
End-to-End Balancing for Causal Continuous Treatment-Effect Estimation

Mohammad Taha Bahadori¹ Eric J. Tchetgen Tchetgen^{1,2} David E. Heckerman¹

Abstract

We study the problem of observational causal inference with continuous treatments in the framework of inverse propensity-score weighting. To obtain stable weights, we design a new algorithm based on entropy balancing that learns weights to directly maximize causal inference accuracy using end-to-end optimization. In the process of optimization, these weights are automatically tuned to the specific dataset and causal inference algorithm being used. We provide a theoretical analysis demonstrating consistency of our approach. Using synthetic and real-world data, we show that our algorithm estimates causal effect more accurately than baseline entropy balancing.

1. Introduction

In many applications, including business, social sciences, and health sciences, we wish to infer the causal effect of a continuous treatment, such as the effect of drug dose or administration duration on a health outcome variable. Often, confounding factors exist that influence both treatment and response variable, making it necessary to account for their impact to yield accurate causal estimation. While methods for doing so have been well studied when treatments are binary, causal inference with continuous treatments is far more challenging and largely understudied (Galagate, 2016; Ai et al., 2021). This is so primarily because continuous treatments induce uncountably many potential outcomes per unit (e.g., a subject), only one of which is observed for each unit and across units.

Propensity score weighting (Robins et al., 2000; Imai & Van Dyk, 2004), stand-alone or combined with regression-based models to achieve double robustness (Díaz & van der Laan, 2013; Kennedy et al., 2017; Nie et al., 2020), has become the state of the art for causal inference. If the weights,

inversely proportional to the conditional distribution of the treatment given the confounders, are correctly modeled, the weighted population will appear to come from a randomized study. However, this approach faces several challenges. One, the weights only balance the confounders in expectation, not necessarily in the given data (Zubizarreta et al., 2011). Two, the weights can be very large for some units, leading to unstable estimation and uncertain inference. As a possible remedy, entropy balancing (Hainmueller, 2012) proposes the use of an alternate set of weights such that they are less extreme but still have the property that the weight data appears to come from a randomized study.

In this work, we note that balanced weights do not directly optimize the accuracy of causal inference, and we introduce an alternative approach that does. We propose *End-to-End Balancing* (E2B), which uses synthetic data resembling the given data and end-to-end optimization to identify these weights. The E2B weights are thus customized for the given dataset and causal inference algorithm. Because we do not know the true treatment response function in the real data, we propose a new approach to generate synthetic training datasets and show that this approach is robust to model misspecification.

To theoretically analyze end-to-end balancing, we define *Generalized Stable Weights* (GSW) for causal inference as a generalization of the stable weights proposed by Robins et al. (2000). We prove that weights learned by entropy balancing for continuous treatments, including E2B weights, are unbiased estimators of GSWs. We also show that E2B weights are asymptotically consistent and efficient estimators of the population weights.

We perform three sets of experiments to demonstrate accuracy improvements by E2B. Two experiments with synthetic data, one with linear and another with non-linear response functions, and show that E2B is more accurate than baseline entropy balancing and inverse propensity score techniques. We also study the impact of mis-specification in the synthetic data generation process. In the experiments on real-world data, we qualitatively evaluate the average treatment effect function learned by E2B. We also qualitatively analyze the patterns in E2B weights.

¹Amazon.com, Inc. ²Wharton School of the University of Pennsylvania. Correspondence to: Mohammad Taha Bahadori <bahadorm@amazon.com>.

2. Problem Definition and Related Work

Problem Statement. Suppose we have the triplet of (\mathbf{x}, a, y) , where $\mathbf{x} \in \mathbb{X} \subset \mathbb{R}^r$, $a \in \mathbb{A} \subset \mathbb{R}$ and $y \in \mathbb{R}$ denote the confounders, treatments, and response variables, respectively, from an observational causal study. In our continuous treatment setting (Galagate, 2016, Ch. 1.2.6), we denote the value of y after intervening by setting the treatment to a as $y^{(a)}$. This quantity is commonly known as a *potential outcome*. Given an i.i.d. sample of size n , $\{(\mathbf{x}_i, a_i, y_i)\}_{i=1}^n$, our objective is to eliminate the impact of the confounders and identify the average treatment effect function $\mu(a) = \mathbb{E}[y^{(a)}]$, which is also called the response function. We make the two classic assumptions: (1) Strong ignorability: $y^{(a)} \perp\!\!\!\perp a \mid \mathbf{x}$. (i.e., no hidden confounders) and (2) Positivity: $0 < P(a|\mathbf{x}) < 1$.

General Causal Inference Literature. The literature on causal inference is vast and we refer the reader to the books for the general inquiry (Pearl, 2009; Imbens & Rubin, 2015; Spirtes et al., 2000; Peters et al., 2017). Instead, we focus on reviewing the inference techniques for *continuous* treatments. In particular, we narrow down our focus on propensity score weighting approaches (Robins et al., 2000; Imai & Van Dyk, 2004), because they can either be used alone or combined with the regression algorithms to create doubly robust algorithms.

Causal Inference via Weighting. In inverse propensity score weighting, a popular approach for causal inference, we create a pseudo-population by weighting data points such that, in this population, the confounders and treatments are independent. Thus, regular regression algorithms can estimate the causal response curve using the pseudo-population, which resembles data from randomized trials. Throughout this paper, we will denote the parameters of the pseudo-population with a tilde mark—for example, $\tilde{f}(a)$ denotes the marginal density of treatments in the weighted population. Multiple forms of propensity scores have been proposed for continuous treatments (Hirano & Imbens, 2004; Imai & Van Dyk, 2004). The commonly-used *stabilized weights* (Robins et al., 2000; Zhu et al., 2015) are defined as the ratio of marginal density over the conditional density of the treatments: $sw = f(a)/f(a|\mathbf{x})$.

Problems with Propensity Scores. Zubizarreta et al. (2011) list two challenges with the propensity scores: One, the weights only balance the confounders in expectation, not necessarily in the given data. (Two, the weights can be very large for some of the data points, leading to unstable estimations. The challenges are amplified in the continuous setting because computing the stabilized weights requires correctly choosing two models, one for the marginal and one for the conditional distributions of the treatments. Kang et al.

(2007) and Smith & Todd (2005) provide multiple pieces of evidence that the propensity score methods can lead to large biases in the estimations. Robins et al. (2007) propose techniques to fix the large weights problems in the binary treatment examples discussed by Kang et al. (2007) (Li et al., 2018; Zhao, 2019) describe techniques to learn more robust propensity scores for binary treatments. The case of continuous treatments, however, has received considerably less attention.

Entropy Balancing. To address the problem of extreme weights, *Entropy Balancing (EB)* (Hainmueller, 2012) estimates weights such that they balance the confounders subject to a measure of dispersion on the weights to prevent extremely large weights. Other loss functions using different dispersion metrics have been proposed for balancing (Zubizarreta, 2015; Chan et al., 2016). Zhao & Percival (2016) show that the entropy balancing is doubly robust. Entropy balancing has been extended to the continuous treatment setting (Fong et al., 2018; Vegetabile et al., 2021), where the balancing condition ensures that weighted correlations between the confounders and the treatment are zero. Ai et al. (2021) propose a method for estimating the counterfactual distribution in the continuous treatment setting.

3. Methodology

To describe our end-to-end balancing algorithm, we first need to describe entropy balancing for continuous treatments with base weights.

