observed in the initial RCT phase are consistent over a longer period and help identify any delayed adverse effects. In addition, it allows for a more ethical study design, as all participants eventually receive active treatment, which is particularly important in studies involving serious or chronic conditions.

However, there are still challenges associated with crossover trials. An issue is the potential for carryover effects, where the impact of the initial treatment phase influences the outcomes in the crossover phase. This can complicate the interpretation of results. Another problem is the increased complexity and duration of the trial, which can lead to higher costs and logistical challenges. Ensuring participant adherence over a longer period can also be difficult, which can affect the reliability of the data (Mills et al., 2009; Heeson, 2020).

One significant challenge associated with crossover trials is the difficulty in accurately estimating the effects of medications during the extended phase for patients who switched from placebo to active treatment (Zhou et al., 2024). This arises because once the crossover occurs, there is no longer a placebo group to serve as a control for comparison. Without ongoing placebo data, it becomes challenging to distinguish between the effects of the medication and other factors such as natural disease progression or placebo effects that might have persisted. This lack of a concurrent control group can complicate the interpretation of the results of the extended phase, which can lead to less robust conclusions about the long-term efficacy and safety of the medication. Consequently, researchers must carefully consider these limitations when designing and analyzing extended crossover trials.

To address the challenge of estimating the effects between switched patients and a potential extended placebo group, researchers can utilize several strategies. An approach is to use external placebo data from historical trials, which can serve as a benchmark for comparison (Letailleur et al., 2023). This method leverages existing data to create a virtual placebo group, allowing a more accurate estimation of treatment effects during the extended phase (Serrano et al., 2023). However, this

solution may not always be feasible, especially for rare diseases where historical data is scarce or nonexistent. In such cases, researchers might consider adaptive trial designs (Pallmann et al., 2018) that allow modifications based on interim results or Bayesian statistical methods (Berry et al., 2010) that incorporate prior knowledge and expert opinion to enhance the analysis. Despite these potential solutions, challenges remain, including ensuring the comparability of external data and addressing biases that may arise from differences in study populations or methodologies. Therefore, careful planning and robust statistical techniques are essential to mitigate these issues and derive meaningful conclusions from extended crossover trials.

The issues described in the crossover trials highlight the broader types and complexity of the design of synthetic control arms in clinical studies. In modern trial settings, generating synthetic data remains a challenging problem that machine learning models have yet to fully address. This paper proposes a model capable of generating synthetic control data by *training solely on active arm data*. The approach seeks to overcome limitations in existing methods that require both control and active data. This model aims to provide a robust alternative for trial designs lacking a conventional control group using advanced learning techniques.

2 RELATED WORK

2.1 LIMITATION OF SCM IN CREATING A SYNTHETIC CONTROL ARM

Generating synthetic control arms in crossover trials is challenging due to the unavailability of control group data after patients switch to active treatment. In a standard RCT, control group data are crucial for comparing outcomes over time. However, in crossover trials—especially those with an open-label extension phase—patients initially assigned to the control group often switch to the active treatment after a certain period (e.g., after 12 weeks). This switch means that there are no remaining patients in the control group during the extended phase (e.g., up to 24 weeks). The absence of extended control data makes it difficult to assess the long-term efficacy and safety of the treatment, as there is no direct comparison group.

Abadie’s synthetic control method (SCM) attempts to approximate the treatment group’s outcomes by creating a weighted combination of control units that closely resembles the treatment-assigned patient data before the intervention (Abadie & Gardeazabal, 2003; Abadie et al., 2010). This method relies on the availability of control data to compute the necessary weights. Applying Abadie’s method becomes problematic in the context of crossover trials lacking extended control data. For example, in a trial with a 12-week RCT followed by a 12-week open-label extension where all patients receive the active treatment, no 24-week control data is available. Without this data, we cannot calculate the weighted sum of the 24-week control outcomes needed to construct a synthetic control arm for the switched patients. This limitation prevents the direct application of Abadie’s algorithm in such scenarios. In Zhou et al. (2024), SCM and difference-in-differences (Abadie, 2005; Sant’Anna & Zhao, 2020) methods successfully measured treatment effects in crossover trials with the availability of external controls; however, this approach is not feasible in studies where securing external controls is challenging.

2.2 S-LEARNER VS. T-LEARNER

In common types of meta-learner, two distinct concepts are prevalent: the S-learner and the T-learner (Künzel et al., 2019; Okasa, 2022). The term ‘S-learner’ is derived from the notion of a single base learner, while ‘T-learner’ is an abbreviation for two base learners (Shalit et al., 2017). In a straightforward scenario involving a treatment group and a control group within a patient data set, the T-learner estimates two separate base learners: one for the treatment group data and another for the control group data. Subsequently, it computes the difference between these two base learners. Estimating individual treatment effects (ITE) based on a Gaussian mixture model exemplifies a T-learner in the context of counterfactual generation (Ahn & Vashist, 2024). The GMM-based counterfactual generator demonstrated superior performance compared to the synthetic control method on simulated low-density lipoprotein (LDL) data (Qian et al., 2021) and showed robustness even when dealing with heavily biased datasets.

However, the algorithm fundamentally operates as a treatment-controlled method, requiring a learned data distribution model on control units to generate counterfactual outcomes. This reliance

poses limitations in scenarios where control data are unavailable, such as in crossover trials with no extended control group. Additionally, its application as a T-learner is confined to data sets with a small, finite number of groups. As the number of groups increases, preparing distinct base learners, each trained on different group data, becomes necessary. A significant limitation of such a T-learner is its inability to accommodate more complex real-world data, where potential treatments can span a multivariate, time-varying, and continuous domain.

Despite this limitation, the algorithm provides an important conceptual foundation for our work. Specifically, it functions as a T-learner, assuming separate distribution models for treatment and control data. Adapting the model into an S-learner framework allows us to develop a unified model capable of counterfactual inference even for unobserved treatment types. This adaptation allows us to generate counterfactual data without relying on control data for training. Therefore, our proposed method builds upon the foundational idea, extending its applicability to control data-unavailable situations.

3 MODEL

In this section, we outline a latent variable model known as static state analysis (SSA), originally used in the paper by (Ahn & Vashist, 2024), and present how our new ideas are incorporated into this model. Building on the methodologies, we introduce two key problems. The first problem involves a brief overview of how to train the SSA model. The second problem concerns what counterfactual predictions to generate using the trained model.

3.1 STATIC STATE ANALYSIS (SSA)

The static state analysis decomposes a longitudinal dataset $\{\mathbf{x}^{(t)}\}$ into a time-varying observer matrix $\mathbf{W}^{(t)}$ and time-independent state vectors \mathbf{s} with additive noises $\boldsymbol{\eta}^{(t)} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Psi}^{(t)})$:

$$\mathbf{x}^{(t)} = \mathbf{W}^{(t)} \mathbf{s} + \boldsymbol{\eta}^{(t)}. \quad (1)$$

The time-independent state \mathbf{s} is referred to as a static state. In this paper, longitudinal data $\mathbf{x}^{(t)}$ generally concatenate all different types of observed data, including baseline data \mathbf{v} (e.g., demographic and genetic information), time-dependent covariates $\mathbf{l}^{(t)}$, results $\mathbf{y}^{(t)}$, and treatments $\mathbf{a}^{(t)}$. $\mathbf{x}^{(t)}$ and $\mathbf{W}^{(t)}$ can be written as $\mathbf{x}^{(t)} = [\mathbf{a}^{(t)}; \mathbf{l}^{(t)}; \mathbf{y}^{(t)}]$ and $\mathbf{W}^{(t)} = [\mathbf{W}_a^{(t)}; \mathbf{W}_l^{(t)}; \mathbf{W}_y^{(t)}]$ for $t \geq 1$. Equation (1) is decomposed into

$$\mathbf{a}^{(t)} = \mathbf{W}_a^{(t)} \mathbf{s} + \boldsymbol{\eta}_a^{(t)}, \quad (2)$$

$\mathbf{l}^{(t)} = \mathbf{W}_l^{(t)} \mathbf{s} + \boldsymbol{\eta}_l^{(t)}$, and $\mathbf{y}^{(t)} = \mathbf{W}_y^{(t)} \mathbf{s} + \boldsymbol{\eta}_y^{(t)}$ with $\mathbf{x}^{(0)} \equiv \mathbf{v} = \mathbf{W}_v^{(t)} \mathbf{s} + \boldsymbol{\eta}_v^{(t)}$. Equation (2) is our new equation that is directly included as observational data, whereas the original SSA model did not include it. Although it may appear to be a mere addition of a simple formula, it is important to note that this single addition fundamentally alters the nature and scope of the algorithm. Equations (1) and (2) also represent a factor model where the noise model follows a diagonal noise covariance matrix. This paper is not the first to model treatments or causes directly using a factor model. Wang & Blei (2019) and Bica et al. (2020a) successfully removed hidden confounding factors by directly modeling multiple causes as latent variables.

