

MODIFIED TREATMENT POLICY EFFECT ESTIMATION WITH WEIGHTED ENERGY DISTANCE

BY ZIREN JIANG¹, JARED D. HULING²

¹*Division of Biostatistics and Health Data Science, University of Minnesota, Minneapolis, MN, U.S.A.*

²*Division of Biostatistics and Health Data Science, University of Minnesota, Minneapolis, MN, U.S.A. huling@umn.edu*

The causal effects of continuous treatments are often characterized through the average dose response function, which is challenging to estimate from observational data due to confounding and positivity violations. Modified treatment policies (MTPs) are an alternative approach that aim to assess the effect of a modification to observed treatment values and work under relaxed assumptions. Estimators for MTPs generally focus on estimating the conditional density of treatment given covariates and using it to construct weights. However, weighting using conditional density models has well-documented challenges. Further, MTPs with larger treatment modifications have stronger confounding and no tools exist to help choose an appropriate modification magnitude. This paper investigates the role of weights for MTPs showing that to control confounding, weights should balance the weighted data to an unobserved hypothetical target population that can be characterized with observed data. Leveraging this insight, we present a versatile set of tools to enhance estimation for MTPs. We introduce a distance that measures imbalance of covariate distributions under the MTP and use it to develop new weighting methods and tools to aid in the estimation of MTPs. Using our methods we study the effect of mechanical power of ventilation on in-hospital mortality.

1. Introduction. Understanding the causal effects of changes in doses or otherwise continuous values of a treatment is important in many scientific disciplines. A typical approach to quantifying the causal effect of varying a continuous treatment is to estimate the causal average dose response function (ADRF), which is the expected potential outcome as a function of all possible or likely treatment values. In our motivating application, the treatment is the mechanical power of a ventilator applied to patients in intensive care who have acute respiratory distress syndrome. In this setting, the ADRF is the expected in-hospital mortality as a function of the power of ventilation. While the ADRF provides an intuitive way to evaluate the causal effect of a continuous treatment, its estimation requires one to assume that every individual could hypothetically receive all possible values of the treatment, which is often unlikely in practice as some values may be clinically unlikely or impossible for subsets of the population. For example, in mechanically-ventilated patients with acute respiratory distress, ventilation guidelines restrict the values of mechanical ventilation settings based on measures of the patient's condition, making it impossible in practice for some patients to receive some ranges of settings. Thus, use of the ADRF in this setting may not have clinical meaning. Moreover, confounding is exacerbated in causal ADRF estimation because the characteristics of subjects with a higher treatment value may differ substantively from those with a lower value. For example, ventilation guidelines suggest explicit actions be taken with regards to mechanical ventilation settings based on patient characteristics, inducing strong

Keywords and phrases: Continuous Treatments, Covariate Balance, Balancing weights, Causal Inference, Observational Studies.

and complex confounding. Additionally, the performance of causal ADRF estimators is often poor (i.e. cannot attain root- n consistency nonparametrically), even with flexible, doubly robust methods (Kennedy et al., 2017).

Alternative definitions of causal effects in this setting have been introduced in part to help alleviate some of these problems while still providing clinically useful results. One of the alternatives is to imagine a hypothetical world where the treatment may be a random or deterministic function of the observed treatment value (Stock, 1989; Robins, Hernán and Siebert, 2004). Muñoz and Van Der Laan (2012) proposed stochastic interventions which focus on the effect of assigning each subject’s intervention based on a random draw from a given distribution depending on their own characteristics. Haneuse and Rotnitzky (2013) further proposed modified treatment policies (MTPs) which generalize the idea of stochastic interventions. Each subject’s counterfactual treatment under a given MTP is defined as a function of their baseline characteristics and their *observed* values of the treatment without the MTP (the “natural” value of the treatment); the MTP thus imagines a slight manipulation of each individual’s treatment value from their actual values. The estimand is then defined as the expected potential outcome under the specified manipulation. For example, in our motivating application, one could estimate the effect of slightly reducing the mechanical power of ventilation, compared with standard of care, on in-hospital mortality among patients in the intensive care unit (ICU) with acute respiratory distress (Neto et al., 2018).

MTPs can help policymakers evaluate hypotheses about the treatment and generate practical interventions that can be later tested via experiments. Another advantage of MTPs is that, by imagining counterfactual worlds where the treatment is modified slightly from reality, the corresponding counterfactual world is not substantially different from reality, making positivity hold by construction. This further allows for studying interventions that do not differ starkly from current practice, potentially helping clinical uptake of findings. MTPs are an attractive tool for causal inference as they can be quite general and can be used in a wide variety of settings. For example, Díaz et al. (2021) generalized MTPs to time-varying treatment settings and Díaz and Hejazi (2020) developed stochastic interventions for causal mediation analysis.

Confounding is a major hurdle in causal inference from observational studies, both for standard estimands and for MTPs. Weighting approaches to confounding control re-weight each subject to balance the distribution of the covariates and do not require the use of outcome information. Traditionally, weights are generated through inverse-probability weighting (IPW) which models the treatment assignment mechanism given covariates and inverts it. However, the performance of IPW relies critically on the choice of the propensity score model and model misspecification can result in severe bias (Kang and Schafer, 2007). Alternative weighting methods have been proposed to mitigate this issue in the setting of estimating standard causal effects of discrete-valued treatments by using weights that encourage (Imai and Ratkovic, 2014) or enforce (Hainmueller, 2012) balance of pre-specified moments of covariates. Particularly relevant to our article is the work by Huling and Mak (2024) who consider weights that minimize the energy distance between the weighted empirical distribution of the intervention groups and the joint population, aiming to balance the full joint distributions of covariates. This nonparametric distributional balancing approach has been shown to work empirically well with minimal need for careful modeling decisions.

Several challenging issues remain unresolved for weighting methods for MTPs. Existing weighting methods for MTPs have focused on the estimation of the conditional density of treatment given the covariates, also known as the generalized propensity score (GPS); weights are a ratio of the GPS at two different values. Muñoz and Van Der Laan (2012) and Haneuse and Rotnitzky (2013) both proposed weighting estimators based on estimation of the GPS. Díaz et al. (2021) proposed weights based on direct modeling of the density ratio.

These methods show great promise, but their success hinges on accurately specifying the conditional density or density ratio model. Yet, estimating a conditional density is highly challenging especially with an increasing number of covariates (Huling, Greifer and Chen, 2023). Even slight misspecification of the estimated density function from the truth can impact the performance of the IPW-style estimators of the ADRF (Naimi et al., 2014); these issues persist in the estimation of MTPs. As the role of covariate balance has not been explored for MTPs, no balancing weight methods exist for them, exacerbating these challenges. Moreover, the impact of weights on finite sample performance is not yet fully understood and no diagnostic tools are available to assess the performance of a given set of weights or to understand whether a given set of weights is likely to yield an unconfounded comparison for a given dataset. As MTPs involve investigating the effect of a shift of treatment values from their observed values, larger shifts are likely to yield analyses with more confounding and with more potential for positivity violations. Yet no tools exist that can assess what range of shifts in a class of MTPs are “safe” (i.e. have confounding that can be fully adjusted using a given weighting method) for a given application. Clear guidance is thus needed for practitioners to determine an appropriate magnitude of the MTP shift.

In this work, we introduce distance-based tools that provide solutions to the aforementioned issues. Based on the extension of the identification result of the modified treatment policy from Haneuse and Rotnitzky (2013), we clearly define the form of the target counterfactual population for an MTP (i.e., MTP-shifted population) and provide a clear mathematical expression of it. Even though the target population is purely hypothetical, we show its distribution can be estimated using the observed data, a fact used extensively in our methods. This builds the connection between the complex problem of MTPs and stochastic interventions and the well-understood problem of covariate balancing. By examining the role of the weights in finite sample error of a generic weighting estimator, we present a novel error decomposition to show that the estimation error of a weighted estimate for the causal effect of MTP is directly related to the imbalance of the weighted empirical distribution (referring to the empirical distribution of the weighted observed sample) and the MTP-shifted empirical distribution. We provide a measure of this distance with a modification of the energy distance (Székely and Rizzo, 2013), a computationally simple measure of the distance between two multivariate distributions, however other distributional measures can also be used. It has been used to measure distributional imbalance in the context of categorical treatments (Huling and Mak, 2024) and to quantify dependence between continuous treatments and confounders (Huling, Greifer and Chen, 2023). In this paper, we provide 1) a novel testing framework for the assessment of a feasible range of shifts for a class of MTPs for arbitrary weighting methods for a given dataset and 2) methods for comparison of different weighting methods for MTPs. Finally, we introduce energy balancing weights for MTPs, a weighting approach for MTP estimation which reduces *distributional* imbalance to the target population, reducing sensitivity of findings to model misspecification of the conditional density function. Although this paper primarily focuses on MTPs, we show that our error decomposition, methods, and setup are straightforwardly extended to stochastic interventions (Kennedy, 2019; Díaz and Hejazi, 2020; Hejazi et al., 2023), allowing the use of our methods for a far broader class of settings beyond MTPs. See Section 2.3 and Supplementary Material Section 3 for more detail. We conclude our manuscript with an analysis of the impact of changes to mechanical power of ventilation from current practice on in-hospital mortality.

2. Setup.

2.1. *Causal framework for modified treatment policies (MTPs).* Consider data $\{\mathbf{X}_i, A_i, Y_i\}_{i=1}^n$ collected from an observational study, where $\mathbf{X}_i \in \mathcal{X} \subseteq \mathbb{R}^p$ denotes a p -dimensional vector

of pre-treatment covariates for subject i , $A_i \in \mathcal{A} \subseteq \mathbb{R}$ denotes the received value of the treatment, and $Y_i \in \mathbb{R}$ denotes the outcome. Here we assume A_i is a continuous treatment variable (e.g., the taken dose of a treatment for individual i). We adopt the potential outcome framework in which we assume that for each subject i , there exists a potential outcome Y_i^a defined as the outcome that subject i would have if subject i is intervened on to take the treatment at value $A = a$. An MTP analysis starts with an analyst-specified treatment policy $q(\mathbf{x}, a)$ that inputs patient characteristics and observed treatment values and outputs a modified value of the treatment. The goal of an MTP analysis is to estimate the mean potential outcome $\mu^q \equiv \mathbb{E}[Y^{q(\mathbf{X}, A)}]$ under this policy; this can then be contrasted with the average *observed* outcome $\mathbb{E}[Y^A] = \mathbb{E}[Y]$ to understand whether the policy improves or harms outcomes on average. To ensure Y_i^a is well-defined (i.e., Y_i^a is unique for a given treatment level a and subject i), we assume the following consistency condition:

- (A0 Consistency): For any subject i in the population, if $A_i = a$ then $Y_i = Y_i^a$.

One of the main strengths of MTPs is that they can satisfy the positivity assumption by construction. In many scientific applications, it is not realistic for some patients to have extreme values of treatment, e.g. extremely low mechanical power of ventilation for patients who have serious respiratory symptoms (see Section 8 for the case study). As such, calculating the causal effect of an arbitrary treatment value a may not be meaningful since this treatment value may not be feasible or applicable to the entire population. Alternatively, as in the work of [Haneuse and Rotnitzky \(2013\)](#), MTPs can be justified with a much weaker assumption on positivity (see assumption A1 below). Specifically, by carefully designing the MTP $q(\mathbf{x}, a)$, we can be more confident that for subject i with covariate \mathbf{x}_i and observed treatment a_i , there is a positive density to observe a subject with the same covariate \mathbf{x}_i and treatment value $q(\mathbf{x}, a)$.