3.1. Entropy Balancing for Continuous Treatments

Causal Inference via Entropy Balancing. Entropy balancing creates a pseudo-population using instance weights $w_i \in \mathbb{R}_{>0}$, $i = 1, \dots, n$, in which the treatment a and the confounders \mathbf{x} are approximately independent from each other. The independence is enforced by first selecting a set of functions on the confounders $\phi_k(\cdot) : \mathbb{X} \mapsto \mathbb{R}$, for $k = 1, \dots, K$, that are dense and complete in L^2 space. Given the ϕ functions, we approximate the independence relationship by $\hat{\mathbb{E}}_n[a\phi_k(\mathbf{x})] = 0$, for $k = 1, \dots, K$, where the empirical expectation $\hat{\mathbb{E}}_n$ is performed on the pseudo-population. Hereafter, we will denote the mapped data points as $\phi(\mathbf{x}_i) = [\phi_1(\mathbf{x}_i), \dots, \phi_K(\mathbf{x}_i)]$. The $\phi_k(\cdot)$ functions can be chosen based on prior knowledge or defined by the penultimate layer of a neural network that predicts (a, y) from \mathbf{x} . Our contributions in this paper are orthogonal to the choice of the $\phi_k(\cdot)$ functions and can benefit from ideas on learning these functions (Zeng et al., 2020). The data-driven choice of the number of bases K is beyond the scope of current paper and left to future work.

Balancing Constraint for the Continuous Treatments.

Following (Fong et al., 2018; Vegetabile et al., 2021), in the case of continuous treatments, we first de-mean the confounders $\phi(x_i)$ and treatments a_i such that without loss of generality they are taken to have mean zero. The balancing objective is to learn a set of weights $w_i, i = 1, \dots, n$ that satisfy $\sum_{i=1}^n w_i \phi(x_i) = \mathbf{0}$, $\sum_{i=1}^n w_i a_i = 0$, and $\sum_{i=1}^n w_i a_i \phi(x_i) = \mathbf{0}$. We can write these three constraints in a compact form by defining a $(2K + 1)$ -dimensional vector $\mathbf{g}_i = [\phi(x_i), a_i, a_i \phi(x_i)]$. The constraints become $\sum_{i=1}^n w_i \mathbf{g}_i = \mathbf{0}$. We stack the \mathbf{g} vectors in a $(2K + 1) \times n$ dimensional matrix \mathbf{G} for compact notation. In this work, without loss of generality, we will present our idea with first order balancing, with neither higher order moments (Galagate, 2016) nor balancing in the kernel space (Wong & Chan, 2018; Kallus & Santacatterina, 2019; Hazlett, 2020).

Primal and Dual EB. A variety of dispersion metrics have been proposed as objective function for minimization, including entropy and variance of the weights (Wang & Zubizarreta, 2020). Hainmueller (2012) originally proposed minimizing the KL-divergence between the weights and a set of *base weights* $q_i, i = 1, \dots, n$. We discuss the details on choice of base weights below, however, note that $q_i = \text{constant}$ leads to maximization of the entropy of weights. Using this dispersion function and the balancing constraints, entropy balancing optimization is as follows:

$$\begin{aligned} \hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} \quad & \sum_{i=1}^n w_i \log \left(\frac{w_i}{q_i} \right), \\ \text{s.t.} \quad & \text{(i) } \mathbf{G}\mathbf{w} = \mathbf{0}, \\ & \text{(ii) } \mathbf{1}^\top \mathbf{w} = 1, \\ & \text{(iii) } w_i \geq 0 \text{ for } i = 1, \dots, n. \end{aligned} \quad (1)$$

The above optimization problem can be solved efficiently using its Lagrangian dual:

$$\hat{\lambda} = \underset{\lambda}{\operatorname{argmin}} \log \left(\mathbf{1}^\top \exp(-\lambda^\top \mathbf{G} + \ell) \right), \quad (2)$$

where $\ell_i = \log q_i$ are the *log-base-weights*. Given the solution $\hat{\lambda}$, the balancing weights can be computed as $\mathbf{w} = \text{softmax}(-\hat{\lambda}^\top \mathbf{G} + \ell)$, where $\text{softmax}(\mathbf{v}) = \exp \mathbf{v} / (\mathbf{1}^\top \exp \mathbf{v})$ for any vector \mathbf{v} . We select the mapping dimension K such that problem (2) is well-conditioned and leave the analysis of the high dimensional setting $K \approx n$ to future work. We can also add an L_1 penalty term to the dual objective in Eq. (2), which corresponds to approximate balancing (Wang & Zubizarreta, 2020).

An observation that is key to our approach is that each log-base-weight ℓ_i is a degree of freedom that we have in the Eq. (2) to improve the quality of causal estimation. In the next section we propose to parameterize the log-base-weights

and optimize them for causal inference using synthetic data. Our analysis in Section 4 shows that with arbitrary base weights, causal estimation using the weights learned in Eq. (2) will be consistent.

3.2. Learning the Base Weights

Hainmueller (2012) suggests two approaches for choosing base weights: (1) weights obtained from a conventional propensity score model and (2), in the context of survey design, using knowledge about the sampling design. We argue that a data-driven approach that learns customized base weights for a given dataset and weighted causal regression algorithm can further improve performance. For example, when we use weighted linear regression for causal inference, appropriate base weights can decrease the condition number of the design matrix and thus improve the quality of the regression. From another point of view, minimizing the KL-divergence between the weights w and the base weights q can act as a regularizer and improve the accuracy of weight estimation. Finally, as we will show in our algorithm, we can embed our prior distribution over possible response functions in the base weights and improve the causal inference accuracy.

To address this problem, we define the log-base-weights ℓ as a parametric function (e.g., a neural network) of the treatments and confounders—that is, $\ell_\theta(\mathbf{x}, a)$. We learn the base weights with the goal of improving the accuracy of the subsequent weighted regression. This is challenging because simply optimizing the weighted regression loss (e.g., weighted MSE) leads to degenerate results. That is, learning ℓ to minimize the loss in prediction of the outcome y will lead to exclusion of the difficult-to-predict data points from the regression, which is undesirable. Thus, we need to find another loss function to optimize, ideally a loss function that directly minimizes the error in estimation of the response function $\mu(a)$.

Our idea for learning the parameters of the base weights is to generate multiple pseudo-responses \bar{y} with randomly generated response functions $\bar{\mu}(a) \sim P_{\bar{\mu}}$. We will elaborate on the possible choices for the distribution $P_{\bar{\mu}}$ later in this section. Now that we know the true response function $\mu(a)$ in the randomly generated data, we can perform causal inference and obtain the estimation of the known response curve $\hat{\mu}(a)$ using our weights. Algorithm 1 outlines our stochastic training of the log-base-weight function. We also schematically visualize the algorithm in Figure 1. In Step 2, we estimate the distribution of noise using the residuals of regressing y over (\mathbf{a}, \mathbf{x}) , capturing the possible heteroskedasticity in the noise. Then, in each iteration, we draw a batch of possible datasets. To generate each dataset, we randomly choose a response function $\bar{\mu}(a)$ and use it to generate the entire dataset (see Section 5.1 for examples of

Algorithm 1 Stochastic Training of ℓ_θ for End-to-End Balancing

Require: Data tuples (\mathbf{x}_i, a_i, y_i) for $i = 1, \dots, n$ with an unknown response function $\mu(a)$.

Require: Representation functions $\phi(\cdot)$ and $\psi(\cdot)$, split size $n_1 < n$ and batch size B .