3.2 PROBLEM I: FITTING OF A MODEL TO A DATA SET

Given an observed data set $\mathbb{D} = \{\mathbf{x}_n^{(t)} | t \in \mathbb{T}_n, n = 1, \dots, N\} \subset \mathcal{D}$, our first problem is to perform SSA by decomposing the data set \mathbb{D} into a time-varying matrix $\mathbf{W}^{(t)}$ and a set of static state vectors $\mathbb{S} = \{\mathbf{s}_n | n = 1, \dots, N\} \subset \mathcal{S}$. To learn the static state space \mathcal{S} from the dataset, we employ a Gaussian mixture model (GMM) probabilistic distribution to model the static state as a vector-valued random variable, which is given by

$$p(\mathbf{s}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{s}; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k). \quad (3)$$

There is no more dependence on treatment notation because treatment data are directly included in the observational data. Now, the problem is to estimate a set of parameters $\mathbb{Q} =$

$\{\mathbf{W}, \pi_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, \boldsymbol{\Psi}\}$ that maximizes the expected complete-data log-likelihood by using the expectation and maximization algorithm:

$$\hat{\mathbb{Q}} = \arg \max_{\mathbb{Q}} \sum_{n=1}^N \sum_{t \in \mathbb{T}_n} \langle \ln p(\mathbf{x}_n^{(t)}, \mathbf{s}_n | \mathbb{Q}) \rangle \quad (4)$$

where $\langle \mathbf{s}_n \rangle = \int d\mathbf{s} p(\mathbf{s} | \mathbf{x}_n, \mathbb{Q}) \times \mathbf{s}$ and $\langle \mathbf{s}_n \mathbf{s}_n^T \rangle = \int d\mathbf{s} p(\mathbf{s} | \mathbf{x}_n, \mathbb{Q}) \times \mathbf{s} \mathbf{s}^T$. Then, we can also get the noise vector as $\langle \boldsymbol{\eta}_n^{(t)} \rangle = \mathbf{x}_n^{(t)} - \widehat{\mathbf{W}}^{(t)} \langle \mathbf{s}_n \rangle$.

3.3 PROBLEM II: COUNTERFACTUAL GENERATION WITH NO BACKWARD CAUSATION

Let $\hat{\mathbb{Q}}$ be a maximizer of the log-likelihood function by Eq. (4) for a factual data set $\mathbb{D}_f = \{\mathbf{x}_{n,f}^{(\cdot)} | n = 1, \dots, N\}$. The main problem in our paper is to generate a synthetic counterfactual longitudinal dataset $\mathbb{D}_{cf}^{(\tau+\delta:\cdot)} = \{\mathbf{x}_{n,cf}^{(\tau+\delta:\cdot)} | n = 1, \dots, N\}$ from using $\hat{\mathbb{Q}}$ when we assume that we had assigned alternative treatments $\mathbf{a}_{n,cf}^{(\tau:\tau+\delta)}$ instead of $\mathbf{a}_{n,f}^{(\tau:\tau+\delta)}$ where we define the colon notation as concatenating time-varying data or parameters:

$$\mathbf{x} \equiv \mathbf{x}^{(\cdot)} \equiv \begin{bmatrix} \mathbf{x}^{(0)} \\ \mathbf{x}^{(1)} \\ \vdots \end{bmatrix}, \quad \mathbf{x}^{(\cdot\tau)} \equiv \mathbf{x}^{(0:\tau)} \equiv \begin{bmatrix} \mathbf{x}^{(0)} \\ \vdots \\ \mathbf{x}^{(\tau-1)} \end{bmatrix}, \quad \mathbf{a}^{(\tau:\tau+\delta)} \equiv \begin{bmatrix} \mathbf{a}^{(\tau)} \\ \vdots \\ \mathbf{a}^{(\tau+\delta-1)} \end{bmatrix}. \quad (5)$$

Counterfactual pretreatment data must align as closely as possible with the provided factual pretreatment data (Abadie & Gardeazabal, 2003; Doudchenko et al., 2021). Thus, equivalently, the above problem can be written as generating an entire dataset $\mathbb{D}_{cf} = \{\mathbf{x}_{n,cf}^{(\cdot)} | n = 1, \dots, N\}$ that satisfies $\mathbf{x}_{n,cf}^{(\cdot\tau)} = \mathbf{x}_{n,f}^{(\cdot\tau)}$.

The subscript ‘‘f’’ stands for factual, and ‘‘cf’’ is derived from the first letters of counter-factual. Data obtained through observation is called factual data. On the other hand, data generated through counterfactual thinking without being observed in the real world is referred to as counterfactual data in this paper. Counterfactual thinking involves considering hypothetical scenarios, such as imagining what would happen if a patient who took medication at time τ had not taken it. Before time τ , the patient actually did not take the medication, so the factual data and counterfactual data are the same, which can be expressed as $\mathbf{x}_{n,cf}^{(\cdot\tau)} = \mathbf{x}_{n,f}^{(\cdot\tau)}$. The goal of this paper is to obtain counterfactual data. This can be interpreted in the context of clinical trials as virtual patient data or a synthetic placebo arm. In the next section, we introduce the basic assumptions of counterfactual thinking and propose a new method for counterfactual prediction.

4 METHODS

To solve the second problem, this section details the specific assumptions and formulas for generating counterfactual predictions.

4.1 ASSUMPTIONS

A probabilistic observation refers to the process of observing outcomes that are not deterministic but rather governed by probability distributions. In other words, instead of a single, fixed outcome, there are multiple possible outcomes, each with a certain probability of occurring. This concept is fundamental in fields like statistics, machine learning, and various branches of science and engineering (Gelman et al., 1995). The process of obtaining data through observation can be also described using the following interpretation of the SSA model.

Probabilistic Observation

Consider a patient with factual data $\mathbf{x}_f^{(\tau)}$ has received a course of treatment $\mathbf{a}_f^{(\tau:\tau+\delta)}$. Post-treatment factual data can be obtained by sampling the patient's state $\mathbf{s}_f \sim p(\mathbf{s}|\mathbf{x}_f^{(\tau)}, \mathbf{a}_f^{(\tau:\tau+\delta)})$ at the realization \mathbf{s}_f :

$$\mathbf{x}_f^{(\tau)} = \mathbf{W}^{(\tau)} \mathbf{s}_f + \boldsymbol{\eta}^{(\tau)} \quad (6)$$

where $\boldsymbol{\eta}^{(\tau)}$ is a measurement noise vector.

This statement implies a probabilistic observation process in which future outcomes are probabilistically determined by a series of treatments given at the current time τ . However, this statement only addresses the probabilistic acquisition of factual data and requires separate counterfactual thinking to handle counterfactual data. It is as follows:

Retrospective Counterfactual Thinking

Suppose the patient with factual data \mathbf{x}_f had instead received an alternative treatment $\mathbf{a}_{cf}^{(\tau:\tau+\delta)} \neq \mathbf{a}_f^{(\tau:\tau+\delta)}$. We would then have obtained counterfactual data through the same observational process at the realization \mathbf{s}_{cf} :

$$\mathbf{x}_{cf}^{(\tau)} = \mathbf{W}^{(\tau)} \mathbf{s}_{cf} + \boldsymbol{\eta}^{(\tau)} \quad (7)$$

$$\mathbf{s}_{cf} \sim p(\mathbf{s}|\mathbf{x}_f, \mathbf{a}_{cf}^{(\tau:\tau+\delta)}). \quad (8)$$

Here, $\boldsymbol{\eta}^{(\tau)}$ represents the same noise term as in the factual data observation from Eq. (6), reflecting the assumption that the same observational and measurement processes apply, albeit under the alternative treatment scenario.

Retrospective counterfactual thinking basically refers to the mental process of imagining alternative outcomes to past events, essentially considering “what might have been” if different actions or circumstances had occurred. The current point in time at which counterfactual thinking is attempted is in the future relative to time τ . In other words, τ is a point in the past. Equations (7) and (8) generate counterfactual data for times after τ , assuming that the counterfactual data before τ are exactly the same as the factual data observed before τ .

4.2 GENERATING A COUNTERFACTUAL STATE

Once we find a maximizer $\hat{\mathbf{Q}}$ estimated by solving problem I, we can sample a factual state \mathbf{s}_f from $p(\mathbf{s}|\mathbf{x}_f)$ or obtain its expectation $\mathbf{s}_f \equiv \langle \mathbf{s}_f \rangle$ for a patient with data \mathbf{x}_f . As the first step in counterfactual data generation for the patient, we can generate a counterfactual state by realizing a vector \mathbf{s}_{cf} from Eq. (8)

$$\mathbf{s}_{cf} \sim \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{s}; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \quad (9)$$

where \mathbf{s}_{cf} is conditioned on

$$\begin{bmatrix} \mathbf{0} \\ \mathbf{a}_{cf}^{(\tau:\tau+\delta)} - \mathbf{a}_f^{(\tau:\tau+\delta)} \end{bmatrix} = \begin{bmatrix} \mathbf{W}^{(\tau)} \\ \mathbf{W}_a^{(\tau:\tau+\delta)} \end{bmatrix} (\mathbf{s}_{cf} - \mathbf{s}_f). \quad (10)$$