MTPs are pertinent to clinical practice, as they can provide insight into how modifications to clinical practice could have changed patient outcomes. By considering the effects of a collection of MTPs, researchers can gain insights into the underlying mechanisms that lead to different outcomes and identify potential strategies for improving patient outcomes. For example, to evaluate a care policy that encourages a slight reduction in intensity of mechanical ventilation, we could estimate the causal effect of the MTP $q(\mathbf{x}, a) = a - 2$ which reduced mechanical power for all subjects by 2 Joules/min compared with current standards. The MTP could be further refined to modify the amount of reduction differently for people with different characteristics or prior ventilation patterns.

For any $(\mathbf{x}, a) \in (\mathcal{X}, \mathcal{A})$, where $(\mathcal{X}, \mathcal{A})$ is the joint support of \mathbf{X} and A , we can define $\mathbb{E}(Y^{q(\mathbf{x}, a)} | \mathbf{X} = \mathbf{x}, A = a)$ to be the expected potential outcome averaging over all subjects who have $\mathbf{X} = \mathbf{x}, A = a$. The mean potential outcome of the MTP $q(\mathbf{x}, a)$ is then the mean of the conditional average potential outcome $\mathbb{E}(Y^{q(\mathbf{x}, a)} | \mathbf{X} = \mathbf{x}, A = a)$ over the entire population, i.e.,

$$(1) \quad \mu^q := \int_{(\mathcal{X}, \mathcal{A})} \mathbb{E}[Y^{q(\mathbf{x}, a)} | \mathbf{X} = \mathbf{x}, A = a] dF_{\mathbf{X}, A}(\mathbf{x}, a)$$

where $F_{\mathbf{X}, A}(\mathbf{x}, a)$ is the distribution function of (\mathbf{X}, A) . In order to make the integral well-defined, we follow [Haneuse and Rotnitzky \(2013\)](#) to assume the continuity of $\mathbb{E}[Y^{q(\mathbf{x}, a)} | \mathbf{X} = \mathbf{x}, A = a]$ over $(\mathbf{x}, a) \in (\mathcal{X}, \mathcal{A})$ and continuity of $Pr(A | \mathbf{X} = \mathbf{x})$ over $\mathbf{x} \in \mathcal{X}$. In order to estimate the causal effect from observational data, μ^q needs to be causally-identifiable. In other words, we need to express μ^q as a function of the distribution of (Y, \mathbf{X}, A) . From [Haneuse and Rotnitzky \(2013\)](#), μ^q can be identified with the following assumptions:

- (A1 Positivity): If $(\mathbf{x}, a) \in (\mathcal{X}, \mathcal{A})$ then $(\mathbf{x}, q(\mathbf{x}, a)) \in (\mathcal{X}, \mathcal{A})$.

- (A2 Conditional exchangeability of related populations) For each $(\mathbf{x}, a) \in (\mathcal{X}, \mathcal{A})$, let $a' = q(\mathbf{x}, a)$, then $Y^{a'} | \mathbf{X} = \mathbf{x}, A = a$ and $Y^{a'} | \mathbf{X} = \mathbf{x}, A = a'$ have the same distribution.

The positivity assumption states that for each subject in the population, it is always possible to find some other subjects who have the same covariates and receive the modified treatment. The second assumption state that the potential outcome of the modified treatment will not be affected by the original treatment allocation. In other words, subjects with $\mathbf{X} = \mathbf{x}$ who received treatment a could have received treatment $q(\mathbf{x}, a)$. The difference between A2 and the usual no-unmeasured confounders assumption is that the latter is much stronger, as it requires the subjects with $\mathbf{X} = \mathbf{x}$ who received treatment a could have received any possible dose $a' \in \mathcal{A}_x$ where $\mathcal{A}_x := \{a : (\mathbf{x}, a) \in (\mathcal{X}, \mathcal{A})\}$ be the support of A given $\mathbf{X} = \mathbf{x}$.

In order to have a closed-form derivation of our proposed estimator, we further require the MTP to have a piecewise differentiable inverse function.

- (C1 Piece-wise smooth invertibility): For each $\mathbf{x} \in \mathcal{X}$, there exists a partition $\{I_{j,\mathbf{x}}\}_{j=1}^{J(\mathbf{x})}$ of \mathcal{A}_x where $q(\mathbf{x}, \cdot)$ is smooth and invertible within the partition. Specifically, let $q_j(\mathbf{x}, \cdot)$ denote the function $q(\mathbf{x}, \cdot)$ on the $I_{j,\mathbf{x}}$ part. Then $q_j(\mathbf{x}, \cdot)$ has differentiable inverse function $h_j(\mathbf{x}, \cdot)$ on the interior of $I_{j,\mathbf{x}}$, such that $h_j(\mathbf{x}, q_j(\mathbf{x}, a)) = a$

Under Condition C1, for a given $\mathbf{x} \in \mathcal{X}$, the MTP $q(\mathbf{x}, a)$ can be expressed as the sum of the functions $\{q_j(\mathbf{x}, a)\}_{j=1}^{J(\mathbf{x})}$ on each interval $\{I_{j,\mathbf{x}}\}_{j=1}^{J(\mathbf{x})}$, $q(\mathbf{x}, a) = \sum_{j=1}^{J(\mathbf{x})} I_{j,\mathbf{x}}(a) q_j(\mathbf{x}, a)$ where $I_{j,\mathbf{x}}(a)$ be the indicator function s.t. $I_{j,\mathbf{x}}(a) = 1$ if $a \in I_{j,\mathbf{x}}$ and $I_{j,\mathbf{x}}(a) = 0$ otherwise. Condition C1 is important, as it allows us to cleanly express the expected potential outcome under the MTP as a function of the observed data distribution. With C1, we have the following identification result that extended from [Haneuse and Rotnitzky \(2013\)](#),

$$(2) \quad \mu^q = \int_{\mathcal{X}} \sum_{j=1}^{J(\mathbf{x})} \int_{h_j(\mathbf{x}, a) \in I_{j,\mathbf{x}}} \mathbb{E}(Y | \mathbf{X} = \mathbf{x}, A = a) dF_{\mathbf{X},A}(\mathbf{x}, h_j(\mathbf{x}, a)).$$

The proof is in the Supplementary Material Section 7.

2.2. Novel error decomposition for weighted MTP estimators. We present a novel error decomposition that allows a clear inspection of the role of sample weights in finite sample errors when estimating μ^q using weighting methods. This decomposition provides insights that enable us to provide a measure of confounding bias as a function of the sample weights. This measure then allows us to provide tools that compare weights and further tools to assess when a given MTP is so far from the observed treatments that measured confounding is excessive and/or too difficult to control with weights.

Denoting $\mu(\mathbf{x}, a) \equiv \mathbb{E}(Y | \mathbf{X} = \mathbf{x}, A = a)$, we have $Y_i = \mu(\mathbf{X}_i, A_i) + \epsilon_i$ where $\epsilon_i \equiv Y_i - \mu(\mathbf{X}_i, A_i)$ is the error term with mean zero. We can express the estimand (2) as

$$(3) \quad \mu^q = \int_{(\mathcal{X}, \mathcal{A})} \mu(\mathbf{x}, a) dF_{\mathbf{X},A}^q(\mathbf{x}, a), \text{ where}$$

$$(4) \quad F_{\mathbf{X},A}^q(\mathbf{x}, a) = \sum_{j=1}^{J(\mathbf{x})} I_{j,\mathbf{x}}(h_j(\mathbf{x}, a)) F_{\mathbf{X},A}(\mathbf{x}, h_j(\mathbf{x}, a))$$

and $F_{\mathbf{X},A}(\mathbf{x}, a)$ is the CDF of (\mathbf{X}, A) . In the Supplementary Materials Section 7.2, we prove that $F_{\mathbf{X},A}^q(\mathbf{x}, a)$ is the CDF of $(\mathbf{X}, q(\mathbf{X}, A))$, as long as $q(\mathbf{x}, a)$ has the property that $\lim_{(\mathbf{x}, a) \rightarrow -\infty} q(\mathbf{x}, a) = -\infty$ and $\lim_{(\mathbf{x}, a) \rightarrow \infty} q(\mathbf{x}, a) = \infty$. The significance of (3) will be made clear in our error decomposition below.

In this paper, we focus on weighted estimators of μ^q which do not require one to use outcome information, allowing the analyst to conduct objective causal inference. A weighting estimator with arbitrary sample weights $\mathbf{w} = (w_1, \dots, w_n)$ can be expressed as:

$$\hat{\mu}_{\mathbf{w}}^q = \frac{1}{n} \sum_{i=1}^n w_i Y_i = \frac{1}{n} \sum_{i=1}^n w_i \mu(\mathbf{X}_i, A_i) + \frac{1}{n} \sum_{i=1}^n w_i \epsilon_i.$$

For the following, we further require that $\sum_{i=1}^n w_i = n$ so that $F_{n,\mathbf{w},\mathbf{X},\mathbf{A}} = \frac{1}{n} \sum_{i=1}^n w_i I(\mathbf{X}_i \leq \mathbf{x}, A_i \leq a)$ can be interpreted as a weighted empirical CDF of the the covariates and treatment values. Now, we can express the error of a weighted estimator $\hat{\mu}_{\mathbf{w}}^q$ with arbitrary weights as:

$$(5) \quad \hat{\mu}_{\mathbf{w}}^q - \mu^q = \underbrace{\int_{(\mathcal{X},\mathcal{A})} \mu(\mathbf{x}, a) d(F_{n,\mathbf{w},\mathbf{X},\mathbf{A}} - F_{n,\mathbf{X},\mathbf{A}}^q)(\mathbf{x}, a)}_{\text{error due to confounding}}$$

$$(6) \quad + \underbrace{\int_{(\mathcal{X},\mathcal{A})} \mu(\mathbf{x}, a) d(F_{n,\mathbf{X},\mathbf{A}}^q - F_{\mathbf{X},\mathbf{A}}^q)}_{\text{sampling error}} + \frac{1}{n} \sum_{i=1}^n w_i \epsilon_i,$$

where $F_{n,\mathbf{X},\mathbf{A}}^q(\mathbf{x}, a) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^{J(\mathbf{X}_i)} I_{j,\mathbf{X}_i}(A_i) I(\mathbf{X}_i \leq \mathbf{x}, q_j(\mathbf{X}_i, A_i) \leq a)$ is the empirical estimator of the shifted distribution $F_{\mathbf{X},\mathbf{A}}^q$. The error decomposition emphasizes the importance of sample weights in mitigating bias. The second term is sampling error. The third term is affected by the variance of the weights, however, it always has mean 0. Thus, bias can be mitigated by minimizing the first term $\int_{(\mathcal{X},\mathcal{A})} \mu(\mathbf{x}, a) d(F_{n,\mathbf{w},\mathbf{X},\mathbf{A}} - F_{n,\mathbf{X},\mathbf{A}}^q)(\mathbf{x}, a)$ which depends on the distance between the two distribution functions and the form of the outcome function $\mu(\mathbf{x}, a)$.