- 1: Generate a random set of indexes $I, |I| = n_1$ and its complement I^c and split the data to S and S^c using them.
- 2: Estimate the distribution of noise in y given (a, \mathbf{x}) as \hat{F}_ε .
- 3: Compute G by stacking $\mathbf{g}_i = [\phi(\mathbf{x}_i), a_i, a_i \phi(\mathbf{x}_i)]$, for $i = 1, \dots, n$.
- 4: **for** Number of Iterations **do**
- 5: Generate B datasets $\{(\mathbf{x}_i, a_i, \bar{y}_{i,b})\}_{i=1}^n$ for $b = 1, \dots, B$ using $\varepsilon \sim \hat{F}_\varepsilon$, and randomly selected $\bar{\mu}(a)_b$ response functions.
- 6: $\ell_i \leftarrow \ell_\theta(\psi(a_i, \mathbf{x}_i))$.
- 7: $\hat{\lambda} \leftarrow \arg\min_{\lambda} \{\log(\mathbf{1}^\top \exp(-\lambda^\top G + \ell)) + \gamma \|\lambda\|_1\}$ using only S data.
- 8: $\mathbf{w} \leftarrow \text{softmax}(-\hat{\lambda}^\top G + \ell)$ using only S^c data.
- 9: $\hat{\mu}(a)_b \leftarrow$ weighting-based causal estimates using $(a_i, \bar{y}_{i,b}, w_i)$ in S^c for $b = 1, \dots, B$.
- 10: Take a step in θ to minimize $\frac{1}{B} \sum_{b=1}^B (\hat{\mu}(a)_b - \bar{\mu}(a)_b)^2$.
- 11: **end for**
- 12: **return** The $\ell_{\hat{\theta}}$ function.

random functions). For the entire batch, we use ℓ_θ to learn the log-base-weights, and subsequently learn the weights in lines 7–8. In line 9 we use a weighted regression algorithm to find our estimation $\hat{\mu}(a)$ of the randomly-generated $\bar{\mu}(a)$. Our loss function is the mean squared error between the latter quantities. Our algorithm, which we call *End-to-End Balancing* (E2B), gets its name from this end-to-end optimization.

Sample Splitting. The E2B procedure involves estimation of two sets of parameters θ in the ℓ_θ and λ for entropy balancing. The joint estimation of θ, λ on a single sample will result in bias (Chernozhukov et al., 2018). Thus, we split the sample to two mutually exclusive parts and perform the optimizations on separate partitions of data.

Distribution of Random Response Functions $P_{\bar{\mu}}$. Ideally, we should rely on domain experts or rough ideas from past studies for choosing the random set of response functions $\bar{\mu}(a)$ that includes the true response function. This will help us embed our prior knowledge in the base weights and improve the accuracy of causal inference. Alternatively, we can choose broad function classes such as random piecewise smooth functions or polynomial functions with random co-

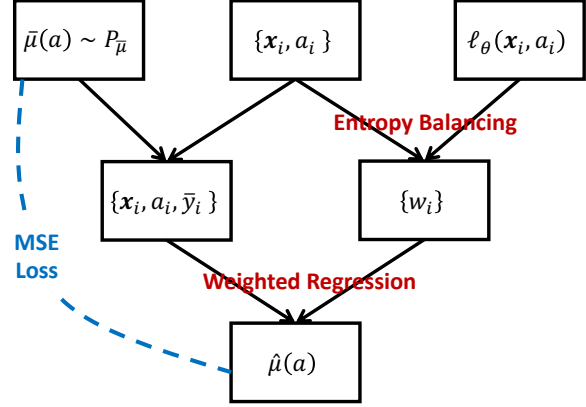


Figure 1: Overview of the steps in End-to-End Balancing.

efficients. See Section 5 for examples of generating random response functions. We can also use generative models to generate data that is more similar to our sample (Athey et al., 2021).

Features Fed to ℓ_θ . We can feed the raw values of the treatments and any handcrafted features, denoted by $\psi(a_i, \mathbf{x}_i)$. We empirically find that $\psi(a_i, \mathbf{x}_i) = (\log p(a_i), \log p(a_i | \mathbf{x}_i))$ makes training the ℓ_θ easier. These features are the logarithms of the numerator and denominator of the stable weights. Given this choice, we can visualize the ℓ_θ function and find its relationship to the marginal and conditional distributions of the treatment. We describe the details of our neural network model for ℓ_θ and our techniques for training in Appendix B.

Weighted Regression Algorithms. To be able to differentiate the loss function with respect to θ , we need weighted regression algorithms whose estimates are differentiable with respect to the weights. In the linear average treatment effect function we choose weighted linear regression and in the non-linear setting we use the weighted polynomial regression and the local kernel regression, as used by Flores et al. (2012).

Double Robustness. Zhao & Percival (2016) show that, in the binary case, entropy balancing is doubly robust. We do not attempt to show double robustness for E2B because we see E2B as a meta algorithm that learns customized weights for each dataset and algorithm. We can either (1) plug-in the E2B weights in the doubly robust algorithm and expect improved accuracy, or (2) learn weights that directly minimize the error of doubly robust algorithms such as (Diaz & van der Laan, 2013; Kennedy et al., 2017). However, if we use an outcome regression model for generating the random response functions $\bar{\mu}(a)$, the E2B weights may no longer provide significant improvements to the doubly robust tech-

niques.

4. Analysis

We prove that for any arbitrary choice of the log-base-weight function ℓ_θ , our approach consistently estimates causal effects. Before proving the consistency results, we characterize the quantity that our solution converges to.

Definition 1. Generalized Stable Weights. Suppose $f(a, \mathbf{x})$ denote the joint probability density function of treatments and confounders in a population. Suppose $\tilde{f}(a)$ and $\tilde{f}(\mathbf{x})$ denote two arbitrary density functions, possibly different from the marginal density functions in our population, that satisfy $\mathbb{E}_{\mathbf{x} \sim \tilde{f}(\mathbf{x})}[\mathbf{x}] = \mathbf{0}$ and $\mathbb{E}_{a \sim \tilde{f}(a)}[a] = 0$. We define the Generalized Stable Weights as follows

$$w_{GSW}(a, \mathbf{x}) = \frac{\tilde{f}(a)\tilde{f}(\mathbf{x})}{f(a, \mathbf{x})}. \quad (3)$$

Remark. Our definition generalizes the stabilized weights defined by Robins et al. (2000) in that their weights require $\tilde{f}(a)$ and $\tilde{f}(\mathbf{x})$ match the marginal probability density functions in the original population, whereas GSW allows these densities to be less constrained. The next proposition motivates the generalization.

Proposition 1. The generalized stable weights w_{GSW} satisfy $\mathbb{E}[w_{GSW} \mathbf{a} \mathbf{x}] = \mathbf{0}$.

Proof.

$$\begin{aligned} \mathbb{E}[w_{GSW} \mathbf{a} \mathbf{x}] &= \mathbb{E}\left[\frac{\tilde{f}(a)\tilde{f}(\mathbf{x})}{f(a, \mathbf{x})} \mathbf{a} \mathbf{x}\right] \\ &= \int \int \frac{\tilde{f}(a)\tilde{f}(\mathbf{x})}{f(a, \mathbf{x})} \mathbf{a} \mathbf{x} dF_{a, \mathbf{x}}(a, \mathbf{x}) \\ &= \int a \tilde{f}(a) da \int \mathbf{x} \tilde{f}(\mathbf{x}) d\mathbf{x} = \mathbf{0}, \end{aligned}$$

where the last equation holds due to the zero mean assumption for the $\tilde{f}(a)$ and $\tilde{f}(\mathbf{x})$ distributions. \square

Now, we can show that with an appropriate choice of the ϕ functions, the solution of Eq. (2) approximates the generalized stable weights. Consider the population version of Eq. (2):

$$\boldsymbol{\lambda}^* = \underset{\boldsymbol{\lambda}}{\operatorname{argmin}} \log(\mathbb{E}[\exp(\mathbf{g}^\top \boldsymbol{\lambda} + \ell)]). \quad (4)$$

The weights corresponding to $\boldsymbol{\lambda}^*$ can be calculated as $w^* = C \exp(\mathbf{g}^\top \boldsymbol{\lambda}^* + \ell)$, where $C = (\int \exp(\mathbf{g}^\top \boldsymbol{\lambda}^* + \ell) dF(a, \mathbf{x}))^{-1}$ is the normalization constant.

Assumptions.