Equivalently, we can write a counterfactual state random vector in the form of

$$\begin{aligned} \mathbf{s}_{cf} &= \mathbf{s}_f + \begin{bmatrix} \mathbf{W}^{(\tau)} \\ \mathbf{W}_a^{(\tau:\tau+\delta)} \end{bmatrix}^\dagger \begin{bmatrix} \mathbf{0} \\ \mathbf{a}_{cf}^{(\tau:\tau+\delta)} - \mathbf{a}_f^{(\tau:\tau+\delta)} \end{bmatrix} + \ker \left(\begin{bmatrix} \mathbf{W}^{(\tau)} \\ \mathbf{W}_a^{(\tau:\tau+\delta)} \end{bmatrix} \right) \boldsymbol{\xi} \\ &\equiv \mathbf{s}_f + \boldsymbol{\delta}_a + \mathbf{N}\boldsymbol{\xi} \end{aligned} \quad (11)$$

where dagger and $\ker()$ denote the pseudo-inverse and a basis of the kernel of $[\mathbf{W}^{(\tau)}; \mathbf{W}_a^{(\tau:\tau+\delta)}]$, respectively, and $\boldsymbol{\xi}$ is a random vector of the following distribution by denoting the treatment difference on the pseudo-inverse basis as $\boldsymbol{\delta}_a$ and the kernel matrix as \mathbf{N} :

$$\xi \sim \sum_{k=1}^K p_k \mathcal{N}(\xi; \mathbf{m}_k, \mathbf{C}_k) \quad (12)$$

with

$$\mathbf{C}_k = (\mathbf{N}^T \Sigma_k^{-1} \mathbf{N})^{-1} \quad (13)$$

$$\mathbf{m}_k = \mathbf{C}_k \mathbf{N}^T \Sigma_k^{-1} (\boldsymbol{\mu}_k - \mathbf{s}_f - \boldsymbol{\delta}_a) \quad (14)$$

$$p_k = \frac{\pi_k \mathcal{N}(\mathbf{N} \mathbf{m}_k + \mathbf{s}_f + \boldsymbol{\delta}_a | \boldsymbol{\mu}_k, \Sigma_k) |\mathbf{C}_k|^{\frac{1}{2}}}{\sum_{k'=1}^K \pi_{k'} \mathcal{N}(\mathbf{N} \mathbf{m}_{k'} + \mathbf{s}_f + \boldsymbol{\delta}_a | \boldsymbol{\mu}_{k'}, \Sigma_{k'}) |\mathbf{C}_{k'}|^{\frac{1}{2}}}. \quad (15)$$

Notice that $\boldsymbol{\delta}_a$ in Equations (11), (14), and (15) is the new term derived by introducing Equation (2), and makes a crucial contribution to this paper.

4.3 GENERATING COUNTERFACTUAL OUTCOMES

From Equation (6), (7), and (11), we can generate a counterfactual data vector based on the Gaussian mixture: $\mathbf{x}_{cf} = \mathbf{x}_f + \mathbf{W}(\boldsymbol{\delta}_a + \mathbf{N}\boldsymbol{\xi})$ where $\boldsymbol{\xi}$ is a realization of Equation (12). More formally, we can state:

Gaussian Mixture Counterfactual Generator

If a patient with factual data \mathbf{x}_f had received a treatment differing by $\boldsymbol{\delta}_a$, the counterfactual data that we would have obtained is represented as a random variable given by

$$\mathbf{x}_{cf} \sim \sum_{k=1}^K p_k \mathcal{N}(\mathbf{x}; \mathbf{x}_f + \mathbf{W}(\boldsymbol{\delta}_a + \mathbf{N} \mathbf{m}_k), \mathbf{W} \mathbf{N} \mathbf{C}_k \mathbf{N}^T \mathbf{W}^T) \quad (16)$$

This corresponds to the *Gaussian Mixture Counterfactual Generator* (GMCG), defined with the parameters $\hat{\mathbf{Q}}$ from Equation (4).

This formulation reflects a probabilistic approach for generating counterfactual outcomes based on a *single* Gaussian mixture model. Note that the GMCG presented in Eq. (16) is not a statically fixed GMM model but a model where $\boldsymbol{\delta}_a$, \mathbf{N} , \mathbf{C}_k , and \mathbf{m}_k dynamically change according to the given treatment $\mathbf{a}_{cf}^{(\tau:\tau+\delta)}$, intervention time τ , interval δ , etc.

Because $\mathbf{W}^{(\tau:\tau)} \boldsymbol{\delta}_a = \mathbf{W}^{(\tau:\tau)} \mathbf{N} = \mathbf{0}$, the random variable $\mathbf{x}_{cf}^{(\tau:\tau)}$ is always the same as $\mathbf{x}_f^{(\tau:\tau)}$ by Equation (16). However, the difference in outcomes made by a particular vector $\boldsymbol{\xi}$ in the probabilistic model in Equation (12) is a time-varying *individualized treatment effect* (ITE) caused by the difference in treatments $\mathbf{a}_{cf}^{(\tau:\tau+\delta)} - \mathbf{a}_f^{(\tau:\tau+\delta)}$ that we can obtain

$$\text{ITE} \equiv \mathbf{y}_{cf}^{(\tau:\tau)} - \mathbf{y}_f^{(\tau:\tau)} = \mathbf{W}_y^{(\tau:\tau)} (\boldsymbol{\delta}_a + \mathbf{N} \boldsymbol{\xi}). \quad (17)$$

The expected ITE (eITE) we can obtain on average is

$$\mathbb{E}[\text{ITE}] \equiv \mathbb{E}[\mathbf{y}^{(\tau:\tau)} | \mathbf{x}_f, \mathbf{a}_{cf}^{(\tau:\tau+\delta)}] - \mathbf{y}_f^{(\tau:\tau)} = \mathbf{W}_y^{(\tau:\tau)} (\boldsymbol{\delta}_a + \mathbf{N} \sum_{k=1}^K p_k \mathbf{m}_k). \quad (18)$$

By using Equations (12), (14), and (15), we can calculate the *conditional averaged treatment effect* (CATE) from Eq. (17)

$$\begin{aligned} \text{CATE} &\equiv \mathbb{E}[\mathbf{y}^{(\tau+\delta:\tau+\delta)} | \mathbf{x}_f^{(\tau:\tau)}, \mathbf{a}_{cf}^{(\tau:\tau+\delta)}] - \mathbb{E}[\mathbf{y}^{(\tau+\delta:\tau+\delta)} | \mathbf{x}_f^{(\tau:\tau)}, \mathbf{a}_f^{(\tau:\tau+\delta)}] \\ &= \mathbf{W}_y^{(\tau:\tau)} \mathbf{N} \sum_{k=1}^K (p_k \mathbf{m}_k - q_k \mathbf{n}_k) \end{aligned} \quad (19)$$

where

$$\mathbf{n}_k = \mathbf{C}_k \mathbf{N}^T \Sigma_k^{-1} (\boldsymbol{\mu}_k - \mathbf{s}_f) \quad (20)$$

$$q_k = \frac{\pi_k \mathcal{N}(\mathbf{N} \mathbf{n}_k + \mathbf{s}_f | \boldsymbol{\mu}_k, \Sigma_k) |\mathbf{C}_k|^{\frac{1}{2}}}{\sum_{k'=1}^K \pi_{k'} \mathcal{N}(\mathbf{N} \mathbf{n}_{k'} + \mathbf{s}_f | \boldsymbol{\mu}_{k'}, \Sigma_{k'}) |\mathbf{C}_{k'}|^{\frac{1}{2}}}. \quad (21)$$

5 EXPERIMENTS

5.1 ILLUSTRATIVE EXAMPLE OF A CROSSOVER TRIAL

5.1.1 LINEAR EFFICACY

This section demonstrates how the proposed GMCG algorithm by Equation (16) can be applied to crossover trial data to generate synthetic control data for illustrative purposes. The data set shown here is not actual but derived from a simulation model, the details of which can be found in Appendix C. In the left panel of Figure 1, data from two arms are depicted using shades of red and blue. The red-shaded arm represents groups that received the assigned doses over an entire 24-week period, while the blue-shaded arm represents groups that received the doses only during the last 12 weeks. Each arm is further subdivided into groups based on different doses of 1 mg and 2 mg, resulting in a total of four distinct dose groups. Notably, no group remains on placebo for the entire duration, necessitating the creation of a synthetic control arm to compare to the treatment arms to estimate the efficacy of the treatment. The GMCG model was trained on 400 patients' outcome data across the four groups, each consisting of 100 patients, without additional covariates.

The four thick lines of different colors on the left are randomly selected and displayed, and they are further dealt with in the four subplots on the right. The first subplot displays data from the 1mg group during the second 12-week period, with the algorithm-generated synthetic data marked by the red dashed line. Comparing these synthetic data with the observed data and the ground-truth control data, respectively, we observe a perfect match during the first 12 weeks and a close alignment in the second half. The difference between the observed 12-week 1mg data and the red dashed synthetic data after the 12th week can be interpreted as the patient's expected individual treatment effect (eITE) estimated by Equation (18). The critical point is that the synthetic control arm created through the proposed algorithm corresponds individually to the active or crossover arm, forming a synthetic individually controlled arm. Generating and utilizing the synthetic individually controlled arm makes it possible to estimate ITEs or eITEs.

This simple example clearly illustrates a problem that does not have a readily available logical solution despite its apparent simplicity in the reference. For example, look at the blue-colored data. Seemingly, synthetic control methods could create synthetic control data for the data. However, traditional synthetic control methods rely on the availability of control units, which are absent here, to approximate it by taking the weighted sum of them. Thus, this algorithm provides an innovative way to infer these control data by circumventing the need for explicit control units, highlighting its capability to solve problems that traditional algorithms cannot address.




Figure 1: Illustrative crossover trial dataset (linear efficacy) and synthetic control data generation

5.2 SIMULATED DATA

5.2.1 LDL CHOLESTEROL DATA

For quantitative analysis, LDL cholesterol simulated data was utilized. Data for the treatment and control groups were created using the PK/PD model (Faltaos et al., 2006; Yokote et al., 2008; Kim et al., 2011) of statin drugs prescribed to hypercholesterolaemic patients. The experiment was conducted using data with the same characteristics as those generated by SyncTwin (Qian et al., 2021) and SSA-GMM (Ahn & Vashist, 2024).