Note that $F_{n,\mathbf{X},\mathbf{A}}^q(\mathbf{x}, a)$ is the empirical CDF of the MTP-shifted random variable $(\mathbf{X}, q(\mathbf{X}, A))$ and $F_{n,\mathbf{w},\mathbf{X},\mathbf{A}}(\mathbf{x}, a) = \frac{1}{n} \sum_{i=1}^n w_i I(\mathbf{X}_i \leq \mathbf{x}, A_i \leq a)$ is the weighted empirical CDF of random variable (\mathbf{X}, A) . Thus, the term $\int_{(\mathcal{X},\mathcal{A})} \mu(\mathbf{x}, a) d(F_{n,\mathbf{w},\mathbf{X},\mathbf{A}} - F_{n,\mathbf{X},\mathbf{A}}^q)(\mathbf{x}, a)$ depends on both the weights \mathbf{w} and the pre-specified modified treatment policy $q(\mathbf{x}, a)$. Larger shifts on the MTP will tend to lead to the larger difference between the distribution of $(\mathbf{X}, q(\mathbf{X}, A))$ and (\mathbf{X}, A) which can make the confounding bias term larger, generally speaking.

The significance of the error decomposition above is that both the key distributional terms, $F_{n,\mathbf{w},\mathbf{X},\mathbf{A}}(\mathbf{x}, a)$ and $F_{n,\mathbf{X},\mathbf{A}}^q(\mathbf{x}, a)$, contributing to the error term in (5) are entirely functions of the observed data. Thus, while the outcome function is unknown, it is possible to compare $F_{n,\mathbf{w},\mathbf{X},\mathbf{A}}(\mathbf{x}, a)$ with $F_{n,\mathbf{X},\mathbf{A}}^q(\mathbf{x}, a)$ to characterize bounds on bias due to measured confounding.

2.3. Extension to stochastic interventions. Although our paper is focused on the modified treatment policy effect, most of the assertions can be easily extended to the context of stochastic intervention which is an alternative causal framework for estimating the effect of continuous intervention (Muñoz and Van Der Laan, 2012; Díaz et al., 2021; Kennedy, 2019). Instead of fixing the shifted intervention $q(\mathbf{X}, A)$ to a single, deterministic value, stochastic interventions take the shifted intervention value as a random draw from a modified intervention distribution. In Supplementary Material Section 3, we establish an error decomposition result analogous to (5) for the estimation of the stochastic intervention effects. As a result, all of our proposed tools are directly applicable to stochastic interventions, allowing for a far more general use of our methods.

3. Weighted Energy Distance for MTPs. With the key component of confounding bias in weighted estimates of the causal effect of an MTP, the next step is developing an outcome-agnostic manner measure or bound of the magnitude of (5) for a given set of weights. Doing so would enable evaluation of a set of weights in the specific context of a given MTP for a given dataset. A challenge in quantifying the bias term is that it depends on 1) an unknown mean function and 2) the distance between a weighted and an MTP-shifted distribution. Since the bias component $\int_{(\mathcal{X}, \mathcal{A})} \mu(\mathbf{x}, a) d(F_{n, \mathbf{w}, \mathbf{X}, A} - F_{n, \mathbf{X}, A}^q)(\mathbf{x}, a)$ is an integral over the difference of the two empirical distributions, a metric that can measure the distance between these two distributions can be used for this purpose. In this paper, we adopt the energy distance (Székely and Rizzo, 2013; Huling and Mak, 2024), a metric based on powers of the Euclidean distance, for use in estimating weights that minimize the key component of the confounding bias. Although there are other measures of distributional distance, we focus on the energy distance due to its simplicity and lack of need for tuning parameters, making it a widely applicable tool for a broad range of investigators. However, other distances can be used in its place, such as the maximum mean discrepancy (MMD) which is the distance of two distributions embedded to the reproducing kernel Hilbert spaces (RKHS). MMDs with the distance kernel have been shown to be equivalent to the energy distance when it is calculated with a semimetric of negative type (Sejdinovic et al., 2013). In Supplementary Material Section 2.2, we show an MMD formulation as an alternative measure.

Following Huling and Mak (2024), we generalize the energy distance (see Supplementary Material Section 2.1 for the definition) to the weighted energy distance which measures the distance between the weighted empirical CDF $F_{n, \mathbf{w}, \mathbf{X}, A}$ and the MTP-shifted empirical CDF $F_{n, \mathbf{X}, A}^q$,

$$\begin{aligned}
 & \mathcal{E}(F_{n, \mathbf{w}, \mathbf{X}, A}, F_{n, \mathbf{X}, A}^q) \\
 &= \frac{1}{n^2} \left\{ 2 \times \sum_{i=1}^n \sum_{k=1}^n \sum_{j=1}^{J(\mathbf{X}_k)} I_{j, \mathbf{X}_k}(A_k) w_i \|(\mathbf{X}_i, A_i) - (\mathbf{X}_k, q_j(\mathbf{X}_k, A_k))\|_2 \right. \\
 & \quad \left. - \sum_{i=1}^n \sum_{k=1}^n w_i w_k \|(\mathbf{X}_i, A_i) - (\mathbf{X}_k, A_k)\|_2 \right. \\
 & \quad \left. - \sum_{i=1}^n \sum_{j=1}^{J(\mathbf{X}_i)} \sum_{k=1}^n \sum_{j'=1}^{J(\mathbf{X}_k)} I_{j, \mathbf{X}_i}(A_i) I_{j', \mathbf{X}_k}(A_k) \|(\mathbf{X}_i, q_j(\mathbf{X}_i, A_i)) - (\mathbf{X}_k, q_{j'}(\mathbf{X}_k, A_k))\|_2 \right\}.
 \end{aligned}
 \tag{7}$$

Since the MTP-shifted empirical CDF $F_{n, \mathbf{X}, A}^q$ can be viewed as the empirical CDF for the sample $\{\mathbf{X}_i, q(\mathbf{X}_i, A_i)\}_{i=1}^n$, we can define its empirical characteristic function as

$$\varphi_{n, \mathbf{X}, A}^q(\mathbf{t}) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^{J(\mathbf{X}_i)} I_{j, \mathbf{X}_i}(A_i) \exp(i\mathbf{t}^T(\mathbf{X}_i, q_j(\mathbf{X}_i, A_i))) = \frac{1}{n} \sum_{i=1}^n \exp(i\mathbf{t}^T(\mathbf{X}_i, Q_i)),
 \tag{8}$$

where $Q_i = q(\mathbf{X}_i, A_i)$. Székely and Rizzo (2013) and Huling and Mak (2024) demonstrate that the energy distance is some (predefined) norm about the two distribution's characteristic functions. We extend the result from Huling and Mak (2024) and prove that, in the context of MTP, the defined weighted energy distance is a distance between the intended distribution, i.e.,

$$\mathcal{E}(F_{n, \mathbf{w}, \mathbf{X}, A}, F_{n, \mathbf{X}, A}^q) = \int_{\mathbb{R}^p} |\varphi_{n, \mathbf{w}, \mathbf{X}, A}(\mathbf{t}) - \varphi_{n, \mathbf{X}, A}^q(\mathbf{t})|^2 v(\mathbf{t}) d\mathbf{t},
 \tag{9}$$

where \mathbf{w} be a vector of weights such that $\sum_{i=1}^n w_i = n$ and $w_i > 0$ for $\forall i$, $v(\mathbf{t}) = 1/(C_d \|\mathbf{t}\|_{1+d}^{1+d})$, C_d is a constant, $d = p + 1$ is the dimension of the variate (\mathbf{X}, A) , $\varphi_{n,\mathbf{X},A}^q(\mathbf{t})$ is the characteristic function of $F_{n,\mathbf{X},A}^q$, and $\varphi_{n,\mathbf{w},\mathbf{X},A}(\mathbf{t}) = \frac{1}{n} \sum_{i=1}^n w_i e^{i\mathbf{t}^T(\mathbf{X}_i, A_i)}$ is the weighted empirical characteristic function (ECHF) of sample $\{\mathbf{X}_i, A_i\}_{i=1}^n$. Therefore, the weighted energy distance can be used to determine whether the weighted distribution is approaching the MTP-shifted distribution (i.e., the empirical distribution of $(\mathbf{X}, q(\mathbf{X}, A))$).

The following shows that the weighted energy distance converges to the energy distance between the limiting CDF and the MTP-shifted CDF $F_{\mathbf{X},A}^q$.

THEOREM 3.1. *Assume $\lim_{n \rightarrow \infty} \varphi_{n,\mathbf{w},\mathbf{X},A}(\mathbf{t}) = \tilde{\varphi}_{\mathbf{X},A}(\mathbf{t})$ almost surely for $\forall \mathbf{t}$, where $\tilde{\varphi}_{\mathbf{X},A}(\mathbf{t})$ is some integrable characteristic function with associated CDF $\tilde{F}_{\mathbf{X},A}$. Then we have almost surely*

$$(10) \quad \lim_{n \rightarrow \infty} \mathcal{E}(F_{n,\mathbf{w},\mathbf{X},A}, F_{n,\mathbf{X},A}^q) = \mathcal{E}(\tilde{F}_{\mathbf{X},A}, F_{\mathbf{X},A}^q).$$

We now build the connection between the key error term (5) and the energy distance. Suppose that $\mu(\cdot) \in \mathcal{H}$ where \mathcal{H} is the native space induced by the radial kernel $\Phi(\cdot) = -\|\cdot\|_2$ on $(\mathcal{X}, \mathcal{A})$, for any weights \mathbf{w} satisfying $\sum_{i=1}^n w_i = n$, $w_i \geq 0$, then from the results of [Mak and Joseph \(2018\)](#) and [Huling and Mak \(2024\)](#), we have

$$(11) \quad \left[\int_{(\mathcal{A}, \mathcal{X})} \mu(\mathbf{x}, a) d(F_{n,\mathbf{X},A}^q - F_{n,\mathbf{w},\mathbf{X},A})(\mathbf{x}, a) \right]^2 \leq C_\alpha \mathcal{E}(F_{n,\mathbf{X},A}^q, F_{n,\mathbf{w},\mathbf{X},A}),$$

where $0 \leq C_\alpha = \langle \mu, \mu \rangle$ is a constant depending only on $\mu(\mathbf{x}, a)$; here, $\langle \cdot \rangle$ is the inner product in \mathcal{H} .

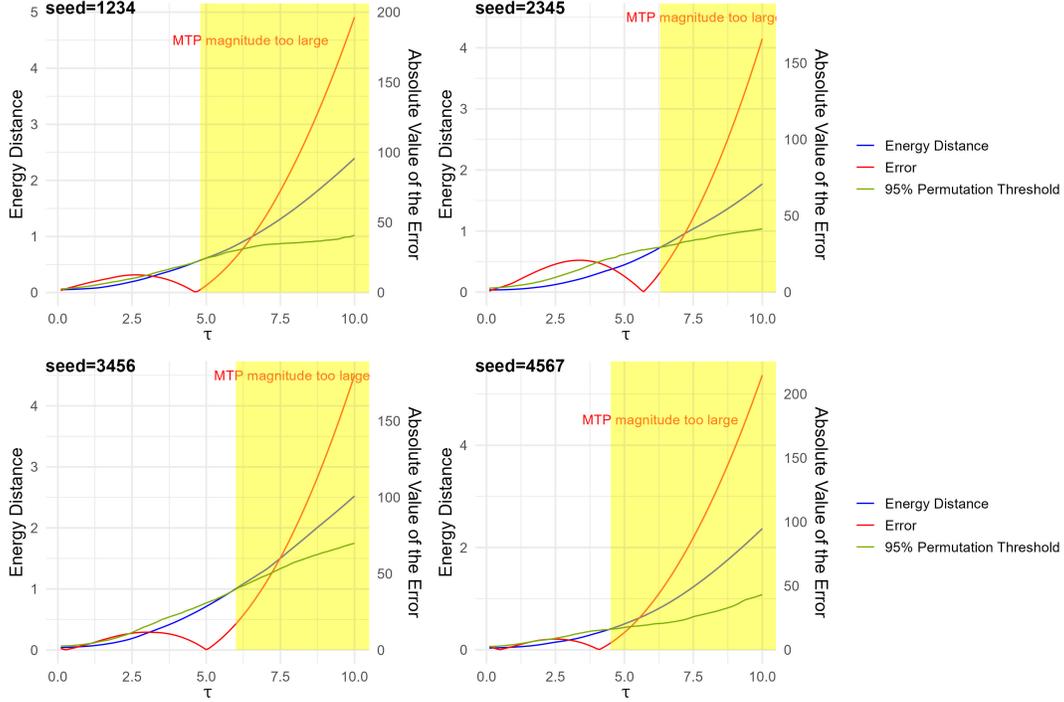
Since (5) is the key component of the bias, this result therefore connects the bias with the distributional balance. It states that under the condition that the conditional mean function $\mu(\mathbf{x}, a)$ belongs to \mathcal{H} , the key component of the bias is bounded by the energy distance. Here we emphasize that \mathcal{H} contains a wide class of functions. For extended discussion and details of the specific form of the native space, readers can refer to [Mak and Joseph \(2018\)](#).