1. $f(a, \mathbf{x}) \geq c > 0$ for all $(a, \mathbf{x}) \in \mathbb{A} \times \mathbb{X}$ pairs, where c is a constant.
2. Suppose the basis functions are dense and rich enough such for some small values of δ_{ϕ_K} that they satisfy:

$$\mathbb{E}[\mathbf{a} \phi(\mathbf{x})] = \mathbf{0} \text{ only if } \sup_{a, \mathbf{x}} |f(a, \mathbf{x}) - f(a)f(\mathbf{x})| = \delta_{\phi_K}.$$
3. Suppose the population problem in Eq. (4) has a unique solution $\boldsymbol{\lambda}^*$ and the corresponding weights are denoted by w^* .

The following theorem shows that the solution to Eq. (4) converges to w_{GSW} :

Theorem 1. Given the assumptions, the solution to the population problem satisfies:

$$\sup_{a, \mathbf{x}} |w^*(a, \mathbf{x}) - w_{GSW}(a, \mathbf{x})| \leq \delta_{\phi_K}/c. \quad (5)$$

If we select the function set ϕ_K such that $\delta_{\phi_K} = o(1)$, the theorem shows that $w^*(a, \mathbf{x})$ is an unbiased estimator of $w_{GSW}(a, \mathbf{x})$. Notice that Assumption 1 is only slightly stronger than the common positivity assumption. Assumption 2 requires us to select the mapping functions such that zero correlation between the mapped confounders and the treatment implies their independence, approximately. We provide the proof in Appendix A.1.

Note that the quality of the ψ features and neural network training does not affect the unbiasedness of the E2B because of the balancing constraint is still satisfied. The flexibility in choice of \tilde{f} distributions in the definition of w_{GSW} is due to the fact that we require only first order balancing, enforcing zero correlation between the first order moments of the treatment and confounders. If we enforce higher order balancing constraints in the form of $\mathbb{E}[w^* \phi_1(a) \phi_2(\mathbf{x})] = \mathbb{E}[w^* \phi_1(a)] \cdot \mathbb{E}[w^* \phi_2(\mathbf{x})]$ for any suitable functions ϕ_1 and ϕ_2 , Theorem 1 in (Ai et al., 2021) shows that $w^* = f(a)/f(a|\mathbf{x})$. The more flexible form of weights in Eq. (3) allows us to pick the marginals $\tilde{f}(a)$ and $\tilde{f}(\mathbf{x})$ with more freedom. In this work, we have chosen a data-driven way to learn them.

Finally, the following theorem establishes the asymptotic consistency and normality result for each individual weight estimated by E2B, under the common regularity conditions for problem (2).

Theorem 2. Suppose $\Lambda \subset \mathbb{R}^{2K+1}$ is an open subset of Euclidean space and the solution $\hat{\boldsymbol{\lambda}}_n \in \Lambda$ to Eq. (2) is

within the subset. The weights estimated by Eq. (2) are asymptotically normal for $i = 1, \dots, n$:

$$\sqrt{n} (\hat{w}_n(a_i, \mathbf{x}_i) - w^*(a_i, \mathbf{x}_i)) \xrightarrow{d} \mathcal{N}(0, \sigma^2(a_i, \mathbf{x}_i)). \quad (6)$$

We provide the population form of $\sigma^2(a_i, \mathbf{x}_i)$ and an unbiased sample estimate for it in Appendix A.2.

5. Experiments

We use two synthetic and one real-world datasets to evaluate the performance of E2B. In the synthetic datasets, we have access to the true treatment effects; thus we measure accuracy of the algorithms in recovering the treatment effects. In the real-world data, we qualitatively evaluate the estimated causal treatment effect curve and inspect the learned log-base-weight function.

Baselines. One baseline in our study is the Inverse Propensity score Weighting (IPW) with Stable Weights (Robins et al., 2000), perhaps the most commonly used technique. To avoid extreme weights and prevent instability, we trim (Winsorize) the weights by [5, 95] percentiles (Cole & Hernán, 2008; Crump et al., 2009; Chernozhukov et al., 2018). Another baseline in our experiments is Entropy Balancing (Vegetabile et al., 2021), which is equal to E2B with $\ell_\theta = \text{constant}$. EB allows us to do an ablation study and see the exact amount of improvement when learning a customized ℓ_θ function. We also include EB with the stabilized weights (SW) as base weights ($\ell_\theta = \log \hat{p}(a) - \log \hat{p}(a|\mathbf{x})$). For fair comparison to the entropy balancing methods, we only include first order balancing methods, as our idea of learning the base-weights can be combined with the higher order and kernel-based balancing methods. Finally, we also include the permutation weighting algorithm (Arbour et al., 2021) that proposes to compute the weights using permutations of the treatments and a classifier that predicts the probability of being permuted. We provide further details on this algorithm in Appendix B.4.

Training Details. All neural networks are trained using Adam (Kingma & Ba, 2014) with early stopping based on validation error. The learning rate and architectural parameters of the neural networks are tuned via hyperparameter search on the validation data. We provide the details of the neural networks used for the ℓ_θ and propensity score estimation for IPW in Appendix B.

Regularization. Proper tuning of the regularization term γ in Line 7 of Algorithm 1 is critical to the performance of EB and E2B. Our first observation is that using L_2 regularization allows more stable tuning of the regularization coefficient γ , though the final accuracies do not significantly

change. We tune the regularization coefficient for EB by measuring the bootstrapped covariate balance, as described by Wang & Zubizarreta (2020). However, for E2B, we have access to the validation error on the pseudo-datasets and use the validation error to tune γ . In the experiments, we observe that tuning with validation error leads to smaller values for the penalization coefficient γ and more accurate causal inference.

5.1. Synthetic Data Experiments

Linear. We use the following steps to generate 100 datasets, each with 1000 data points.

1. Generate confounders $\mathbf{x} \in \mathbb{R}^5$, $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Sigma)$, where Σ is a tridiagonal covariance matrix with diagonal and off-diagonal elements equal to 1.0 and 0.2, respectively.
2. $\mathbf{a} \sim \mathcal{N}(\mu_a, 0.3^2)$, where $\mu_a = \sin(\beta_{xa}^\top \mathbf{x})$ and $\beta_{xa,k} \sim \text{Unif}(-1, 1)$ for $k = 1, \dots, 5$.
3. $y \sim \mathcal{N}(\mu_y, 0.5^2)$, where $\mu_y = \beta_{xy}^\top \mathbf{x} + \beta_{ay} \mathbf{a}$, where $\beta_{ax}, \beta_{xy,k} \sim \mathcal{N}(0, 1)$ for $k = 1, \dots, 5$.

We use weighted least squares as the regression algorithm and report the average $|\widehat{\beta_{ay}} - \beta_{ay}|$ over all 100 datasets.