As explained in Section 2 on the motivation of this paper, the proposed algorithm was developed by converting SSA-GMM into the S-learner framework. While this greatly enhances its versatility by allowing application to various types of treatments, its performance should remain equivalent to the original GMM-based counterfactual generator. However, we achieved better results by finding optimized parameters in the GMCG algorithm. As shown in Table 1 below, the MAE of SSA-GMM reached a level of 0.07, but GMCG was able to reduce the error to a much lower level. When compared to the SyncTwin, SCM (Synthetic Control Method), and CRN (Bica et al., 2020b) models, the error gap becomes even larger.

Table 1: LDL Cholesterol Data Experiment Results

Method	$N = N_0 + N_1 = 200 + 200 = 400$			$N = N_0 + N_1 = 1000 + 200 = 1200$		
	$p_0 = 0.1$	$p_0 = 0.25$	$p_0 = 0.5$	$p_0 = 0.1$	$p_0 = 0.25$	$p_0 = 0.5$
GMCG	0.042	0.038	0.043	0.040	0.039	0.035
SSA-GMM	0.072	0.072	0.076	0.070	0.073	0.069
SyncTwin	0.308	0.150	0.116	0.178	0.106	0.094
SCM	0.341	0.151	0.150	0.231	0.172	0.158
CRN	0.530	0.631	0.343	0.456	0.404	0.371

5.2.2 LDL CHOLESTEROL DATA WITH NO CONTROLS

The reason for using LDL cholesterol data is not just to show superiority over SSA-GMM and the other algorithms. As sketched in the introduction and illustratively shown in Section 5.1, it is intended to be used for more quantitative analysis of crossover trials. However, the composition of the data differs from that in Table 1. In Table 1, $N_0 = 200$ control data and $N_1 = 200$ treatment data were used. In addition, $N_0 = 1000$ control data and $N_1 = 200$ treatment data were also used as training data. But now, we aim to evaluate a much more challenging task: using no control data at all. This is logically unsound because comparing the results of the treatment data with the control data is the golden standard for determining the treatment effect. Without comparison, the effect of the treatment cannot be discussed. Instead, we assumed that we would test the experimental data at various doses. The total duration of the experiment is $T = 30$. Until time $t = 25$, no patient receives the drug. After time $t = 25$, each patient takes a random dose of statin between 0mg and 10mg. This marks a crossover point at time $t = 25$.

After training the model using random doses data, we tested it on the 10mg test data used in Table 1 and obtained the results shown in Table 2. When comparing the results for $N = 1200$, the performance was similar to SCM. Increasing N slightly improved the results. However, the data for $N = 400$ did not train properly. The results in Table 1 were obtained using control data, while the results in Table 2 were obtained by training the model without control data, highlighting the notable performance of GMCG. As explained in the Related Work section, models like SCM and SyncTwin do not function without control data because they create synthetic control by finding linear or nonlinear combinations of control units. Although recurrent network models like CRN can be tested, considering the poor results in Table 1, it is not possible to obtain proper results without training with control data. The results in Tables 1 and 2 are numerical results in terms of MAE, and examples for visual reference are provided in the Appendix D.

GMM-type algorithms use the EM algorithm to find the local maximum of the given likelihood function. The EM algorithm for GMCG training also finds the local maximum of the likelihood function (see Appendix B for the EM algorithm for GMCG). This means that results are influenced

Table 2: LDL Cholesterol Data Experiment Results (No Control Group)

Method	$N = 1200$			$N = 2000$		
	$p_0 = 0.1$	$p_0 = 0.25$	$p_0 = 0.5$	$p_0 = 0.1$	$p_0 = 0.25$	$p_0 = 0.5$
GMCG	0.251	0.199	0.168	0.184	0.150	0.152

by the initial values. In the GMCG algorithm for crossover trials, the choice of initial values affected performance. Applying the EM algorithm from random values may make it difficult to achieve results like those in Table 2. The data structure of crossover trials is fundamentally curved and requires multiple mixture components. The synthetic control data by zero dose is particularly influenced by mixture components derived from low-dose data. Therefore, the GMCG-EM algorithm should be executed using values obtained by first applying a GMM to data created by low doses. The results in Table 2 used two mixture components for dose data between 0mg and 3mg, one for 3mg to 7mg, and one for 7mg to 10mg. Thus, a total of $K = 4$ mixture components were used.

Although the MAE is larger compared to the GMCG results in Table 1, this method has potentially powerful implications. First, to the best of the authors’ knowledge, this is the first attempt to solve a causal inference problem without using control data. In the field of clinical trials, there are single-arm trials (Wang et al., 2024), but they make control groups using historical or external controls. However, the GMCG method structurally learns the dose response from the experimental data alone and extrapolates the prediction for zero dose. Second, clinical trials without any control data present a new form of clinical trials that are more economical, easier to recruit patients, and more ethically suitable for patients in many cases where historical or external controls cannot be sufficiently obtained, thus opening up broader possibilities.

6 CONCLUSION

This paper proposes a counterfactual generator algorithm based on the Gaussian mixture model to create synthetic control arms. It can generate counterfactual data under typical experimental conditions where drug dosages vary continuously over time. We demonstrated this algorithm as a method to generate synthetic placebo arms for clinical trials, such as crossover trials, where a placebo control arm is absent. The methodology of this algorithm has potential applications in various causal inference problems, not only in the healthcare field but also in econometrics and other domains.

REFERENCES

- Alberto Abadie. Semiparametric difference-in-differences estimators. *The review of economic studies*, 72(1):1–19, 2005.
- 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, 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.
- Jong-Hoon Ahn and Akshay Vashist. A linear algebraic framework for counterfactual generation. In *The Twelfth International Conference on Learning Representations*, 2024.
- Scott M Berry, Bradley P Carlin, J Jack Lee, and Peter Muller. *Bayesian adaptive methods for clinical trials*. CRC press, 2010.
- Ioana Bica, Ahmed Alaa, and Mihaela Van Der Schaar. Time series deconfounder: Estimating treatment effects over time in the presence of hidden confounders. In *International conference on machine learning*, pp. 884–895. PMLR, 2020a.
- Ioana Bica, Ahmed M Alaa, James Jordon, and Mihaela van der Schaar. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. *arXiv preprint arXiv:2002.04083*, 2020b.
- Nick Doudchenko, Khashayar Khosravi, Jean Pouget-Abadie, Sebastien Lahaie, Miles Lubin, Vahab Mirrokni, Jann Spiess, et al. Synthetic design: An optimization approach to experimental design with synthetic controls. *Advances in Neural Information Processing Systems*, 34:8691–8701, 2021.
- Kerry Dwan, Tianjing Li, Douglas G Altman, and Diana Elbourne. Consort 2010 statement: extension to randomised crossover trials. *bmj*, 366, 2019.
- Diana R Elbourne, Douglas G Altman, Julian PT Higgins, Francois Curtin, Helen V Worthington, and Andy Vail. Meta-analyses involving cross-over trials: methodological issues. *International journal of epidemiology*, 31(1):140–149, 2002.
- Demiana William Faltaos, Saïk Urien, Valérie Carreau, Marina Chauvenet, Jean Sebastian Hulot, Philippe Giral, Eric Bruckert, and Philippe Lechat. Use of an indirect effect model to describe the ldl cholesterol-lowering effect by statins in hypercholesterolaemic patients. *Fundamental & clinical pharmacology*, 20(3):321–330, 2006.
- Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 1995.
- Changran Geng, Harald Paganetti, and Clemens Grassberger. Prediction of treatment response for combined chemo-and radiation therapy for non-small cell lung cancer patients using a bi-mathematical model. *Scientific reports*, 7(1):13542, 2017.
- Jeff Heaton. Ian goodfellow, yoshua bengio, and aaron courville: Deep learning: The mit press, 2016, 800 pp, isbn: 0262035618. *Genetic programming and evolvable machines*, 19(1):305–307, 2018.
- P Heeson. Crossover trials: what are they and what are their advantages and limitations. *Students 4 Best Evidence*, 2020.
- Richard Arnold Johnson, Dean W Wichern, et al. Applied multivariate statistical analysis. 2002.
- Jimyon Kim, Byung-Jin Ahn, Hong-Seok Chae, Seunghoon Han, Kichan Doh, Jeongeun Choi, Yong K Jun, Yong W Lee, and Dong-Seok Yim. A population pharmacokinetic–pharmacodynamic model for simvastatin that predicts low-density lipoprotein-cholesterol reduction in patients with primary hyperlipidaemia. *Basic & clinical pharmacology & toxicology*, 109(3):156–163, 2011.