Of import for practical use, we note that the weighted energy distance depends on the scales of both \mathbf{X} and A . Therefore, in order to make the distance unaffected by the measurement scale of \mathbf{X} or A , we suggest standardizing covariates and treatment to have unit variance.

4. Using weighted energy distance for MTPs. In this section, we introduce several important tools that facilitate and enhance the application of MTP techniques. In particular, these tools help provide critical and rigorous guidance on important decisions necessary within the MTP framework. Each of the tools is based on the weighted energy distance due to its connection with estimation bias for MTPs. The first tool centers on determining a reasonable range of treatment policies that can be reliably addressed with a given weighting approach. The second involves assessing arbitrary weighting approaches in their ability to control for measured confounding for a specific dataset and MTP, enabling researchers to determine the most effective weighting method for confounding control.

4.1. A framework for choosing a feasible MTP scale for a given weighting method. When exploring the causal effect of an MTP, the magnitude of the modification then determines how different the treatment under the MTP is compared with assigned treatments in reality. Larger magnitudes or scales of an MTP's modification to treatment values may allow researchers to explore a wider variety of effects, yet more extreme policies differ more

FIG 1. An illustration of our method for selecting a feasible MTP scale based on four simulated datasets. The X-axis represents the MTP scale, with larger values of τ indicating larger shifts in the treatment value. The blue curve shows the energy distance after weighting, while the green curve represents the calculated permutation threshold. The yellow area of the plot indicates values of τ for which the weighted energy distance exceeds the permutation threshold, indicating uncontrolled measured confounding. The estimation error is plotted as a function of τ (red curve). The error tends to increase substantially when employing MTP magnitudes with identified risk of bias.



starkly from reality, making confounding more severe and thus more difficult to control adequately. Thus, the scale of the intervention or modification must be chosen carefully. In this section we develop a formal hypothesis testing framework to provide explicit guidance for this choice.

From the error decomposition in Section 2.2, bias operates only through the term (5) which depends on the difference between the weighted empirical CDF for the observed population and the MTP-shifted empirical CDF for the target population. If the MTP scale is too large, a given set of weights or weighting approach may not properly balance these two populations to have the same distribution, leading to a greater potential for bias in the estimator. To address this issue, we propose a method based on the weighted energy distance to identify a scale of intervention for a given class of MTPs for which a given weighting approach can adequately balance the two populations and thus adequately control for confounding. By doing so, researchers can ensure that the MTP scale is not too large to be estimated well, reducing potential for bias.

In the ideal scenario that the balancing weights perfectly balance the observed population $\{\mathbf{w}_i \cdot (\mathbf{X}_i, A_i)\}_{i=1}^n$ to the MTP-shifted population $\{\mathbf{X}_i, q(\mathbf{X}_i, A_i)\}_{i=1}^n$ for an MTP, measured confounding is completely controlled for. Thus, denoting $F_{\mathbf{w}, \mathbf{X}, A}(\mathbf{x}, a) = \mathbb{E}[F_{n, \mathbf{w}, \mathbf{X}, A}(\mathbf{x}, a)]$ our method therefore focuses on the following null and alternative hypotheses:

$$H_0 : F_{\mathbf{w}, \mathbf{X}, A}(\mathbf{x}, a) = F_{\mathbf{X}, A}^q(\mathbf{x}, a) \quad \forall (\mathbf{x}, a) \in (\mathcal{X}, \mathcal{A}) \quad \text{v.s.} \quad H_\alpha : F_{\mathbf{w}, \mathbf{X}, A}(\mathbf{x}, a) \neq F_{\mathbf{X}, A}^q(\mathbf{x}, a).$$

Rejecting H_0 indicates that the given balancing weights do not perfectly balance the observed population’s joint distribution of covariates and treatment to that of the counterfactual MTP-shifted population; thus, the alternative hypothesis represents scenarios where confounding has not been adequately controlled, because the given weights are inadequate for confounding control for the given MTP. Some weighting methods may work well for smaller MTP magnitudes, but fail when the MTP’s modification is more extreme. Carrying out a test to assess the above hypotheses is especially challenging, as unlike [Rizzo and Székely \(2016\)](#) where they use the energy statistic to test the equal distribution of two *independent* samples, in our scenario, the observed sample and the MTP-shifted sample are pair-to-pair correlated (since they are derived from the same sample points). To handle this issue, we randomly partition the sample data $\{\mathbf{X}_i, A_i\}_{i=1}^n$ into two subsamples $\{\mathbf{X}_{I_k}, A_{I_k}\}_{k=1}^{n/2}$ and $\{\mathbf{X}_{J_k}, A_{J_k}\}_{k=1}^{n/2}$ where I and J are complementary index sets with equal size, without loss of generality. We use the first subsample to form the weighted observed subsample $w_{I_k} \cdot \{\mathbf{X}_{I_k}, A_{I_k}\}_{k=1}^{n/2}$ and the second to form the MTP-shifted subsample $\{\mathbf{X}_{J_k}, q(\mathbf{X}_{J_k}, A_{J_k})\}_{k=1}^{n/2}$. Due to the random sampling, we have that $\int_{(\mathcal{X}, \mathcal{A})} F_{n, w, \mathbf{X}, A}(\mathbf{x}, a) dF_{\mathbf{X}, A} = \int_{(\mathcal{X}, \mathcal{A})} F_{n/2, w, \mathbf{X}_I, A_I}(\mathbf{x}, a) dF_{\mathbf{X}, A}$ and $F_{\mathbf{X}, A}^q = \int_{(\mathcal{X}, \mathcal{A})} F_{n, \mathbf{X}, A}^q(\mathbf{x}, a) dF_{\mathbf{X}, A} = \int_{(\mathcal{X}, \mathcal{A})} F_{n/2, \mathbf{X}_J, A_J}^q(\mathbf{x}, a) dF_{\mathbf{X}, A}$. We can then use the permutation test described in [Lehmann, Romano and Casella \(1986\)](#) Section 15.2 using the weighted observed subsample and the MTP-shifted subsample. Since for any $k_1 \neq k_2$, $(\mathbf{X}_{I_{k_1}}, A_{I_{k_1}})$ and $(\mathbf{X}_{I_{k_2}}, A_{I_{k_2}})$ are independent, our partition ensures that the two samples are mutually independent, thus satisfying the “randomization hypothesis” required in Theorem 15.2.1 from [Lehmann, Romano and Casella \(1986\)](#) under permutation.

We use the weighted energy distance between the two subsamples as the test statistic $T = \mathcal{E}(F_{n/2, w_I, \mathbf{X}_I, A_I}, F_{n/2, \mathbf{X}_J, A_J}^q)$. ([Székely, Rizzo and Bakirov, 2007](#)) have shown that, for two independent samples, this test statistic converges in distribution to a quadratic form of independent standard normal random variables under the null hypothesis of equal distribution and tends to infinity stochastically under the alternative hypothesis. From the Theorem 15.2.1 in [Lehmann, Romano and Casella \(1986\)](#), the permutation test will have level α under H_0 . The critical value for the hypothesis testing is determined by the permutation randomization described in the following algorithm:

Determining whether the MTP scale is feasible for a given weighting approach

Input: Data $\{\mathbf{X}_i, A_i\}_{i=1}^n$, MTP q , and balancing weights w .

Randomly partition the data into two parts indexed as I and J and calculate the test statistic:

$$T = \mathcal{E}(F_{n/2, w_I, \mathbf{X}_I, A_I}, F_{n/2, \mathbf{X}_J, A_J}^q)$$

Combine the data as $\{\tilde{w}_i, \tilde{\mathbf{X}}_i, \tilde{A}_i\}_{i=1}^n = (\{w_{I_k}, \mathbf{X}_{I_k}, A_{I_k}\}_{k=1}^{n/2}, \{w_{J_k}, \mathbf{X}_{J_k}, q(\mathbf{X}_{J_k}, A_{J_k})\}_{k=1}^{n/2})$ where $w_{J_k} = 1$ for $k = 1, \dots, n/2$.

For $b = 1$ **to** B -th iteration:

Randomly partition the combined data into two parts indexed as I^b and J^b

Calculate the test statistic: $T_b = \mathcal{E}(F_{n/2, \tilde{w}_{I^b}, \tilde{\mathbf{X}}_{I^b}, \tilde{A}_{I^b}}, F_{n/2, \tilde{w}_{J^b}, \tilde{\mathbf{X}}_{J^b}, \tilde{A}_{J^b}})$

End For

Calculate the critical value T^p as the upper 0.95th quantile of the permutation sample $\{T_b\}_{b=1}^B$.

Reject H_0 if $T^p < T$.

We illustrate this new approach through a simulated example using our novel energy balancing weights (described in more detail in Sections 5 and 6). Simulated data are generated using the first condition of our simulation experiments described in Section 7 with $n = 200$ and dimensionality $p = 10$. We adopt the shift function $q(\mathbf{x}, a) = \tau \tilde{q}(\mathbf{x}, a)$ where $\tilde{q}(\mathbf{x}, a)$ is the shift function we used in our simulation study and τ ranges from $[0, 10]$ in increments of

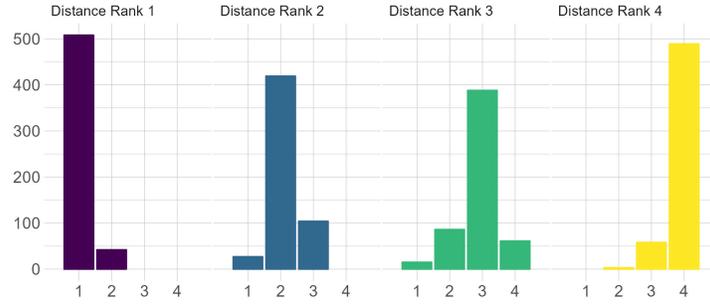


FIG 2. The histogram depicts the association between the rank of estimation error and the rank of weighted energy distance using the same data as in Figure 2. The x-axis of each plot is the rank of each method in terms of performance. The figure shows that, in most cases, the method with the lowest weighted energy distance yields the best performance.