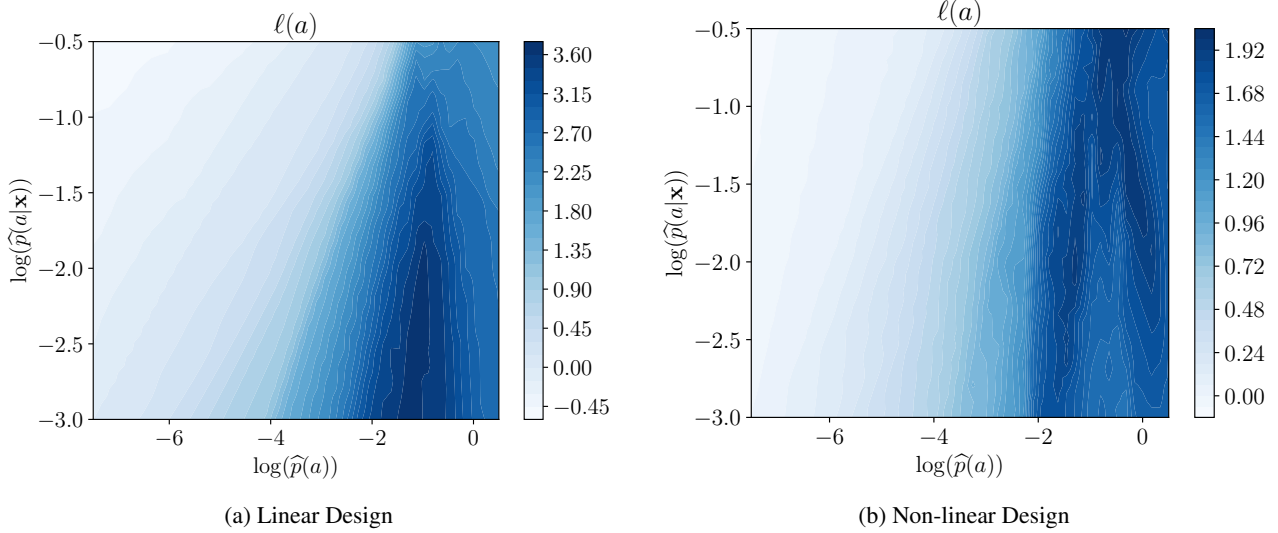
Nonlinear We first generate confounders \mathbf{x} and treatments \mathbf{a} similar to steps 1 and 2 of the linear case. Then, we generate the response variable according to $y \sim \mathcal{N}(\mu_y, 0.5^2)$, where $\mu_y = \beta_{xy}^\top \mathbf{x} + h_{\gamma_{ay}}(\mathbf{a})$, where $\beta_{xy,k} \sim \mathcal{N}(0, 1)$ for $k = 1, \dots, 5$, and $h_\gamma(z) = \gamma_0 + \gamma_1 z + \gamma_2 (z^2 - 1) + \gamma_3 (x^3 - 3x)$ are Hermit polynomials. Similar to the linear case, we generate 100 samples of size 1000. We use the weighted polynomial regression as the regression algorithm to estimate $\hat{\gamma}$ and report the average RMSE between true γ and $\hat{\gamma}$. We report the mean and standard error of errors on 100 datasets in Table 1.

As shown in Table 1, in both linear and non-linear datasets, the E2B is significantly more accurate in uncovering the true treatment response functions. Both constant and IPW base weights perform worse than the base-weights learned by end-to-end balancing. Comparing EB and E2B, we see that the optimal penalization coefficients that we find are 100 and 16, respectively. This indicates that hyperparameter tuning with validation loss —available only to E2B— leads to smaller penalization coefficients. This decreased penalization produces stronger covariate balance for the E2B, compared to EB, and thus contributes to the higher accuracy of E2B.

To gain more insights, in Figure 2, we plot the log-base weight function that we learn as a function of $\log(\hat{p}(a))$ and $\log(\hat{p}(a|\mathbf{x}))$. We align all curves at their starting point and

Table 1: Average RMSE for estimation of the response functions. The results are in the form of “mean (std. err.)” from 100 runs.

Algorithm	Linear	Non-linear
Inverse Propensity Weighting (SW)	2.057 (0.437)	0.530 (0.025)
Permutation Weighting	1.1543 (6.580)	0.525 (0.250)
Entropy Balancing (Const.)	0.626 (0.041)	0.417 (0.024)
Entropy Balancing (SW)	0.412 (0.029)	0.399 (0.021)
End-to-End Balancing	0.299 (0.029)	0.249 (0.011)


 Figure 2: The estimated log-base-weight function ℓ_θ as a function of logarithms of the empirical density of the treatment $\log(\hat{p}(a))$ and conditional distribution $\log(\hat{p}(a|x))$. We perform the experiment 100 times and report the median and the inter-quantile range. We align all curves by normalizing their value at the beginning to zero.

plot the median of 100 runs. Looking at the y-axis, especially in the linear case, the smaller conditional probability leads to larger base weights. This in effect regularizes the weights to be closer to $\hat{p}(a|x)$, with a scaling factor. Both figures, show variations in the $\log(\hat{p}(a))$ -axis. This means that the algorithm finds the high-density regions to be more trustworthy for causal inference. Finally, the complexity of the plots emphasizes the need for end-to-end methods for learning weights. Given the results in Figure 2, the learned log-base-weight does not seem to be a convex combination of uniform and SW base weights.

Note that, as [Robins et al. \(2007\)](#) caution, synthetic data evaluation might exacerbate the extreme weights issue because, unlike when analyzing real data, usually no manual inspection of weights is done.

Mis-specification of Random Response Functions. To assess the sensitivity of the E2B to the choice of random response functions $\bar{\mu}(a)$, we perform an additional experi-

ment in which the class of random response functions does not include the true response function. In our experiment, the true response function is a nonlinear function (a Hermite polynomial of degree 3) but we use linear random response functions $\bar{\mu}(a)$ for training ℓ_θ . We generate 100 datasets with different randomly selected response functions and report the mean and standard error of RMSE in estimation of the response functions. We also measure the validation error on the datasets generated using the linear response functions.

In our experiments with misspecification, the RMSE is 0.254 (0.013). It shows that misspecification of the random response function does not significantly impact the accuracy of E2B and robustness of E2B to misspecification. The result also shows that the gain is not due to the prior knowledge used for generation of the random response functions. Note that the hyperparameter tuning selects the same regularization parameter because during validation we validate with the misspecified model.

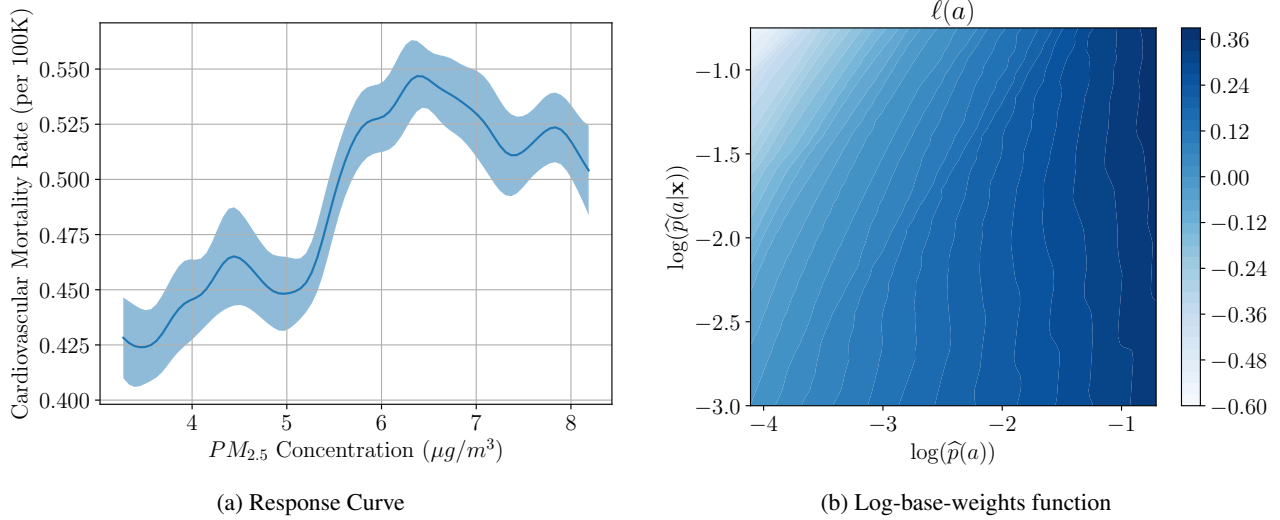


Figure 3: (a) The average treatment effect curve for measuring the impact of $PM_{2.5}$ concentration on the cardiovascular mortality rate. We perform the experiment 100 times and report the mean and \pm std range. (b) The estimated log-base-weight function ℓ_θ as a function of logarithm of the empirical density of the treatment $\log(\hat{p}(a))$.