- Sören R Künnel, Jasjeet S Sekhon, Peter J Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences*, 116(10):4156–4165, 2019.
- Valentin Letailleur, Isabelle Chaillol, Fanny Cherblanc, Herve Ghesquieres, Frederic Peyrade, Fabrice Jardin, Stephanie Guidez, Fontanet Delices Bijou, Pierre Sesques, Cédric Rossi, et al. Synthetic control arm from clinical trials and real-world data from lisa group for untreated diffuse large b cell lymphoma patients aged over 80 years: A bona fide strategy for innovative clinical trials. *Blood*, 142:72, 2023.
- Bryan Lim. Forecasting treatment responses over time using recurrent marginal structural networks. *Advances in neural information processing systems*, 31, 2018.
- Geoffrey J McLachlan and Thriyambakam Krishnan. *The EM algorithm and extensions*. John Wiley & Sons, 2007.
- Edward J Mills, An-Wen Chan, Ping Wu, Andy Vail, Gordon H Guyatt, and Douglas G Altman. Design, analysis, and presentation of crossover trials. *Trials*, 10:1–6, 2009.
- Sarah Jane Nolan, Ian Hambleton, and Kerry Dwan. The use and reporting of the cross-over study design in clinical trials and systematic reviews: a systematic assessment. *PLoS One*, 11(7): e0159014, 2016.
- Gabriel Okasa. Meta-learners for estimation of causal effects: Finite sample cross-fit performance. *arXiv preprint arXiv:2201.12692*, 2022.
- Philip Pallmann, Alun W Bedding, Babak Choodari-Oskoei, Munyaradzi Dimairo, Laura Flight, Lisa V Hampson, Jane Holmes, Adrian P Mander, Lang’o Odondi, Matthew R Sydes, et al. Adaptive designs in clinical trials: why use them, and how to run and report them. *BMC medicine*, 16:1–15, 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.
- Pedro HC Sant’Anna and Jun Zhao. Doubly robust difference-in-differences estimators. *Journal of econometrics*, 219(1):101–122, 2020.
- César Serrano, Sara Rothschild, Guillermo Villacampa, Michael C Heinrich, Suzanne George, Jean-Yves Blay, Jason K Sicklick, Gary K Schwartz, Sameer Rastogi, Robin L Jones, et al. Rethinking placebos: embracing synthetic control arms in clinical trials for rare tumors. *Nature medicine*, 29(11):2689–2692, 2023.
- Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International conference on machine learning*, pp. 3076–3085. PMLR, 2017.
- Minyan Wang, Huan Ma, Yun Shi, Haojie Ni, Chu Qin, and Conghua Ji. Single-arm clinical trials: design, ethics, principles. *BMJ Supportive & Palliative Care*, 2024.
- Yixin Wang and David M Blei. The blessings of multiple causes. *Journal of the American Statistical Association*, 114(528):1574–1596, 2019.
- Koutaro Yokote, Hideaki Bujo, Hideki Hanaoka, Masaki Shinomiya, Keiji Mikami, Yoh Miyashita, Tetsuo Nishikawa, Tatsuhiko Kodama, Norio Tada, and Yasushi Saito. Multicenter collaborative randomized parallel group comparative study of pitavastatin and atorvastatin in japanese hypercholesterolemic patients: collaborative study on hypercholesterolemia drug intervention and their benefits for atherosclerosis prevention (chiba study). *Atherosclerosis*, 201(2):345–352, 2008.
- Xiner Zhou, Herbert Pang, Christiana Drake, Hans Ulrich Burger, and Jiawen Zhu. Estimating treatment effect in randomized trial after control to treatment crossover using external controls. *Journal of Biopharmaceutical Statistics*, pp. 1–29, 2024.

A NOTATIONS

Numbers and Arrays

a	A scalar (integer or real)
\mathbf{a}	A vector
\mathbf{A}	A matrix
\mathbf{I}_n	Identity matrix with n rows and n columns
\mathbf{I}	Identity matrix with dimensionality implied by context
$\text{diag}(\mathbf{a})$	A square, diagonal matrix with diagonal entries given by \mathbf{a}
\mathbf{a}	A scalar random variable
\mathbf{a}	A vector-valued random variable
\mathbf{A}	A matrix-valued random variable
$[\mathbf{a}; \mathbf{a}] \equiv \begin{bmatrix} \mathbf{a} \\ \mathbf{a} \end{bmatrix}$	Vectors concatenation
$[\mathbf{A}; \mathbf{A}] \equiv \begin{bmatrix} \mathbf{A} \\ \mathbf{A} \end{bmatrix}$	Matrices concatenation

Sets and Graphs

\mathbb{A}	A set
\mathbb{R}	The set of real numbers
$\{0, 1\}$	The set containing 0 and 1
$\{0, 1, \dots, n\}$	The set of all integers between 0 and n

Probability and Information Theory

$P(\mathbf{a})$	A probability distribution over a discrete variable
$p(\mathbf{a})$	A probability distribution over a continuous variable, or over a variable whose type has not been specified
$\mathbf{a} \sim P$	Random variable \mathbf{a} has distribution P
$\mathbb{E}_{\mathbf{x} \sim P}[f(\mathbf{x})]$ or $\mathbb{E}f(\mathbf{x})$	Expectation of $f(\mathbf{x})$ with respect to $P(\mathbf{x})$
$\text{Var}(f(\mathbf{x}))$	Variance of $f(\mathbf{x})$ under $P(\mathbf{x})$
$\text{Cov}(f(\mathbf{x}), g(\mathbf{x}))$	Covariance of $f(\mathbf{x})$ and $g(\mathbf{x})$ under $P(\mathbf{x})$
$\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$	Gaussian distribution over \mathbf{x} with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$

B EXPECTATION-MAXIMIZATION ALGORITHMS FOR GMCG

For the case of having noise $\boldsymbol{\eta}^{(t)} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Psi}^{(t)})$ with a diagonal matrix $\boldsymbol{\Psi}^{(t)}$, Eq. (1) implies a probability distribution over $\mathbf{x}^{(t)}$ -space for a given $\mathbf{s} \in \mathbb{R}^M$ of the form

$$p(\mathbf{x}^{(t)}|\mathbf{s}) = |2\pi\boldsymbol{\Psi}^{(t)}|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2}(\mathbf{x}^{(t)} - \mathbf{W}^{(t)}\mathbf{s})^T (\boldsymbol{\Psi}^{(t)})^{-1} (\mathbf{x}^{(t)} - \mathbf{W}^{(t)}\mathbf{s}) \right\}. \quad (22)$$

With a Gaussian mixture prior over \mathbf{s} defined by Eq. (3), we obtain the marginal distribution of $\mathbf{x}^{(t)}$ in the form

$$p(\mathbf{x}^{(t)}) = \sum_{k=1}^K \pi_k \cdot |2\pi\mathbf{V}_k^{(t)}|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2}(\mathbf{x}^{(t)} - \mathbf{W}^{(t)}\boldsymbol{\mu}_k)^T \mathbf{V}_k^{(t)-1} (\mathbf{x}^{(t)} - \mathbf{W}^{(t)}\boldsymbol{\mu}_k) \right\}, \quad (23)$$

where the model covariance is $\mathbf{V}_k^{(t)} = \boldsymbol{\Psi}^{(t)} + \mathbf{W}^{(t)}\boldsymbol{\Sigma}_k\mathbf{W}^{(t)T}$.

We can also use the colon notation $\mathbf{x} \equiv \mathbf{x}^{(\cdot)}$ and $\mathbf{W} \equiv \mathbf{W}^{(\cdot)}$ defined by Eq. (5). Then, we have the conditional distribution and the marginal distribution in the form of

$$p(\mathbf{x}|\mathbf{s}) = |2\pi\boldsymbol{\Psi}|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2}(\mathbf{x} - \mathbf{W}\mathbf{s})^T \boldsymbol{\Psi}^{-1} (\mathbf{x} - \mathbf{W}\mathbf{s}) \right\} \quad (24)$$

$$p(\mathbf{x}) = \sum_{k=1}^K \pi_k \cdot |2\pi\mathbf{V}_k|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2}(\mathbf{x} - \mathbf{W}\boldsymbol{\mu}_k)^T \mathbf{V}_k^{-1} (\mathbf{x} - \mathbf{W}\boldsymbol{\mu}_k) \right\}. \quad (25)$$

where $\mathbf{V}_k = \boldsymbol{\Psi} + \mathbf{W}\boldsymbol{\Sigma}_k\mathbf{W}^T$. By Bayes' rule, this leads to the posterior distribution of the form

$$\begin{aligned} p(\mathbf{s}|\mathbf{x}) &= \sum_k p(\mathbf{s}|\mathbf{x}, k) p(k|\mathbf{x}) = \sum_k p(k|\mathbf{x}) \cdot p(\mathbf{x}|\mathbf{s}) p(\mathbf{s}|k) / p(\mathbf{x}|k) \\ &= \sum_k p(k|\mathbf{x}) \cdot |2\pi\mathbf{M}_k|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} \left((\mathbf{s} - \boldsymbol{\mu}_k) - \mathbf{M}_k \mathbf{W}^T \boldsymbol{\Psi}^{-1} (\mathbf{x} - \mathbf{W}\boldsymbol{\mu}_k) \right)^T \right. \\ &\quad \left. \times \mathbf{M}_k^{-1} \left((\mathbf{s} - \boldsymbol{\mu}_k) - \mathbf{M}_k \mathbf{W}^T \boldsymbol{\Psi}^{-1} (\mathbf{x} - \mathbf{W}\boldsymbol{\mu}_k) \right) \right\} \end{aligned} \quad (26)$$

where $\mathbf{M}_k = (\boldsymbol{\Sigma}_k^{-1} + \mathbf{W}^T \boldsymbol{\Psi}^{-1} \mathbf{W})^{-1}$.