0.1 and controls the magnitude of our MTP. Larger values of τ represent the larger discrepancies from the original treatment. For each τ , we calculate the weighted energy distance as the test statistic, the error of the energy balancing weights estimator, and the 95% upper quantiles of the permutation sample. Figure 1 displays four general scenarios with different random seeds. The energy distance for the energy balancing weights increases with a larger MTP scale τ ; along with this increase, the absolute error of the estimator also increases mostly monotonically. The vertical dashed line indicates the intersection of the energy distance (the test statistic) and the permutation threshold of the null distribution, to the right of which the values of τ are large enough that there is evidence that the weights do not adequately control for confounding. We can see that for τ values to the left of the vertical lines, the estimation error is generally small, indicating that this procedure is a reasonable criterion for choosing a feasible MTP scale.

We also conduct a full simulation study to verify correct type I error rate control under H_0 and further explore alternative permutation strategies for the test. These results are presented in the Supplementary Material Section 4.1.

4.2. Evaluation of arbitrary weights for a given MTP and dataset. In the previous section, we evaluate different MTP scales with a fixed weighting method. Similarly, we can also use the weighted energy distance to evaluate the performance of different weighting methods for a fixed MTP. Since smaller weighted energy distances indicate a better balance between the weighted sample and the shifted sample, a natural criterion is to select the method that yields the smallest weighted energy distance. We demonstrate this approach to weight selection using the simulation results with sample sizes $n = 800$ and covariate dimensionality $p = 80$ under the first simulation scenario described in Section 7.

Figure 2 displays the distribution of the rank of the absolute estimation error for a given dataset compared with the rank of the weighted energy distance for the four weighting methods. From this, the weighting method with the smallest energy distance (rank 1) tends to perform best in most cases. In the Supplementary Material Section 4.2, we further explore the validity of this method under broader conditions.

5. Energy balancing weights for MTP estimation. As demonstrated in the error decomposition in Section 2.2, weighted distributional imbalance is a critical component of the

estimation bias of weighted estimators for MTPs. Further, we showed in (9) and Theorem 3.1 that the weighted energy distance characterizes weighted distributional imbalance. In this section we propose a new set of weights designed as the optimizer of this measure. By minimizing the distance between the weighted empirical distribution and target distribution, our proposed weights, which we call energy balancing weights (EBWs), control for confounding in a robust and flexible manner. Specifically, these weights minimize the energy distance between the observed sample and the MTP-shifted sample. Thus, the energy balancing weights balance the weighted empirical distribution to the MTP-shifted empirical distribution.

Extensively highlighted in the weighting literature, extreme weights can inflate the variance of the estimator (Li, Morgan and Zaslavsky, 2018; Chattopadhyay, Hase and Zubizarreta, 2020). To ensure root- n consistency of our energy balancing estimator, it is necessary to impose a restriction to avoid extreme weights. One approach is to introduce an additional penalization term for the sum of squared weights, resulting in penalized energy balancing weights. In this section, we introduce the penalized energy balancing weights w_n^p . Under mild conditions, we show these weights result in root- n consistency in estimating the mean potential outcome under an MTP.

5.1. Penalized energy balancing estimator. The penalized energy balancing weights account for the magnitude of the weights in addition to the weighted energy distance. Specifically, we define the penalized energy balancing weights w_n^p with a user-specified parameter $\lambda > 0$ as follows:

$$(12) \quad w_n^p \in \arg \min_{\mathbf{w}=(w_1, \dots, w_n)} \mathcal{E}(F_{n, \mathbf{w}, \mathbf{X}, A}, F_{n, \mathbf{X}, A}^q) + \frac{\lambda}{n^2} \sum_{i=1}^n w_i^2 \text{ s.t. } \sum_{i=1}^n w_i = n, w_i \geq 0.$$

In the Supplementary Material Section 5.2, we showed that (12) is equivalent to a quadratic optimization problem that can be solved efficiently by the existing algorithms. The corresponding penalized energy balancing estimator is then $\hat{\mu}_{w_n^p}^q = n^{-1} \sum_{i=1}^n w_{i,n}^p Y_i$. Analogous to Huling and Mak (2024), we prove the following properties of the energy balancing weights. Theorem 5.1 shows that the penalized energy balancing weights make the weighted empirical CDF of the observed data converge to the true MTP-shifted CDF.

THEOREM 5.1. *Assume the assumptions in the previous theorems hold. Let w_n^p be the penalized energy balancing weights. Then, we have $\lim_{n \rightarrow \infty} F_{n, w_n^p, \mathbf{X}, A}(\mathbf{x}, a) = F_{\mathbf{X}, A}^q(\mathbf{x}, a)$ almost surely for every continuity point $(\mathbf{x}, a) \in \mathcal{X}$. Furthermore, the following holds almost surely*

$$(13) \quad \lim_{n \rightarrow \infty} \mathcal{E}(F_{n, w_n^p, \mathbf{X}, A}, F_{n, \mathbf{X}, A}^q) = 0.$$

We now show that the penalized energy balancing estimator achieves root- n consistency.

THEOREM 5.2. *Assume the conditions in Theorem 2. Let \mathcal{H} be the native space induced by the radial kernel $\Phi(\cdot) = -\|\cdot\|_2$ on $(\mathcal{X}, \mathcal{A})$. Suppose the following mild conditions hold:*

- **CP-1:** $\mu(\cdot, \cdot) \in \mathcal{H}$
- **CP-2:** $\text{Var}[\mu(\mathbf{X}, A)] < \infty$
- **CP-3:** $\text{Var}[Y|X = \mathbf{x}, A = a]$ are bounded over $(\mathbf{x}, a) \in (\mathcal{X}, \mathcal{A})$.
- **CP-4:** $\mathbb{E}[g^2(\mathbf{W}, \mathbf{W}', \mathbf{W}'', \mathbf{W}''')] \leq \infty$ where $\mathbf{W}, \mathbf{W}', \mathbf{W}'', \mathbf{W}''' \stackrel{i.i.d.}{\sim} F_{\mathbf{X}, A}^q$ be a vector of (\mathbf{X}, A) and, with $h(\mathbf{w}) = \frac{f_{\mathbf{X}, A}^q(\mathbf{x}, a)}{f_{\mathbf{X}, A}(\mathbf{x}, a)}$, the kernel function $g(\cdot)$ is defined as:

$$g(\mathbf{w}, \mathbf{w}', \mathbf{w}'', \mathbf{w}''') = h(\mathbf{w})\|\mathbf{w} - \mathbf{w}''\|_2 + h(\mathbf{w}')\|\mathbf{w}' - \mathbf{w}'''\|_2 - h(\mathbf{w})h(\mathbf{w}')\|\mathbf{w} - \mathbf{w}'\|_2 - \|\mathbf{w}''' - \mathbf{w}''\|_2$$

Then, the proposed penalized EBW estimator $\hat{\mu}_{q, w_n^p}$ is root- n consistent in that

$$(14) \quad \sqrt{\mathbb{E}_{\mathbf{X}, A, Y} [(\hat{\mu}_{w_n^p}^q - \mu^q)^2]} = O_p(n^{-1/2}).$$

Theorem 5.2 shows that the penalized energy balancing weights yield a root- n consistent estimate of μ^q under mild conditions on the data-generating process even though the energy balancing weights are not shown to be consistent for the true density ratio weights in any sense.

We briefly comment on the conditions required for Theorem 5.2. Condition **CP-1** requires that the outcome regression function be contained in the native space \mathcal{H} . As demonstrated in Huling and Mak (2024), this condition is a smoothness assumption on the true outcome regression function $\mu(\mathbf{x}, a)$. **CP-2** requires that the outcome regression function have finite variance and **CP-3** requires that the conditional variance function be bounded. These two conditions are both mild and fairly weak in practice. **CP-4** is a requirement that certain moments of the covariates and treatments be bounded.

The selection of λ can be important in practice, as it determines the trade-off between bias and variance. Larger λ values result in estimators with reduced variance but increased bias, due to greater imbalance between the two distributions after weighting. In this context, bias cannot ever be exactly characterized, so choosing an optimal value for λ is a major challenge in practice. However, as bias is related to the imbalance of covariates, we can at least characterize what range of values of λ still result in acceptable covariate balance. Based on the method proposed in Section 4, which tests whether the weights effectively balance the two distributions, one approach for choosing λ is to select the largest λ for which the permutation test indicates no significant imbalance remains after weighting.

6. Augmented energy balancing estimator. Augmented estimator is a widely-used approach in the causal inference literature that combine balancing weights and an estimated outcome model to improve the estimation of a causal effect, potentially increasing efficiency and reducing sensitivity to model misspecification. In this context, we propose an augmented version of the energy balancing estimator that incorporates an estimated outcome model. Because the energy balancing weights were shown to yield consistent estimates of μ^q without the need for an outcome regression model, this allows the analyst to utilize an outcome regression model primarily for the purpose of variance reduction.

Let $\hat{\mu}(\mathbf{x}, a)$ be an estimate of the outcome regression function $\mu(\mathbf{x}, a)$. The augmented estimator is constructed by subtracting $\int_{\mathcal{X}} \int_{\mathcal{A}} \hat{\mu}(\mathbf{x}, a) d(F_{n, \mathbf{X}, A}^q - F_{n, w_n^p, \mathbf{X}, A})(\mathbf{x}, a)$ from the weighted estimator $\hat{\mu}_{w_n^p}^q = \sum_{i=1}^n w_i^p Y_i$. The resulting augmented energy balancing estimator is

$$(15) \quad \begin{aligned} \hat{\mu}_{AG}^q &= \hat{\mu}_{w_n^p}^q - \int_{\mathcal{X}} \int_{\mathcal{A}} \hat{\mu}(\mathbf{x}, a) d(F_{n, \mathbf{X}, A}^q - F_{n, w_n^p, \mathbf{X}, A})(\mathbf{x}, a) \\ &= \frac{1}{n} \sum_{i=1}^n w_i^p (Y_i - \hat{\mu}(\mathbf{X}_i, A_i)) + \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^{J(\mathbf{X}_i)} I_{j, \mathbf{X}_i}(A_i) \hat{\mu}(\mathbf{X}_i, q_j(\mathbf{X}_i, A_i)). \end{aligned}$$

In the second form, the weighted residuals can be view as a bias correction term for the outcome regression-based estimate of the MTP effect. We show its asymptotic normality and propose a statistically-valid and computationally-efficient multiplier bootstrap for inference.

6.1. Asymptotic normality. In this section we demonstrate the asymptotic distribution of augmented energy balancing estimator $\hat{\mu}_{AG}^q$. The required conditions for asymptotic normality are:

- **CA-1:** $\hat{\mu}(\mathbf{x}, a)$ is such that $\hat{\mu} - \mu \in \mathcal{H}$.
- **CA-2:** w_n^e satisfy $1 \leq w_i^e \leq BN^{1/3}$ for some constant B where N is the total sample size.
- **CA-3:** $\max_i E(|\epsilon_i|^3) < \infty$.
- **CA-4:** $\hat{\mu}(\mathbf{x}, a)$ is a consistent estimator of $\mu(\mathbf{x}, a)$.