5.2. Real Data Experiments

We study the impact of $PM_{2.5}$ particle level on the cardiovascular mortality rate (CMR) in 2132 counties in the US using the data provided by the National Studies on Air Pollution and Health (Rappold, 2020). The data is publicly available under U.S. Public Domain license. The $PM_{2.5}$ particle level and the mortality rate are measured by $\mu g/m^3$ and the number of annual deaths due to cardiovascular conditions per 100,000 people, respectively. We use only the data for 2010 to simplify the experiment setup; thus we measure the same year impact of $PM_{2.5}$ particle level. Other than the treatment and response variables, the data includes 10 variables such as poverty rate, population, and household income, which we use as confounders. We provide the descriptive statistics and the histograms for the treatment and effect in Appendix C.

To train E2B, we create the random dataset (Line 5 in Algorithm 1) using Hermite polynomials of max degree 3, $\mu_y = |h_{\gamma_{xy}}(\beta_{xy}^\top \mathbf{x} / \|\beta_{xy}^\top \mathbf{x}\|_2) + h_{\gamma_{ay}}(a)|$. We use absolute value to capture the positivity of our response variable. The data also shows heteroskedasticity; we model the noise as a zero mean Gaussian variable (Line 2 of the algorithm) with variance $\sigma^2(\hat{y}) = 6.00\hat{y}$. For regression, we use the non-parametric local kernel regression algorithm. We measure the uncertainty in the curves using the deep ensembles technique (Lakshminarayanan et al., 2017) with 100 random ensembles. That means, in each experiment, we initialize the neural network with different random values. To further improve the uncertainty estimation, in each training, we resample the dataset as well.

Figure 3a shows the average treatment effect curve for the impact of $PM_{2.5}$ on CMR. We show the one standard deviation interval using the shaded areas. Starting around $PM_{2.5} = 5.3\mu g/m^3$ the curve increases with a steep slope, confirming the previous studies that increased $PM_{2.5}$ levels increase the probability of cardiovascular mortality. We can see that after $PM_{2.5} = 6.4\mu g/m^3$ the curve plateaus and mortality rate stays at elevated levels. Looking at the histogram of the treatments in Figure 4a in the appendix, we observe that most counties have $PM_{2.5}$ between 6 and 8. This might justify the fluctuations that we see in this interval and may allude to existence of potential unmeasured confounders.

Figure 3b shows the log-base-weight function that we learn in this data. Similar to the synthetic experiments, we show the median of 100 runs. While the plot shows smaller variations, it is generally inline with the observations we had in the synthetic data.

6. Discussion

Causal inference is a well-studied problem. Its main goal is to remove biases due to confounding by balancing the population to look similar to randomized controlled trials. Removing the impact of confounders can play a critical role in reducing and possibly eliminating bias in our decision making leading to potentially positive societal impacts. Our results rely on two classical assumptions: (1) no unmeasured confounders and (2) positivity. While these assumptions are sometimes reasonable in practice, their violation might lead to biased causal inferences. For example, the positivity

assumption might be violated if we do not collect any data for a sub-population. Overall, the debiasing property of causal inference should not relieve us from rigorous data collection and analysis. In our experiments, we have been careful to quantify uncertainty in our causal estimation and be wary of over-confidence in our results. We performed our experiments on a CPU machine with 16 cores on AWS, in a region that uses hydroelectric power.

7. Conclusion

We observe that, in the entropy balancing framework, the base weights provide an extra degree of freedom to optimize the accuracy of causal inference. We propose end-to-end balancing (E2B) as a technique to learn the base weights such that they directly improve the accuracy of causal inference using end-to-end optimization. In our theoretical analysis we find that E2B weights are approximating Generalized Stable Weights and discuss E2B's statistical consistency. Using synthetic and real-world data, we show that our proposed algorithm outperforms entropy balancing in terms of causal inference accuracy.

References

- Ai, C., Linton, O., and Zhang, Z. Estimation and inference for the counterfactual distribution and quantile functions in continuous treatment models. *Journal of Econometrics*, 2021.
- Arbour, D., Dimmery, D., and Sondhi, A. Permutation weighting. In *ICML*, 18–24 Jul 2021.
- Athey, S., Imbens, G. W., Metzger, J., and Munro, E. Using wasserstein generative adversarial networks for the design of monte carlo simulations. *Journal of Econometrics*, 2021.
- Chan, K. C. G., Yam, S. C. P., and Zhang, Z. Globally efficient non-parametric inference of average treatment effects by empirical balancing calibration weighting. *Journal of the Royal Statistical Society. Series B, Statistical methodology*, 78(3):673, 2016.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. Double/debiased machine learning for treatment and structural parameters: Double/debiased machine learning. *The Econometrics Journal*, 21(1), 2018.
- Cole, S. R. and Hernán, M. A. Constructing inverse probability weights for marginal structural models. *American journal of epidemiology*, 168(6):656–664, 2008.
- Crump, R. K., Hotz, V. J., Imbens, G. W., and Mitnik, O. A. Dealing with limited overlap in estimation of average treatment effects. *Biometrika*, 96(1):187–199, 2009.
- Díaz, I. and van der Laan, M. J. Targeted data adaptive estimation of the causal dose–response curve. *Journal of Causal Inference*, 1(2):171–192, 2013.
- Flores, C. A., Flores-Lagunes, A., Gonzalez, A., and Neumann, T. C. Estimating the effects of length of exposure to instruction in a training program: the case of job corps. *Review of Economics and Statistics*, 94(1): 153–171, 2012.
- Fong, C., Hazlett, C., Imai, K., et al. Covariate balancing propensity score for a continuous treatment: application to the efficacy of political advertisements. *The Annals of Applied Statistics*, 12(1):156–177, 2018.
- Galagate, D. *Causal Inference With a Continuous Treatment and Outcome: Alternative Estimators for Parametric Dose-response Functions With Applications*. PhD thesis, University of Maryland, 2016.
- Hainmueller, J. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political analysis*, pp. 25–46, 2012.
- Hazlett, C. Kernel balancing: A flexible non-parametric weighting procedure for estimating causal effects. *Statistica Sinica*, 2020.
- Hirano, K. and Imbens, G. W. The propensity score with continuous treatments. *Applied Bayesian modeling and causal inference from incomplete-data perspectives*, 226164:73–84, 2004.
- Imai, K. and Van Dyk, D. A. Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association*, 99(467): 854–866, 2004.
- Imbens, G. W. and Rubin, D. B. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- Kallus, N. and Santacatterina, M. Kernel optimal orthogonality weighting: A balancing approach to estimating effects of continuous treatments. *arXiv preprint arXiv:1910.11972*, 2019.
- Kang, J. D., Schafer, J. L., et al. Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical science*, 22(4):523–539, 2007.
- Kennedy, E. H., Ma, Z., McHugh, M. D., and Small, D. S. Nonparametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, 79(4): 1229, 2017.

- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In *NeurIPS*, pp. 6405–6416, 2017.
- Li, F., Morgan, K. L., and Zaslavsky, A. M. Balancing covariates via propensity score weighting. *Journal of the American Statistical Association*, 113(521):390–400, 2018.
- Nie, L., Ye, M., Nicolae, D., et al. Vcnet and functional targeted regularization for learning causal effects of continuous treatments. In *International Conference on Learning Representations*, 2020.
- Pearl, J. *Causality*. Cambridge university press, 2009.
- Peters, J., Janzing, D., and Schölkopf, B. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- Rappold, A. Annual PM2.5 and cardiovascular mortality rate data: Trends modified by county socioeconomic status in 2, 132 US counties, 2020. URL <https://edg.epa.gov/metadata/catalog/search/resource/details.page?uuid=https://doi.org/10.23719/1506014>.
- Robins, J., Hernán, M., and Brumback, B. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000.
- Robins, J., Sued, M., Lei-Gomez, Q., and Rotnitzky, A. Comment: Performance of double-robust estimators when "inverse probability" weights are highly variable. *Statistical Science*, 22(4):544–559, 2007.
- Smith, J. A. and Todd, P. E. Does matching overcome lalonde’s critique of nonexperimental estimators? *Journal of econometrics*, 125(1-2):305–353, 2005.
- Spirtes, P., Glymour, C. N., Scheines, R., and Heckerman, D. *Causation, prediction, and search*. MIT press, 2000.
- Van der Vaart, A. W. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- Vegetabile, B. G., Griffin, B. A., Coffman, D. L., Cefalu, M., Robbins, M. W., and McCaffrey, D. F. Nonparametric estimation of population average dose-response curves using entropy balancing weights for continuous exposures. *Health Services and Outcomes Research Methodology*, 21(1):69–110, 2021.
- Wang, Y. and Zubizarreta, J. R. Minimal dispersion approximately balancing weights: asymptotic properties and practical considerations. *Biometrika*, 107(1):93–105, 2020.
- Wong, R. K. and Chan, K. C. G. Kernel-based covariate functional balancing for observational studies. *Biometrika*, 105(1):199–213, 2018.
- Zeng, S., Assaad, S., Tao, C., Datta, S., Carin, L., and Li, F. Double robust representation learning for counterfactual prediction. *arXiv:2010.07866*, 2020.
- Zhao, Q. Covariate balancing propensity score by tailored loss functions. *The Annals of Statistics*, 47(2):965–993, 2019.
- Zhao, Q. and Percival, D. Entropy balancing is doubly robust. *Journal of Causal Inference*, 5(1), 2016.
- Zhu, Y., Coffman, D. L., and Ghosh, D. A boosting algorithm for estimating generalized propensity scores with continuous treatments. *Journal of causal inference*, 3(1): 25–40, 2015.
- Zubizarreta, J. R. Stable weights that balance covariates for estimation with incomplete outcome data. *Journal of the American Statistical Association*, 110(511):910–922, 2015.
- Zubizarreta, J. R., Reinke, C. E., Kelz, R. R., Silber, J. H., and Rosenbaum, P. R. Matching for several sparse nominal variables in a case-control study of readmission following surgery. *The American Statistician*, 65(4):229–238, 2011.

A. Proofs of the Theorems

Whenever the context of an expectation operation is not clear, we disambiguate it by specifying the variable that the expectation is taken over and its distribution $\mathbb{E}_{\mathbf{x} \sim f(\mathbf{x})}[\cdot]$.

A.1. Proof of Theorem 1

Proof. Given that the logarithm function is a strictly increasing function, we can omit it in the optimization; i.e., $\boldsymbol{\lambda}^* = \operatorname{argmin}_{\boldsymbol{\lambda}} \mathbb{E} [\exp(\mathbf{g}^\top \boldsymbol{\lambda} + \ell)]$. Because this is an unconstrained optimization, the optimal solution occurs when the gradient is equal to zero.

$$\begin{aligned} \mathbb{E}[\mathbf{g} \exp(\mathbf{g}^\top \boldsymbol{\lambda}^* + \ell)] &= \mathbf{0}, \\ \mathbb{E}[\mathbf{g} w^*(\mathbf{a}, \mathbf{x})] &= \mathbf{0}, \end{aligned} \quad (7)$$

where the last equation is due to the equation of the weights in the population optimization problem (Eq. 4).

Using the definition for the \mathbf{g} vector, Eq. (7) implies that $\mathbb{E}[w^*(\mathbf{a}, \mathbf{x}) \mathbf{a} \phi(\mathbf{x})] = \mathbf{0}$. Thus, we conclude that in the weighted population (with distribution \tilde{F}), the \mathbf{a} and $\phi(\mathbf{x})$ are uncorrelated:

$$\mathbb{E}_{(\mathbf{a}, \mathbf{x}) \sim \tilde{F}}[\mathbf{a} \phi(\mathbf{x})] = \mathbf{0} \quad (8)$$

For every set $\mathcal{B} \subset \mathbb{A} \times \mathbb{X}$, we can write:

$$\tilde{F}(\mathcal{B}) = \int_{\mathcal{B}} w^*(a, \mathbf{x}) dF(a, \mathbf{x}). \quad (9)$$

The Radon-Nikodym theorem implies that $w^*(a, \mathbf{x})$ is the Radon-Nikodym derivative:

$$w^*(\mathbf{x}, a) = \frac{d\tilde{F}(\mathbf{x}, a)}{dF(\mathbf{x}, a)} = \frac{\tilde{f}(\mathbf{x}, a)}{f(\mathbf{x}, a)} \quad (10)$$

$$= \frac{\tilde{f}(\mathbf{x})\tilde{f}(a) + \{\tilde{f}(\mathbf{x}, a) - \tilde{f}(\mathbf{x})\tilde{f}(a)\}}{f(\mathbf{x}, a)}, \quad (11)$$

$$= w_{GSW}(a, \mathbf{x}) + \frac{\tilde{f}(\mathbf{x}, a) - \tilde{f}(\mathbf{x})\tilde{f}(a)}{f(\mathbf{x}, a)} \quad (12)$$

Thus, using Eq. (8) and Assumptions 1 and 3 we can write

$$\sup_{a, \mathbf{x}} |w^*(a, \mathbf{x}) - w_{GSW}(a, \mathbf{x})| \leq \delta_{\phi_K}/c.$$

□

A.2. Theorem 2

Proof. Given that the logarithm function is a strictly increasing function, we can omit it in the optimizations. Thus the sample and population solutions are:

$$\begin{aligned} \hat{\boldsymbol{\lambda}}_n &= \operatorname{argmin}_{\boldsymbol{\lambda}} \frac{1}{n} \sum_{i=1}^n \exp(\mathbf{g}_i^\top \boldsymbol{\lambda} + \ell_i), \\ \boldsymbol{\lambda}^* &= \operatorname{argmin}_{\boldsymbol{\lambda}} \mathbb{E} [\exp(\mathbf{g}^\top \boldsymbol{\lambda} + \ell)]. \end{aligned}$$

The estimator is an M-estimator and given our sample-splitting, the proof follows the asymptotic normality of the estimator (Van der Vaart, 2000, Chapter 5.3).

$$\sqrt{n} (\hat{\boldsymbol{\lambda}}_n - \boldsymbol{\lambda}^*) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{V}), \quad (13)$$

To obtain the value of \mathbf{V}_1 , note that the optimal sample solution occurs at the solution of the following equation (Z-estimator equation):

$$\sum_{i=1}^n \mathbf{g}_i \exp(\mathbf{g}_i^\top \hat{\boldsymbol{\lambda}}_n) = \mathbf{0}.$$

Thus, the score function is $\psi_{\lambda} = \mathbf{g}_i \exp(\mathbf{g}_i^{\top} \lambda)$. We denote the matrix of derivatives of the score function by $\dot{\psi}_{\lambda}$ whose elements are defined as $\dot{\psi}_{\lambda, kk'} = \partial \psi_{\lambda, k} / \partial \lambda_{k'}$. Using the theorem in (Van der Vaart, 2000, Chapter 5.3), we can write:

$$\mathbf{V} = \mathbb{E}[\dot{\psi}_{\lambda^*}]^{-1} \mathbb{E}[\psi_{\lambda^*} \psi_{\lambda^*}^{\top}] \mathbb{E}[\dot{\psi}_{\lambda^*}]^{-1}. \quad (14)$$

In the above equation we have assumed that $\mathbb{E}[\dot{\psi}_{\lambda^*}]$ matrix is invertable. An unbiased sample estimation of \mathbf{V} can be obtained by substituting $\hat{\lambda}_n$ in place of λ^* and taking empirical expectations.