Now we have the joint distribution of observational data \mathbf{x}_n and the latent variables \mathbf{s}_n in the form of

$$\begin{aligned} p(\mathbf{x}_n, \mathbf{s}_n) &= |2\pi\boldsymbol{\Psi}|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2}(\mathbf{x}_n - \mathbf{W}\mathbf{s}_n)^T \boldsymbol{\Psi}^{-1} (\mathbf{x}_n - \mathbf{W}\mathbf{s}_n) \right\} \times \\ &\quad \sum_{k=1}^K \pi_k |2\pi\boldsymbol{\Sigma}_k|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2}(\mathbf{s}_n - \boldsymbol{\mu}_k)^T \boldsymbol{\Sigma}_k^{-1} (\mathbf{s}_n - \boldsymbol{\mu}_k) \right\}. \end{aligned} \quad (27)$$

By calculating the expectation values

$$\gamma_{nk} \equiv p(k|\mathbf{x}_n) \quad (28)$$

$$\langle \mathbf{s}_n \rangle_k = \int p(\mathbf{s}|\mathbf{x}_n, k) \mathbf{s} |d\mathbf{s}| \quad (29)$$

$$\langle \mathbf{s}_n \rangle = \int p(\mathbf{s}|\mathbf{x}_n) \mathbf{s} |d\mathbf{s}| \quad (30)$$

$$\langle (\mathbf{s}_n - \langle \mathbf{s}_n \rangle) (\mathbf{s}_n - \langle \mathbf{s}_n \rangle)^T \rangle = \sum_k p(k|\mathbf{x}_n) \mathbf{M}_k, \quad (31)$$

we can obtain the expectation of the complete data log-likelihood in the form of

$$\begin{aligned}
\langle \mathcal{L}_C \rangle = & \sum_{n=1}^N \sum_{k=1}^K \gamma_{nk} \{ \ln \pi_k - \frac{1}{2} \ln |\Sigma_k| - \frac{1}{2} \text{tr}(\Sigma_k^{-1} \langle s_n s_n^T \rangle) + \mu_k^T \Sigma_k^{-1} \langle s_n \rangle - \frac{1}{2} \mu_k^T \Sigma_k^{-1} \mu_k \\
& - \frac{1}{2} \ln |\Psi| - \frac{1}{2} x_n^T \Psi^{-1} x_n + x_n^T \Psi^{-1} W \langle s_n \rangle - \frac{1}{2} \text{tr}(W^T \Psi^{-1} W \langle s_n s_n^T \rangle) \} \\
& - \lambda \left(\sum_{k=1}^K \pi_k - 1 \right)
\end{aligned} \tag{32}$$

where the expectation values are given by

E-steps:

$$\gamma_{nk} = \frac{\pi_k \mathcal{N}(x_n | W \mu_k, V_k)}{\sum_{k'} \pi_{k'} \mathcal{N}(x_n | W \mu_{k'}, V_{k'})} \tag{33}$$

$$\langle s_n \rangle_k = \mu_k + M_k W^T \Psi^{-1} (x_n - W \mu_k) \tag{34}$$

$$\langle s_n \rangle = \sum_k \gamma_{nk} \langle s_n \rangle_k \tag{35}$$

$$\langle s_n s_n^T \rangle = \sum_k \gamma_{nk} (M_k + (\langle s_n \rangle_k - \langle s_n \rangle)(\langle s_n \rangle_k - \langle s_n \rangle)^T) + \langle s_n \rangle \langle s_n \rangle^T \tag{36}$$

with $V_k = \Psi + W \Sigma_k W^T$ and $M_k = (\Sigma_k^{-1} + W^T \Psi^{-1} W)^{-1}$.

Equation (32) is maximized by the following M-step formulas

M-steps:

$$\pi_k = \frac{\sum_n \gamma_{nk}}{N} \tag{37}$$

$$\mu_k = \frac{1}{\sum_n \gamma_{nk}} \sum_n \gamma_{nk} \langle s_n \rangle \tag{38}$$

$$\Sigma_k = \frac{1}{\sum_n \gamma_{nk}} \sum_n \gamma_{nk} (\langle s_n s_n^T \rangle - \mu_k \langle s_n \rangle^T - \langle s_n \rangle \mu_k^T + \mu_k \mu_k^T) \tag{39}$$

$$W = \left[\sum_n x_n \langle s_n \rangle^T \right] \left[\sum_n \langle s_n s_n^T \rangle \right]^{-1} \tag{40}$$

$$\Psi = \frac{1}{N} \sum_{n=1}^N x_n x_n^T - 2W \langle s_n \rangle x_n^T + W \langle s_n s_n^T \rangle W^T. \tag{41}$$

C ADDITIONAL MATERIALS FOR CROSSOVER TRIAL EXAMPLE




Figure 2: A simple illustrative example of creating crossover trials data.

Data Generation. The linear efficacy data presented in Figure 1 were created using the method shown in Figure 2. First, as illustrated in the left panel, outcome data for 400 patients were created by adding noise mainly within the range of -0.2 to 0.2. Figure 1 used these data as the ground-truth control. The middle panel shows the assumed efficacy for four groups based on the duration and dosage of the drug used. It was assumed that all patients in each group exhibited the same efficacy. The assumed curve shape was generated as a function of time, specifically using $\log(t + 1)$ for $0 \leq t < 24$. The two blue-shaded curves are horizontal translations of the two red-shaded curves. The final synthetic crossover trial data, shown in the right panel, were created by combining the data from panels I and II for each patient. This way represents the simplest form of data and does not accurately reflect accurate clinical trial data. Clinical data include discontinuous values, missing values, and various complex covariate structures. While covariate data can provide additional helpful information for experimental results, they were excluded from this example.

Nonlinear Efficacy. In Figure 1, the synthetic control data arm could be created with only two different dose groups per arm due to the assumption that the efficacy is proportional to the dose administered. However, in a more general case, where the relationship between dose and efficacy is nonlinear, the generation of synthetic control data requires more dose groups. Figure 3 presents an example that addresses this question, showing four dose groups for each arm: shades of red representing four dose groups with 0.5, 1.0, 1.5, and 2.0 mg over the full 24-week period, and shades of blue representing four other groups with the same doses during only the last 12 weeks.

The nonlinear relationship assumed here is modeled by $(\text{treatment effect by a dose}) = (\text{dose} + 0.2 \times \text{dose}^2) \times (\text{treatment effect by 1mg dose})$. Each group consists of 100 individuals, totaling 800 training data. The four small plots on the right show synthetic control data for four hypothetical patients taking 1.0 mg and 2.0 mg in each arm, with the generated synthetic mean control data closely approximating the ground truth control.




Figure 3: Simulated crossover trial dataset (nonlinear efficacy) and synthetic control data generation.

This problem’s nonlinear effect cannot be adequately captured with only two dose groups per arm. The Appendix provides additional results related to this issue. This situation is similar to needing at least three points, rather than just two, to determine a quadratic function’s position, direction, and curvature. Therefore, more diverse dose levels are necessary in the experimental data to uncover complex or hidden effects in counterfactual predictions. Therefore, more diverse dose levels are likely necessary in the experimental data to uncover more complex or hidden effects in counterfactual predictions.

This section demonstrates how the proposed GMCG algorithm can be applied to crossover trial data and the types of results it can produce. However, it is important to note that this paper does not claim guaranteed success for all types of crossover trials. As an algorithm within the GMM family, the GMCG algorithm’s clear and concise mathematical model highlights its potential to tackle new types of synthetic data generation problems. Discussing this potential alone represents a meaningful achievement.

Non-trivial Control The ground-truth control data in Figure 1 or 2 follows a simple linear trend, which may seem trivial. However, the proposed method does not fit the shape or distribution of the ground-truth control. Figure 3 demonstrates that the ground-truth control data is well predicted even when following a sine wave pattern.

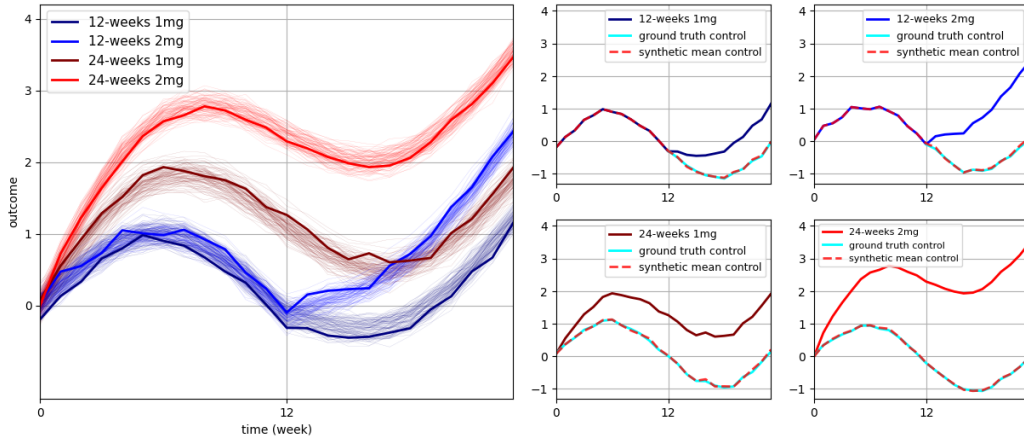


Figure 4: Simulated crossover trial dataset (linear efficacy and sinusoidal ground truth control) and synthetic control data generation.

The problem of estimating the nonlinear effect of dosage shown in Figure 3 cannot be adequately addressed with only two dose groups per arm. Therefore, Figure 3 uses four dose groups per arm. Figure 5 shows the bad result when only two dose groups are used. It demonstrates that the synthetic control for the 12-week treatment group is pulled upward due to the nonlinear effect, and the 24-week treatment group shows unstable predictions with significant fluctuations. Increasing the diversity of dosages in the training data, i.e., using more dose groups per arm, can achieve more stable and accurate counterfactual predictions. For instance, if a hypothetical experiment is conducted where each patient receives slightly different dosages, as shown in Figure 6, the results improve even in more complex nonlinearities in efficacy.




Figure 5: Simulated crossover trial dataset (nonlinear efficacy) and incorrect synthetic control data generation.