Condition **CA-1** requires that the outcome regression model minus the true outcome regression model be contained in the native space \mathcal{H} . This condition somewhat limits the complexity of the error function of the estimated outcome regression model. Condition **CA-2** has been used before in the literature; see [Athey, Imbens and Wager \(2018\)](#); [Huling, Greifer and Chen \(2023\)](#). This condition can be imposed directly in the weight optimization procedure without changing the empirical performance of the weights or any of the asymptotic results of the weights, as in [Huling, Greifer and Chen \(2023\)](#). It should be noted that, if we use the penalized energy balancing weights, this condition will no longer be needed as adding the penalty term in the optimization problem helps to sufficiently restrict the variability of the weights. Condition **CA-3** indicates that the error term has finite third moments, and condition **CA-4** requires consistency of the outcome regression model; **CA-4** is only required for studying the asymptotic distribution of $\hat{\mu}_{AG}^q$. With these conditions, we have the following theorem:

THEOREM 6.1. *Under Conditions CA1-CA4, let $J_n = n^{1/2}(\hat{\mu}_{AG}^q - \mu^q)$ and $J_n^* = [\text{Var}(\mu(\mathbf{X}, A))]^{1/2}F + \sigma n^{-1/2} \sum_{j=1}^n w_j^e G_j$ where F, G_1, \dots, G_n are independently and identically distributed standard normal random variables independent of $\mathbf{X}_1, \dots, \mathbf{X}_n, A_1, \dots, A_n$ and $\epsilon_1, \dots, \epsilon_n$. Let ψ_n and ψ_n^* be the corresponding characteristic functions of J_n and J_n^* . Then as $n \rightarrow \infty$, $|\psi_n(t) - \psi_n^*(t)| \rightarrow 0, t \in \mathbb{R}$ where $\psi_n^*(t)$ is twice differentiable, and*

$$(16) \quad \limsup_n \text{Var}(J_n) \leq \text{Var}(\mu(\mathbf{X}, Q)) + \sigma^2 V,$$

where $\mu(\mathbf{X}, Q) = \sum_{j=1}^{J(\mathbf{X})} I_{j, \mathbf{X}}(A) \mu(\mathbf{X}, q_j(\mathbf{X}, A))$ is the conditional mean under the MTP and $V = \mathbf{E} \left(\frac{f_{\mathbf{X}, A}^q(\mathbf{X}, A)}{f_{\mathbf{X}, A}(\mathbf{X}, A)} \right)^2 + D$ is the variance of the Radon-Nikodym weights plus a positive number.

Theorem 6.1 demonstrates that the energy balancing weights, when paired with a consistent outcome regression model, achieve asymptotic normality and may be efficient, however, an exploration of the conditions under which efficiency is guaranteed requires additional exploration.

6.2. Inference. The asymptotic normality of our augmented estimator $\hat{\mu}_{AG}^q$ allows us to use a standard bootstrap method to estimate its uncertainty and conduct statistical inference for the estimator. However, in practice, the standard nonparametric bootstrap method has some limitations. First, it can be computationally intensive because both the penalized energy balancing weights and the outcome regression estimator must be re-computed during each bootstrap iteration $r = 1, \dots, R$. This can be especially time-consuming for large values of R (e.g., $R = 10000$).

We thus propose a multiplier bootstrap method originally presented by [Wu \(1986\)](#) that allows for quick computation. Similar to [Matsouaka, Liu and Zhou \(2023\)](#) (which they name a wild bootstrap), the method perturbs the influence function of the estimator to estimate its variance. The influence function for our augmented estimator is

$$\varphi_i = \frac{f_{\mathbf{X}, A}^q(\mathbf{X}, A)}{f_{\mathbf{X}, A}(\mathbf{X}, A)} (Y_i - \mu(\mathbf{X}_i, A_i)) + \sum_{i=1}^{J(\mathbf{X}_i)} I_{j, \mathbf{X}_i}(A_i) \mu(\mathbf{X}_i, q_j(\mathbf{X}_i, A_i)) - \mu^q.$$

We estimate the influence function by plugging in the estimators $\hat{\mu}^q$, $\hat{\mu}(\mathbf{X}, A)$, and w_i^p for their population counterparts:

$$(17) \quad \hat{\varphi}_i = w_i^p (Y_i - \hat{\mu}(\mathbf{X}_i, A_i)) + \sum_{i=1}^{J(\mathbf{X}_i)} I_{j, \mathbf{X}_i}(A_i) \hat{\mu}(\mathbf{X}_i, q_j(\mathbf{X}_i, A_i)) - \hat{\mu}^q.$$

Even though w_i^p does not necessarily estimate the true inverse propensity score, we show that the following procedure is still valid. Then, the multiplier bootstrap estimator $\hat{\Sigma}$ can be constructed as displayed in Algorithm 6.2. A 95% Wald-type confidence interval of $\hat{\mu}_{AG}^q$ can be calculated as $\hat{\mu}_{AG}^q \pm 1.96 \times n^{-1/2} \hat{\Sigma}^{1/2}$.

Multiplier bootstrap

For $r = 1$ **to** $r = R$

Generate the random variate $\boldsymbol{\xi} = (\xi_1, \dots, \xi_n)$ with $\xi_i \sim \exp(1)$ for $i = 1, \dots, n$.

Compute $\hat{q}_r^q = \frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i \hat{\varphi}_i$.

Calculate $\hat{\Sigma}^{1/2} = \text{IQR}(\hat{q}^q) / (z_{0.75} - z_{0.25})$ where $\text{IQR}(\hat{q}^q)$ is the interquartile range of $\{\hat{q}_r^q\}_{r=1}^R$ and $z_{0.75} - z_{0.25} = 1.349$ be the interquartile range of standard normal distribution.

Here, we prove the validity of the proposed multiplier bootstrap in the following theorem.

THEOREM 6.2. *Under Conditions CA1-CA4, and \hat{q}_r^q defined in Algorithm 1, we have that $\hat{q}_r^q = \frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i \hat{\varphi}_i \rightarrow n^{1/2}(\hat{\mu}_{AG}^q - \mu^q)$ as $n \rightarrow \infty$.*

Theorem 6.2 shows that one can estimate the asymptotic distribution of $\hat{\mu}_{AG}^q$ using the multiplier bootstrap procedure in Algorithm 1. An advantage of the multiplier bootstrap procedure over the standard bootstrap is that the weights and outcome regression model only need to be estimated once, not R times, resulting in a substantial computational speedup.

7. Simulation studies. In this section, we study the empirical properties of the proposed energy balancing weights and further to assess their relative performance to other commonly used estimators for MTPs. The aim of the simulation study is to evaluate the performance of our proposed estimators under different simulation conditions.

7.1. Data generating mechanisms. We have two main simulation conditions: a moderately complex data-generating mechanism (simulation study # 1) and a highly complex (simulation study # 2) data-generating mechanism. For each simulation setting, we vary the following settings: the sample size ($n = 100, 200, 400, 800$), the number of covariates ($p = 10, 20, 40, 80$) for each individual, and the treatment type. This allows us to inspect the joint effects of dimensionality and increasing sample size on the empirical performance of all methods. We replicate all simulation experiments under two different conditions on the treatment variable: one condition where the treatment is continuously-distributed, and the other where the treatment has discrete, but countably infinite, support. Here, we briefly illustrate the data-generating mechanism. Readers are referred to Section 6.1 of the Supplementary Material for a detailed description of the data-generating mechanism. 1) The covariates are generated randomly following either a uniform distribution or a binomial distribution. 2) The treatment mean for a participant follows a cubic function of the covariates. 3) The moderately complex simulation condition assumes the outcome is a function of a quadratic term of the

treatment times a cubic term of the covariates. The highly complex simulation condition further assumes interaction terms of the covariates. 4) The MTP function q shifts the treatment more aggressively when the observed treatment is small. We emphasize that the estimation error (i.e., the observed sample mean effect minus the true MTP effect) is designed to be relatively stable with the increase of the number of covariates p , so that the complexity of the data-generating processes does not explode with dimension.

7.2. Estimands and estimators under evaluation. In each simulation replication, the estimand of interest is the mean potential outcome under the MTP. We evaluate three proposed energy balancing weight methods, including the penalized energy balancing weights described in Section 5.1, an unpenalized version of the energy balancing weights, and a kernelized energy balancing weights approach using the Gaussian kernel (see Supplementary Material Section 2.2). [The penalty factors \$\lambda\$ for the penalized energy balancing weights are set to be \$\lambda = 1\$ for all simulation scenarios.](#) Various alternative estimators for balancing weights are considered, such as the naive, unadjusted method that assigns equal weights to all subjects, density ratio weights (IPW) with generalized propensity score estimated using a Poisson density (for the discrete intervention settings), and the classification method proposed by (Díaz et al. (2021) Section 5.4) with classification models estimated by either 1) logistic regression or 2) a random forest. The corresponding augmented estimators for each weighting method are implemented using the same outcome model, the ensemble method SuperLearner (Van der Laan, Polley and Hubbard, 2007), which incorporates the lasso regularized generalized linear models (`SL.glmnet`). Also included is a targeted minimum-loss based estimator (TMLE) method (Díaz et al., 2021) with all nuisance parameters estimated using SuperLearner.

7.3. Simulation results. In every simulation scenario, the true causal effect is determined using Monte Carlo with a sample size of 100,000. For each simulation setting, we repeat the experiment independently 1,000 times, applying each comparator method, allowing for the calculation of the mean squared error (MSE), bias, and coverage rate for 95% confidence intervals. The coverage rate is calculated based on Wald-type intervals with SE for each estimator estimated by the non-parametric bootstrap with 100 replications. Additionally, we also use the multiplier bootstrap method proposed in Section 6.2 for the augmented energy balancing estimator. The simulation results are displayed in Figures 3 and 4. Since the results of the two simulation studies are similar, we only display the MSE and coverage rate for the 95% confidence interval for the first simulation study. Readers can refer to Supplementary Material Section 6.2 for the results of the highly complex setting.

In all simulation conditions, our three energy balancing methods (energy balancing, penalized energy balancing, and kernel energy balancing method) consistently outperform other methods in terms of both the bias and the coverage rate, indicating their robustness and ability to handle various situations. This pattern holds both for pure weighting (not augmented) estimators and for augmented estimators. The penalized energy balancing weights exhibit slightly worse performance in terms of bias compared to the energy balancing weights, which is probably due to the additional penalty term in the estimation. The kernel energy balancing weights (using the MMD described in the Supplementary Material Section 2.2) have the smallest bias in most cases when the sample size is relatively small; however, their performance varies significantly among different conditions and the bias does not improve as much with the increase of the sample sizes. This unstable performance may be due to the choice of bandwidth parameters in the Gaussian kernel function, which is the median heuristic that is known to not necessarily perform well in all scenarios. Therefore, although we believe the kernel energy balancing method has the potential to be very flexible and perform well

in practice, its performance may depend critically on the choice of tuning parameters. Consequently, we recommend using penalized energy balancing methods that are more stable across different conditions.

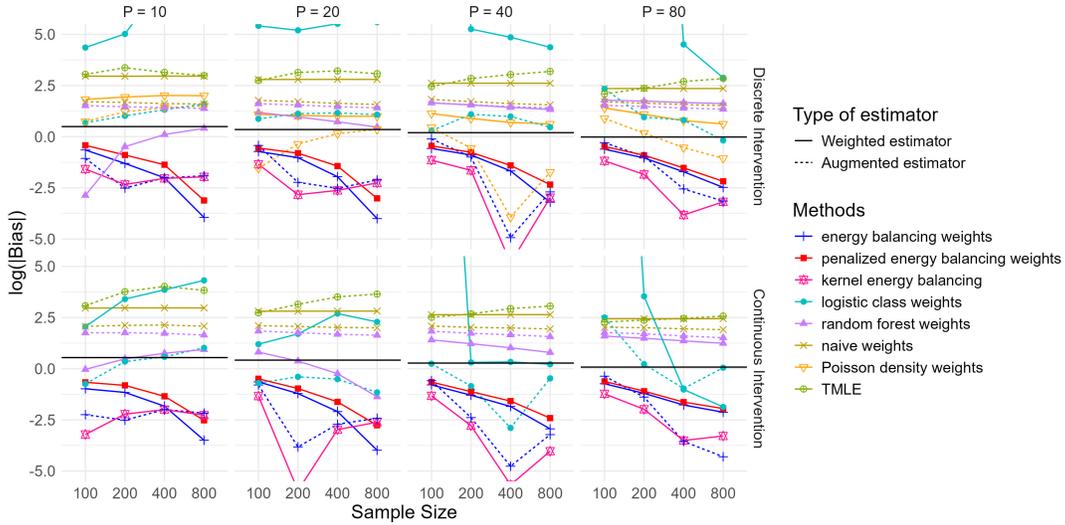


FIG 3. Simulation results in terms of the logarithm of the absolute value of bias across different sample sizes, type of treatment/intervention variable, and the dimensionality of covariates. The data-generating mechanism is moderately complex (simulation # 1). Balancing methods are displayed in different colors and shapes of points. Weighted estimators are displayed in solid lines and the augmented estimators are in dashed lines.