An application of the delta method on Eq. (13) yields:

$$\sqrt{n} \left(\frac{\exp(\mathbf{g}_i^{\top} \hat{\lambda}_n + \ell_i)}{\frac{1}{n} \sum_{i=1}^n \exp(\mathbf{g}_i^{\top} \hat{\lambda}_n + \ell_i)} - \frac{\exp(\mathbf{g}_i^{\top} \lambda^* + \ell_i)}{\frac{1}{n} \sum_{i=1}^n \exp(\mathbf{g}_i^{\top} \lambda^* + \ell_i)} \right) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \sigma^2), \quad (15)$$

$$\sqrt{n} \left(\hat{w}_n(a_i, \mathbf{x}_i) - \frac{\exp(\mathbf{g}_i^{\top} \lambda^* + \ell_i)}{\mathbb{E}[\exp(\mathbf{g}^{\top} \lambda^* + \lambda)]} \right) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \sigma^2), \quad (16)$$

$$\sqrt{n} (\hat{w}_n(a_i, \mathbf{x}_i) - w^*(a_i, \mathbf{x}_i)) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \sigma^2), \quad (17)$$

where Eq. (16) is due to Slutsky's theorem and Eq. (17) is obtained by substitution of the definition for $w^*(a_i, \mathbf{x}_i)$. The variance is obtained by defining the Softmax function $s(\lambda) = \frac{\exp(\mathbf{g}_i^{\top} \lambda + \ell_i)}{\frac{1}{n} \sum_{i=1}^n \exp(\mathbf{g}_i^{\top} \lambda + \ell_i)}$. We denote the gradient of the Softmax function by $\nabla s(\lambda)$. We can write (Van der Vaart, 2000, Chapter 3):

$$\sigma^2(a_i, \mathbf{x}_i) = \nabla s(\lambda^*)^{\top} \mathbf{V} \nabla s(\lambda^*).$$

Substituting the value of \mathbf{V} from (14), we conclude:

$$\sigma^2(a_i, \mathbf{x}_i) = \nabla s(\lambda^*)^{\top} \mathbb{E}[\dot{\psi}_{\lambda^*}]^{-1} \mathbb{E}[\psi_{\lambda^*} \psi_{\lambda^*}^{\top}] \mathbb{E}[\dot{\psi}_{\lambda^*}]^{-1} \nabla s(\lambda^*).$$

Note that the value of the softmax function depends on the value of (a_i, \mathbf{x}_i) at each point. □

B. Neural Network and Training Details

B.1. Details of the ℓ_{θ} Neural Network

The ℓ_{θ} network is defined as follows:

$$\ell_{\theta}(z) = cz + \text{dense3}(\text{elu}(\text{layer_norm}(\text{dense2}(\tanh(\text{dense1}(z))))))$$

The linear term cz acts as a skip connection. The input and output dimensions for the dense linear layers are as follows:

$$\text{dense1} : 1 \mapsto h,$$

$$\text{dense2} : h \mapsto h,$$

$$\text{dense3} : h \mapsto 1,$$

where h denotes the hidden dimension. Because the softmax function is invariant to the constant shifts, we do not have any bias terms for dense3 and the skip connection. dense2 also does not have the bias because of the proceeding layer normalization. The dimension h has been tuned as a hyperparameter on a validation data and set to 10.

B.2. Details of the Propensity Score Computation for IPW

We model both $f(a)$ and $f(a|\mathbf{x})$ as univariate normal distributions. This is the correct assumption in our synthetic data. The marginal distribution $f(a)$ is estimated by simply finding the mean and standard deviation of the observed treatment values. For the conditional distribution, we write $a|\mathbf{x} \sim \mathcal{N}(\mu_a(\mathbf{x}), \sigma_{a|\mathbf{x}}^2)$, where $\mu_a(\mathbf{x})$ is modeled using a feedforward neural network with two layers and $\sigma_{a|\mathbf{x}}^2$ is estimated using the residuals of the neural network predictions. The dimension of the neural network has been tuned as a hyperparameter on validation data and set to 30.

B.3. Further Training Details

We used PyTorch to implement E2B. For reproducibility purposes, we provide the final settings used for training:

- Learning algorithm: Adam with learning rate 0.001, no AMSGrad.
- Batch size: 1000
- Max epochs: 400
- Weight decay: 2.5×10^{-5} .
- Validation on a dataset of size 400, every 10 steps.

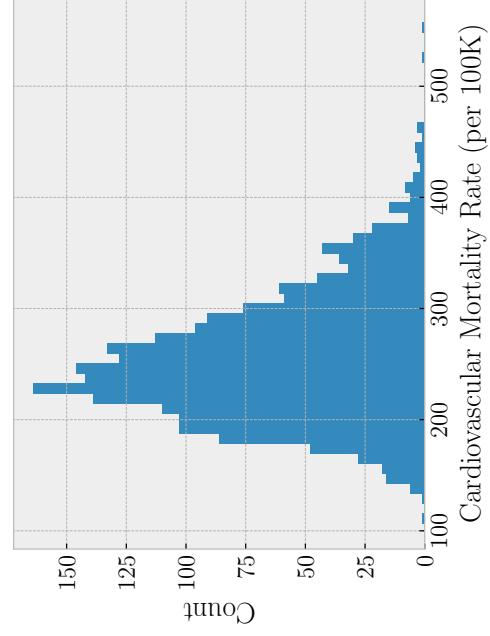
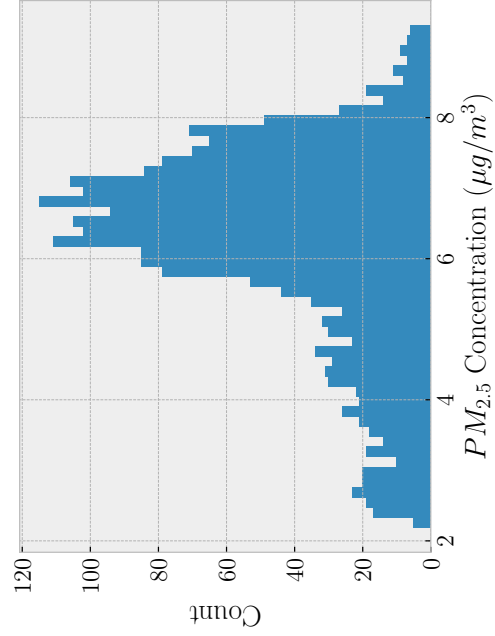
B.4. Detail of Permutation Weighting

We created the stacked data by stacking $\{(\mathbf{x}_i, a_i, a_i \odot \mathbf{x}_i)\}_{i=1}^n$ and $\{(\mathbf{x}_i, \tilde{a}_i, \tilde{a}_i \odot \mathbf{x}_i)\}_{i=1}^n$, where \tilde{a} are permutations of the original treatments. We trained a random forest classifier to predict whether each data is from the permuted or the original set. We tried both random forests and neural networks and obtained better results with the former. We also calibrated the predicted probabilities of the classifier before computation of the weights.

C. Data and Preprocessing Description

Table 2: NSAPH Data Description

	PM2.5	CMR	healthfac	population	ses	unemploy	HH_inc	femaleHH	vacant	owner_occ	eduattain	pctfam_pover
count	2132.0	2132.0	2132.0	2132.0	2132.0	2132.0	2132.0	2132.0	2132.0	2132.0	2132.0	2132.0
mean	6.17	255.25	0.18	10.78	0.0	7.85	10.69	11.92	14.25	71.44	35.03	11.25
std	1.45	56.76	0.5	1.26	0.96	2.83	0.24	3.94	8.71	7.76	7.07	5.2
min	2.19	106.14	-2.85	6.2	-1.84	0.0	9.91	2.1	3.8	19.3	9.4	0.0
25%	5.51	215.38	0.0	10.04	-0.67	6.0	10.54	9.3	8.8	67.7	30.4	7.6
50%	6.43	248.16	0.14	10.62	-0.14	7.6	10.67	11.2	11.65	72.7	35.4	10.55
75%	7.15	288.77	0.34	11.46	0.47	9.3	10.83	13.6	16.6	76.7	39.9	13.82
max	9.3	557.43	3.33	16.07	6.46	30.9	11.66	38.0	74.0	89.7	54.6	44.9

(a) Histogram of $PM_{2.5}$

(b) Histogram of Cardiovascular Mortality Rate

Figure 4: The Histograms of Data