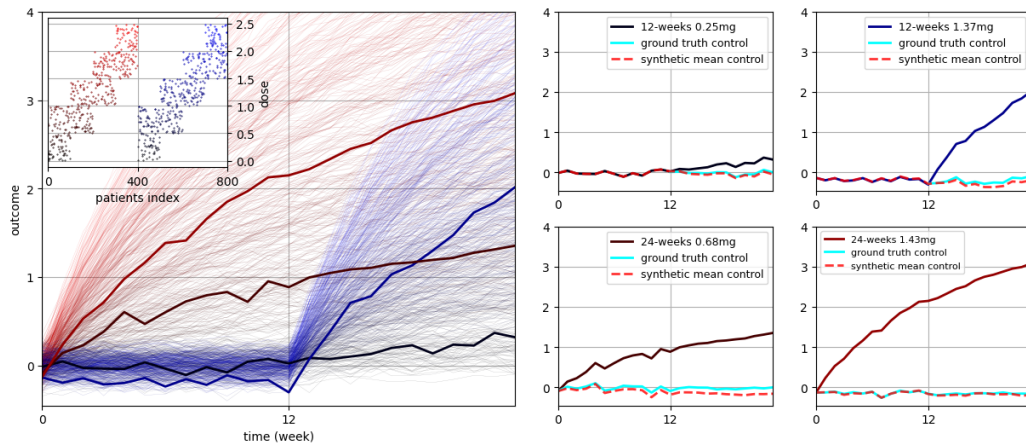


Figure 6: Simulated crossover trial dataset (nonlinear efficacy with diverse doses) and synthetic control data generation.

D ADDITIONAL RESULTS OF LDL DATA EXPERIMENTS

In this experiment, we did not use dimensionality reduction. Thus, \mathbf{W} becomes an identity matrix, making the EM algorithm of section B equivalent to the standard GMM algorithm. Specifically, we used one component for each of the treatment and control groups, which trivially means performing single Gaussian density estimation for each group. At this point, there is still an undetermined noise model: $\Psi = \sigma^2 \mathbf{I}$ (see Appendix B to check how Ψ affects the parameters of Eq. (16)), and we manually substituted various σ values to obtain the counterfactual predictions given by Eq. (16). The results are shown in Fig. 7. The figure shows that very small or very large sigma values resulted in high error values, while values between 0.1 and 1 yielded the lowest error values. Considering that we added noise with a variance of 0.1 when simulating LDL data, it explains why the lowest error value in terms of MAE was obtained around 0.1.




Figure 7: Counterfactual prediction error (MAE) as a function of σ^2 in the noise covariance model $\Psi = \sigma^2 \mathbf{I}$.

Figure 8 plots 15 different patients' examples obtained by applying the GMCG algorithm to the LDL cholesterol data. Specifically, for a total time length of $T = 30$, the simulation was performed at $t = 25$ when statin medication was taken, and the factual data, ground truth counterfactual data, and counterfactual prediction obtained through the GMCG algorithm were plotted on each axis. There are two points of observation here: the pretreatment data before $t = 25$ align exactly with the factual data, and the post-treatment data after $t = 25$ closely approximate the ground-truth counterfactual data.

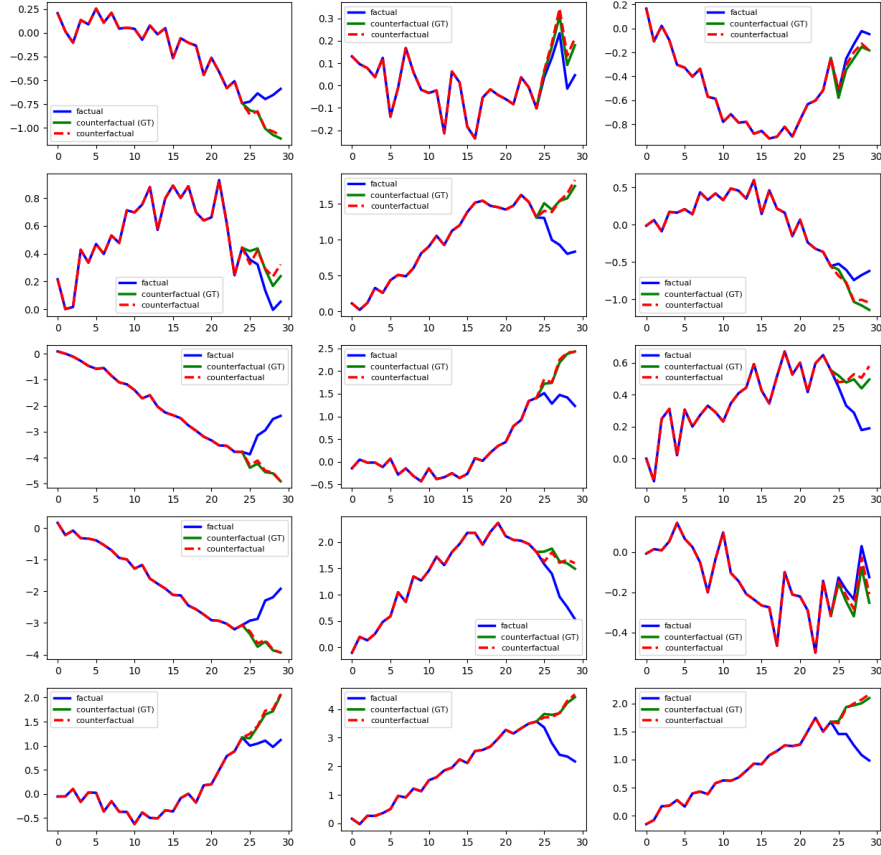


Figure 8: Several plots derived from the results of $p_0 = 0.1$ and $N = 400$ presented in Table 1. Each subplot represents the data of an individual patient. There are two key points to note here: First, the pre-treatment data before $t = 25$ matches exactly between the factual data and counterfactual prediction. Second, the counterfactual prediction after $t = 25$ closely approximates the given ground truth post-treatment data. This serves as further evidence of the excellent results shown in Table 1.

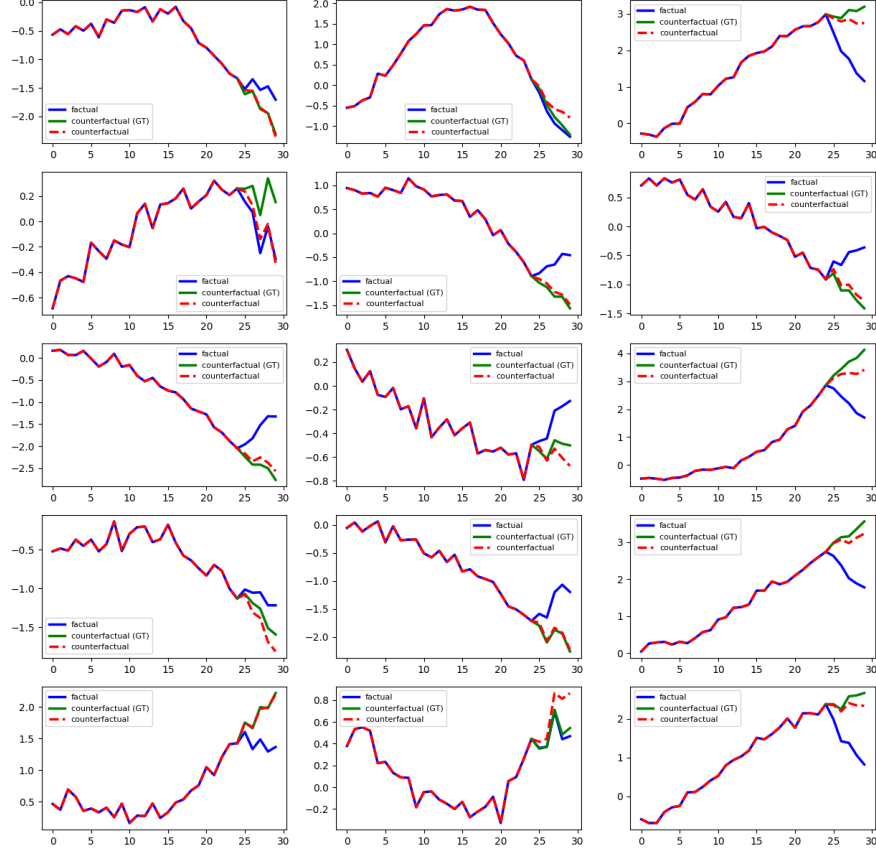


Figure 9: Several plots derived from the results of $p_0 = 0.5$ and $N = 1200$ presented in Table 2. Each subplot represents the data of an individual patient. There are three key points to note here: First, these results were generated by a GMCG model trained without control data, generating control data. Second, the pre-treatment data before $t = 25$ matches exactly between the factual data and counterfactual prediction. Third, the counterfactual prediction after $t = 25$ approximates the given ground truth post-treatment data. This serves as further evidence of the excellent results shown in Table 1.

E SIMULATED TUMOR GROWTH DATA

Although the proposed algorithm is designed to generate synthetic control arms, it can also perform general causal inference, as defined in Section 3.3. This definition aligns with the problem definitions specified in the RMSN and CRN algorithms (Lim, 2018; Bica et al., 2020b). Therefore, in this section, we applied the algorithm to tumor growth data (Geng et al., 2017), which is commonly used for validating general causal inference algorithms. We used data from approximately 10,000 virtual patients and applied it to unbiased and biased data for $T = 10, 20, 30$. The number of mixture components used in the GMCG model is $K = 30$. We compared the state-of-the-art CRN model to the same data and obtained the results. Detailed information is provided in the Appendix.