The augmented energy balancing estimator generally has a smaller bias with a similar MSE (See Supplementary Material Figure 5) in most of the simulation conditions, which indicates the benefit of the inclusion of the outcome model. Although using the same outcome model, the augmented estimators with other balancing weights perform significantly worse than the energy balancing estimators. The coverage rate results (Figure 4) align with our observations for bias and MSE. All three energy balancing methods exhibit good coverage rate performance. Owing to its small bias, the coverage rate for the kernel energy balancing method approaches 95% most closely.

8. Analysis of Mechanical Ventilation Data. Critically ill patients often require mechanical ventilation to support their breathing. The ventilator delivers a controlled amount of air (tidal volume) at a specific pressure and rate to ensure adequate gas exchange in the lungs. In order to overcome the resistance of the airway and expand the thorax wall, the ventilator needs to transfer a certain amount of energy to the patient’s respiratory system (Neto et al., 2018; Neto, Amato and Schultz, 2016; Cressoni et al., 2016). Mechanical power (MP) is a comprehensive parameter that quantifies the amount of energy transferred to the patient’s respiratory system during each breath. It takes into account tidal volume, respiratory rate, peak airway pressure, and positive end-expiratory pressure (PEEP) (Gattinoni et al., 2016). Excessive ventilatory settings can result in ventilator-induced lung injury (VILI), which can worsen the patient’s clinical condition and increase the risk of mortality (Neto, Amato and Schultz, 2016; Cressoni et al., 2016). Previous research has suggested that factors like high tidal volume, high airway pressures, and high respiratory rates could be associated with an

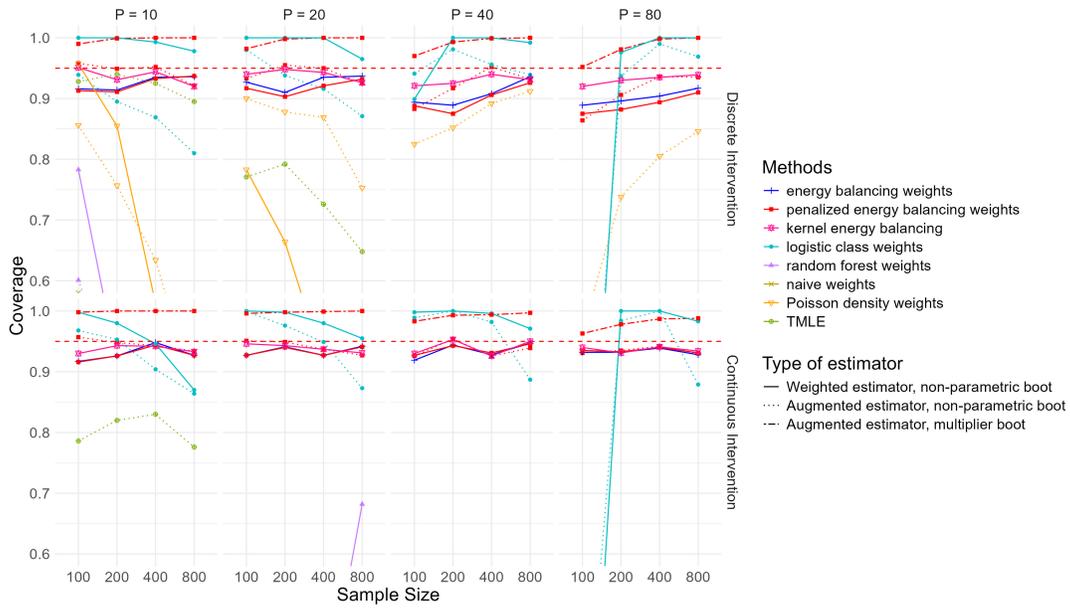


FIG 4. Simulation result for the 95% coverage rate across different sample sizes and dimensionality of covariates. The data-generating mechanism is moderately complex (simulation # 1). The red dashed line indicates nominal 95% coverage. Different weighting methods are displayed in different colors and shapes of points. Weighted estimators are displayed in solid lines and the augmented estimators are in dashed lines. The multiplier bootstrap method is displayed in a larger size of the red points.

increased risk of lung injury and poor outcomes (Nieman et al., 2016). However, these individual factors do not capture the overall impact of ventilation settings on the lungs. Mechanical power can provide a more comprehensive assessment of the risk associated with different ventilatory settings.

A recent study (Neto et al., 2018) examined data from the high-resolution database, the Medical Information Mart for Intensive Care (MIMIC-III) (Johnson et al., 2016). The MIMIC-III data contain information on critically ill patients who required mechanical ventilation in the intensive care unit (ICU). All included patients received invasive ventilation for at least 48 consecutive hours. Time-varying variables were collected at 8-hour intervals for the entire 48-hour period. (Neto et al., 2018) defined the exposure as the mean between the highest and lowest value of MP in the second 24 hours and concluded that high MP is independently associated with higher in-hospital mortality. In this case study, we reanalyze data from Neto et al. (2018) to explore the causal relationship between a postulated MTP of MP and in-hospital mortality.

We excluded patients with zero MP in Joules/min or extreme MP values (MP greater than 150 Joules/min) from the analysis. The exposure of interest is the mean value between the highest and lowest MP during the second 24 hours. The potential confounding variables we aim to balance are measured during or prior to the first 24 hours. As in Neto et al. (2018), our primary focus is on the in-hospital mortality of the included participants. The dataset contains a total of 5,011 participants with 97 identified covariates; among these covariates are many of the important factors such as the PaO₂/FiO₂ Ratio that are used in guidelines for determining mechanical ventilator settings. We use the following MTP to explore the causal effect of MTP that decreases the MP value based on its original value. This MTP reduces MP for individuals with high MP more aggressively, as earlier studies have indicated potential harms of high MP. The shifted MP value $q(x, a)$ is defined as $q(x, a) = a$ if $0 < a < 5$,

$q(\mathbf{x}, a) = a - 5\tau$ if $5 < a < 10$, $q(\mathbf{x}, a) = a - 10\tau$ if $10 < a < 20$, $q(\mathbf{x}, a) = a - 15\tau$ if $20 < a < 40$, $q(\mathbf{x}, a) = a - 30\tau$ if $40 < a$. Here a is the original MP value and τ is the parameter that controls the magnitude of the MTP (a larger value of τ results in a more dramatic decrease of MP).

With the τ ranging from 0 to 1 for the MTPs, our analysis results are displayed in Figure 5. We adopt the unpenalized version of the energy balancing estimator. To identify a range of the shift magnitude τ for which confounding control via our energy balancing weights is feasible, we calculate the sampling variation of the energy distance under the corresponding MTP as described in Section 4.1. The upper 95% percentiles of the permutation sample are plotted with smoothed lines in the figure. Again, this value represents the intrinsic tails of variation of the energy distance when the two distributions are identical (i.e., the population is perfectly balanced). If the actual energy distance after balancing is larger than these upper bounds, the validity of the hypothesis that the population is balanced after weighting becomes questionable. The energy distance between the weighted sample and the target sample under the MTP is displayed as the blue curve for each value of τ . From the results, all MTPs with τ range between $[0, 1]$ have energy distances below the 95% permutation threshold, indicating that the shifts are within a reasonable range and have a low risk of bias due to measured confounding after weighting.

The red curve in Figure 5 represents the estimated potential in-hospital mortality rate under the MTP. The pink curves around it are the pointwise upper and lower bound for the 95% confidence interval of the difference between the observed mortality and the estimated mortality under the MTP using the non-parametric bootstrap. To make the confidence interval more visually apparent, we have centered it around the estimated mortality under the MTP. From Figure 5, we can see that the MTP results in a significant reduction in mortality even when the shift of treatment is small. When the MTP shift increases, it is more challenging to balance the distributions of covariates, which makes the confidence interval wider. However, the overall trend still shows that the mortality rate decreases with increases in the MTP shift. From Figure 5, as τ ranges between 0 to 1, the energy distances after balancing are consistently lower than the permutation thresholds, suggesting that measured confounding has been adequately controlled after weighting. This strengthens the validity of the result that a lower MP than is used in practice would likely decrease in-hospital mortality.

Compared to the results from (Neto et al., 2018; Hong et al., 2021), we adopted an MTP analysis, strengthening the evidence about the potential harms of high MP due to the weakened assumptions under which our MTP estimators operate. Our set of tools provides a more convincing conclusion in terms of 1) the permutation test confirms that the measured confounding is not an issue after covariate balancing, and 2) the causal effect is significant even when the shift of MP is small, which further confirms the overall trend in our results.

9. Discussion. Estimating causal effects of treatments from observational data is particularly challenging due to confounding. This issue is further compounded when the treatment takes continuous values. An MTP analysis is well-suited for this scenario and defines a hypothetical intervention on a treatment variable based on the baseline characteristics and the observed treatment value. However, the choice of the hypothetical intervention in an MTP is still subjective and it faces the dilemma that, 1) if the definition of MTP is too conservative (i.e., the shift of the intervention is very small), the causal effect of the MTP may be too small to have material clinical implications and on the other hand 2) if the definition of MTP is too aggressive, then the observed population and the population under MTP will be intrinsically different, which makes confounding a major issue again, negating many of the conceptual and practical benefits of MTPs. There is a lack of methods that can measure the magnitude of the potential for measured confounding and help to define a reasonable MTP that can be reliably estimated from the data at hand.

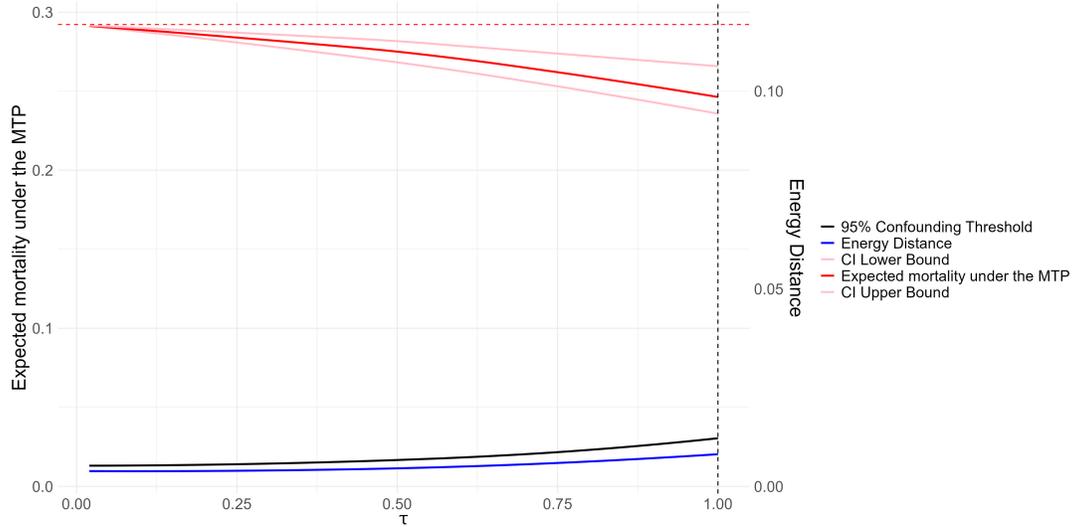


FIG 5. Exploring different magnitudes of MTP shifts to mechanical power of ventilation. The X-axis represents the τ value which controls the amount of the MTP shift. For each MTP shift, the blue curve displays the energy distance of the penalized energy balancing weights described in Section 5. The black line represents the 95% Permutation Threshold estimated using the method described in Section 4.1. The red curve along with the pink curves is the estimated mortality and the corresponding 95% bootstrap confidence interval.

In this paper, we demonstrate the connection between the estimation bias of weighted estimators of the causal effect under an MTP and weighted distributional imbalance. By explicitly defining the target population under MTP, we propose an error decomposition that highlights the role of covariate balance as a critical component of the estimation bias. We then introduce a distance, based on the weighted energy distance of [Huling and Mak \(2024\)](#), that explicitly characterizes this imbalance. As a result, the performance of any arbitrary set of weights in controlling for confounding for MTPs can be gauged through the weighted energy distance. Two methods are proposed to enhance the estimation of MTP effects: 1) we propose a method for assessing the performance of various balancing weights according to their weighted energy distance; this approach enables researchers to choose different balancing techniques suitable for specific datasets, 2) we then propose a method to detect whether the MTP shift is too large for a balancing method. We use a permutation method to measure the variability of the energy distance under the null hypothesis of no measured confounding. Comparing this with the energy distance after applying the balancing weights enables an assessment of whether the weights can effectively balance the covariates under the current MTP and thereby control for measured confounding. This comparison also sheds light on the extent of confounding challenges associated with an MTP paired with a given dataset. [For the proposed permutation test, it is important to note that, as with all statistical tests, the power to reject the null hypothesis \(and conclude that the two populations are unbalanced\) decreases with smaller sample sizes. Therefore, we recommend that researchers relax the permutation threshold to 90% or lower when sample sizes are small. Since there is no strong reason to strictly control the type I error rate to a small level in such cases, lowering the threshold can increase the power to detect imbalance after weighting.](#)

Our second contribution is a novel covariate balancing weight method based on the energy distance. These energy balancing weights are generated by minimizing the weighted energy distance, thereby minimizing a measure of distributional imbalance and thus potential for bias. We additionally propose an augmented energy balancing estimator that integrates an

outcome model with the balancing weights to help further reduce variance in the estimates. We establish the validity of our methods by demonstrating the root- n consistency of these estimators and further asymptotic normality of the augmented estimator. Through simulation studies involving both moderately complex and highly complex scenarios, we showcase the stability and performance of our methods in terms of bias, mean squared error (MSE), and coverage rate across various conditions, reflecting their capacity to handle diverse real-world situations.

We further proposed a multiplier bootstrap method to estimate the uncertainty of the augmented energy balancing estimator. The multiplier bootstrap approach offers a more computationally efficient means for statistical inference compared to the classical nonparametric bootstrap, as it does not require calculating weights within each bootstrap iteration. We prove the large-sample properties of this multiplier bootstrap and, through simulation, illustrate its similarity to the nonparametric bootstrap in terms of the coverage rates.

There remain some statistical properties that warrant further exploration. The first is to show whether the energy balancing weights converge to the true density ratio, as doing so may enable clearer studies of the potential efficiency of our proposed estimators. Secondly, as suggested in the simulation study, the coverage rate for the energy balancing estimator and the penalized energy balancing estimator is approximately 95%, implying that both estimators exhibit asymptotic normal distribution. However, establishing the asymptotic distribution for the estimators using EBWs and penalized EBWs is highly challenging. Moreover, our methods can be extended by incorporating alternative distance measures in the estimation of energy balancing weights and studying their corresponding statistical properties. [While this paper has focused on a broad set of tools for dealing with measured confounding, unmeasured confounding is always a possibility for any observational data analysis. To our knowledge, as of yet no tools to help assess sensitivity of findings to unmeasured confounding exist in the context of MTPs and we believe this is a critical area of future work, as it would help further enhance MTP analyses. Finally, for treatments that take ordinal values but are not continuous, our method \(along with the corresponding R code\) can be applied directly without modification. An ordinal treatment with finite support inherently satisfies the ‘piece-wise smooth invertibility’ condition outlined in Section 2, making it applicable to all theoretical results presented in this manuscript. However, for ordinal treatments with infinite support, this condition does not hold as stated, though it may be possible to extend our results to such cases.](#)

REFERENCES

- ATHEY, S., IMBENS, G. W. and WAGER, S. (2018). Approximate residual balancing: debiased inference of average treatment effects in high dimensions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **80** 597–623.
- CHATTOPADHYAY, A., HASE, C. H. and ZUBIZARRETA, J. R. (2020). Balancing vs modeling approaches to weighting in practice. *Statistics in Medicine* **39** 3227–3254.
- CRESSONI, M., GOTTI, M., CHIURAZZI, C., MASSARI, D., ALGIERI, I., AMINI, M., CAMMAROTO, A., BRIONI, M., MONTARULI, C., NIKOLLA, K. et al. (2016). Mechanical power and development of ventilator-induced lung injury. *Anesthesiology* **124** 1100–1108.
- DÍAZ, I. and HEJAZI, N. S. (2020). Causal mediation analysis for stochastic interventions. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **82** 661–683.
- DÍAZ, I., WILLIAMS, N., HOFFMAN, K. L. and SCHENCK, E. J. (2021). Nonparametric causal effects based on longitudinal modified treatment policies. *Journal of the American Statistical Association* 1–16.
- GATTINONI, L., TONETTI, T., CRESSONI, M., CADRINGER, P., HERRMANN, P., MOERER, O., PROTTI, A., GOTTI, M., CHIURAZZI, C., CARLESSO, E. et al. (2016). Ventilator-related causes of lung injury: the mechanical power. *Intensive Care Medicine* **42** 1567–1575.
- HAINMUELLER, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis* **20** 25–46.

- HANEUSE, S. and ROTNITZKY, A. (2013). Estimation of the effect of interventions that modify the received treatment. *Statistics in Medicine* **32** 5260–5277.
- HEJAZI, N. S., RUDOLPH, K. E., VAN DER LAAN, M. J. and DÍAZ, I. (2023). Nonparametric causal mediation analysis for stochastic interventional (in) direct effects. *Biostatistics* **24** 686–707.
- HONG, Y., CHEN, L., PAN, Q., GE, H., XING, L. and ZHANG, Z. (2021). Individualized mechanical power-based ventilation strategy for acute respiratory failure formalized by finite mixture modeling and dynamic treatment regimen. *EClinicalMedicine* **36**.
- HULING, J. D., GREIFER, N. and CHEN, G. (2023). Independence weights for causal inference with continuous treatments. *Journal of the American Statistical Association* **in press** 1–25.
- HULING, J. D. and MAK, S. (2024). Energy Balancing of Covariate Distributions. *Journal of Causal Inference* **12** 20220029. <https://doi.org/doi:10.1515/jci-2022-0029>
- IMAI, K. and RATKOVIC, M. (2014). Covariate balancing propensity score. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76** 243–263.
- JOHNSON, A. E., POLLARD, T. J., SHEN, L., LI-WEI, H. L., FENG, M., GHASSEMI, M., MOODY, B., SZOLOVITS, P., CELI, L. A. and MARK, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data* **3** 1–9.
- KANG, J. D. and SCHAFER, J. L. (2007). Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science* **22** 523–539.
- KENNEDY, E. H. (2019). Nonparametric causal effects based on incremental propensity score interventions. *Journal of the American Statistical Association* **114** 645–656.
- KENNEDY, E. H., MA, Z., MCHUGH, M. D. and SMALL, D. S. (2017). Non-parametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **79** 1229–1245.
- LEHMANN, E. L., ROMANO, J. P. and CASELLA, G. (1986). *Testing statistical hypotheses* **3**. Springer.
- LI, F., MORGAN, K. L. and ZASLAVSKY, A. M. (2018). Balancing covariates via propensity score weighting. *Journal of the American Statistical Association* **113** 390–400.
- MAK, S. and JOSEPH, V. R. (2018). Support points. *The Annals of Statistics* **46** 2562–2592.
- MATSOUAKA, R. A., LIU, Y. and ZHOU, Y. (2023). Variance estimation for the average treatment effects on the treated and on the controls. *Statistical Methods in Medical Research* **32** 389–403. <https://doi.org/10.1177/09622802221142532>
- MUÑOZ, I. D. and VAN DER LAAN, M. (2012). Population intervention causal effects based on stochastic interventions. *Biometrics* **68** 541–549.
- NAIMI, A. I., MOODIE, E. E., AUGER, N. and KAUFMAN, J. S. (2014). Constructing inverse probability weights for continuous exposures: a comparison of methods. *Epidemiology* 292–299.
- NETO, A., AMATO, M. and SCHULTZ, M. (2016). Dissipated energy is a key mediator of VILI: rationale for using low driving pressures. *Annual update in intensive care and emergency medicine* 2016 311–321.
- NETO, A. S., DELIBERATO, R. O., JOHNSON, A. E., BOS, L. D., AMORIM, P., PEREIRA, S. M., CAZATI, D. C., CORDIOLI, R. L., CORREA, T. D., POLLARD, T. J. et al. (2018). Mechanical power of ventilation is associated with mortality in critically ill patients: an analysis of patients in two observational cohorts. *Intensive Care Medicine* **44** 1914–1922.
- NIEMAN, G. F., SATALIN, J., ANDREWS, P., HABASHI, N. M. and GATTO, L. A. (2016). Lung stress, strain, and energy load: engineering concepts to understand the mechanism of ventilator-induced lung injury (VILI). *Intensive Care Medicine Experimental* **4** 1–6.
- RIZZO, M. L. and SZÉKELY, G. J. (2016). Energy distance. *wiley interdisciplinary reviews: Computational statistics* **8** 27–38.
- ROBINS, J. M., HERNÁN, M. A. and SIEBERT, U. (2004). Effects of multiple interventions. *Comparative quantification of health risks: global and regional burden of disease attributable to selected major risk factors* **1** 2191–2230.
- SEJDINOVIC, D., SRIPERUMBUDUR, B., GRETTON, A. and FUKUMIZU, K. (2013). Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The Annals of Statistics* **41** 2263–2291. <https://doi.org/10.1214/13-AOS1140>
- STOCK, J. H. (1989). Nonparametric Policy Analysis. *Journal of the American Statistical Association* **84** 567–575. <https://doi.org/10.1080/01621459.1989.10478805>
- SZÉKELY, G. J., RIZZO, M. L. and BAKIROV, N. K. (2007). Measuring and testing dependence by correlation of distances. *The Annals of Statistics* **35** 2769–2794.
- SZÉKELY, G. J. and RIZZO, M. L. (2013). Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference* **143** 1249–1272.
- VAN DER LAAN, M. J., POLLEY, E. C. and HUBBARD, A. E. (2007). Super learner. *Statistical applications in genetics and molecular biology* **6**.
- WU, C.-F. J. (1986). Jackknife, bootstrap and other resampling methods in regression analysis. *The Annals of Statistics* **14** 1261–1295.