Table 3 shows the mean absolute error (MAE) obtained by varying δ while keeping τ fixed for data with $T = 10$ and $T = 20$. The error measure was averaged between $t = \tau + 1$ and $t = \tau + \delta$. Surprisingly, the GMM-based GMCG algorithm outperformed the RNN-based CRN regarding MAE. Results for unbiased data are also summarized in the Appendix.

(T, τ)	(10, 5)					(20, 15)				
δ	1	2	3	4	5	1	2	3	4	5
GMCG	0.0436	0.0484	0.0520	0.0505	0.0497	0.0044	0.0051	0.0054	0.0054	0.0056
CRN	0.0587	0.0596	0.0627	0.0633	0.0644	0.0325	0.0283	0.0283	0.0292	0.0305

Table 3: Mean absolute errors (MAE) comparison for biased tumor growth data.

As the duration T increases, there is a tendency for the error values to decrease. However, this cannot be solely attributed to the more extended data being learned. Fundamentally, with longer durations, the tumor is exposed to more therapy sessions, and the tumor size has already reduced. Consequently, the ground truth counterfactual tumor size for other arbitrary therapy alternatives does not differ significantly from factual tumor size. While the GMCG results closely approximated the ground truth counterfactual data or factual data, the CRN’s counterfactual predictions showed some differences from both of them.

How can a GMM-based algorithm yield better results than an RNN-based adversarial network algorithm? Firstly, it is important to recall that GMM can act as an excellent universal approximator as the number of mixture components increases (Heaton, 2018). Additionally, Problem I defined in Section 3.2 is an unsupervised learning problem, and GMM is stably optimized through the EM algorithm (McLachlan & Krishnan, 2007). The most crucial factor lies in the *conditional* probability calculated through Equation (12). Even in regions with sparse training data, the conditional probability calculates a *normalized* probabilistic density distribution conditional on the counterfactual assumption, which creates a robust characteristic against bias. This explains the successful generation of a synthetic control arm in the crossover trials example where training data did not exist.

F ADDITIONAL MATERIALS FOR TUMOR DATA EXAMPLE

F.1 USED HYPERPARAMETERS

To train the CRN for our experiments, we used the default hyperparameters. The exact values used are as follows. These values fall within the optimized hyper-parameters range in Tables 6 and 7 of Appendix J in the CRN paper (Bica et al., 2020b).

The encoder’s used hyperparameters:

```
num_epochs = 100
{ rnn_hidden_units : 24,
  br_size : 12,
  fc_hidden_units : 36,
  learning_rate : 0.01,
  batch_size : 128,
  rnn_keep_prob : 0.9 }
```

The decoder’s used hyperparameters:

```
num_epochs = 100
{ rnn_hidden_units : 12,
  br_size : 18,
  fc_hidden_units : 36,
  learning_rate : 0.001,
  batch_size : 1024,
  rnn_keep_prob : 0.9 }
```

where (br_size) is (balancing representation size) and (rnn_keep_prob) = 1 - (RNN dropout probability).

F.2 ADDITIONAL TABLE

(T, τ)	(10, 5)					(20, 15)				
δ	1	2	3	4	5	1	2	3	4	5
GMCG	0.0548	0.0636	0.0697	0.0694	0.0682	0.0057	0.0061	0.0071	0.0071	0.0073
CRN	0.0831	0.0795	0.0773	0.0788	0.0776	0.0429	0.0390	0.0348	0.0327	0.0296

Table 4: Mean absolute errors (MAE) comparison for unbiased tumor growth data for $T = 10$ and $T = 20$.

(T, τ)	(30, 15)					(30, 25)				
δ	1	2	3	4	5	1	2	3	4	5
GMCG	0.0075	0.0072	0.0077	0.0079	0.0081	0.0015	0.0017	0.0019	0.0019	0.0020
CRN	0.0179	0.0226	0.0262	0.0293	0.0321	0.0128	0.0151	0.0167	0.0192	0.0209

Table 5: Mean absolute errors (MAE) comparison for unbiased tumor growth data for $T = 30$.

G ASSUMPTIONS

G.1 STANDARD ASSUMPTIONS

This paper also adopts Rubin’s potential outcome framework for LDL cholesterol data used in Table 1 and is based on standard assumptions commonly found in many studies, such as consistency, overlap, and unconfoundedness. In addition, important new assumptions are discussed in Section 4.1, where we explore the concept of counterfactual thinking. There are two types of counterfactual thinking: retrospective counterfactual thinking, which examines the effects of altering past treatments at the current time, and prospective counterfactual thinking, which forecasts the outcomes of choosing among various possible future treatments at the present moment.

Assumption 1: Consistency. If treatment $\mathbf{a}^{(:t)} = \mathbf{a}_f^{(:t)}$ was administered to a given patient, then the potential outcome for treatment $\mathbf{a}_f^{(:t)}$ is the same as the observed factual outcome: $\mathbf{y}_f^{(t+1)} = \mathbf{y}^{(t+1)}[\mathbf{a}_f^{(:t)}]$.

Assumption 2: Positivity. If $p(\mathbf{x}^{(:t)} = \mathbf{x}^{(:t)}) \neq 0$, then $p(\mathbf{a}^{(t+1)} = \mathbf{a}^{(t+1)} | \mathbf{x}^{(:t)} = \mathbf{x}^{(:t)}) > 0$ for all $\mathbf{a}^{(t)}$.

Assumption 3: Sequential Strong Ignorability. $\mathbf{y}^{(t+1)}[\mathbf{a}^{(t)}] \perp\!\!\!\perp \mathbf{a}^{(t)} | \mathbf{x}^{(:t)} \quad \forall \mathbf{a}^{(t)}, \forall t$.

G.2 ALTERNATIVE ASSUMPTIONS FOR POSITIVITY VIOLATION

When the positivity assumption is violated, as with the data used in Table 2, an alternative assumption is needed. Generally, it can be assumed that the conditional expectation of the outcome variable is a sufficiently smooth function with respect to the treatment variable. This continuity allows for the estimation of outcomes for treatment values that are not directly observed, i.e.,

Assumption 4: Strong Continuity. $\mathbb{E}[\mathbf{y} | \mathbf{x}_f = \mathbf{x}_f, \mathbf{a}_{cf} = \mathbf{a}_{cf}]$ is continuous and differentiable with respect to treatment to an appropriate order.

Alternatively, it can be assumed that the relationship between the outcome \mathbf{y} , the treatment \mathbf{a}_{cf} , and the covariates \mathbf{x}_f follows a specific functional form (e.g., linear or quadratic regression model). This assumption enables extrapolation to estimate causal effects:

Assumption 5: Functional Form of Outcome Model. $\mathbb{E}[\mathbf{y} | \mathbf{x}_f = \mathbf{x}_f, \mathbf{a}_{cf} = \mathbf{a}_{cf}] = f(\mathbf{x}_f, \mathbf{a}_{cf}; \theta)$ where f is a known functional form, and θ represents parameters to be estimated.

Specifically, the results in Figure 1 and the various examples in the Appendix were obtained based on this assumption.

H PROOF OF GAUSSIAN MIXTURE COUNTERFACTUAL GENERATOR

Lemma 1. *Let the random vector ξ follow a mixture of multivariate normal distributions with the k -th mean \mathbf{m}_k and the k -th covariance matrix \mathbf{C}_k for $k = 1, \dots, K$. Let \mathbf{A} be a full-rank matrix and \mathbf{b} be a translation vector. Then the random vector \mathbf{x} defined by*

$$\mathbf{x} = \mathbf{A}\xi + \mathbf{b}$$

has a mixture of multivariate normal distributions with the k -th mean

$$\mathbf{A}\mathbf{m}_k + \mathbf{b} \tag{42}$$

and the k -th covariance matrix

$$\mathbf{A}\mathbf{C}_k\mathbf{A}^T \tag{43}$$

without affecting proportions.

See Johnson et al. (2002) for the proof.

Proof of GMCG.

Subtracting Eq. (6) from Eq. (7), we obtain the following equation:

$$\mathbf{x}_{\text{cf}}^{(\tau:\cdot)} - \mathbf{x}_{\text{f}}^{(\tau:\cdot)} = \mathbf{W}^{(\tau:\cdot)}(\mathbf{s}_{\text{cf}} - \mathbf{s}_{\text{f}}).$$

Combining this with Eq. (10), we can write it as follows:

$$\mathbf{x}_{\text{cf}} = \mathbf{x}_{\text{f}} + \mathbf{W}(\mathbf{s}_{\text{cf}} - \mathbf{s}_{\text{f}}).$$

Using Eq. (11) to simplify the expression and rearranging the equation with respect to ξ , we can organize the equation in the form of \mathbf{A} and \mathbf{b} :

$$\begin{aligned} \mathbf{x}_{\text{cf}} &= \mathbf{x}_{\text{f}} + \mathbf{W}(\delta_a + N\xi) \\ &= (\mathbf{x}_{\text{f}} + \mathbf{W}\delta_a) + (\mathbf{W}N)\xi \\ &\equiv \mathbf{b} + \mathbf{A}\xi. \end{aligned}$$

Using Eq. (42) and (43) of Lemma 1, we finally obtain Eq. (16):

$$\mathbf{x}_{\text{cf}} \sim \sum_{k=1}^K p_k \mathcal{N}(\mathbf{x}; \mathbf{x}_{\text{f}} + \mathbf{W}(\delta_a + N\mathbf{m}_k), \mathbf{W}N\mathbf{C}_kN^T\mathbf{W}^T).$$