Optimal Transport Weights for Causal Inference

Eric A. Dunipace[*]

David Geffen School of Medicine at UCLA
Department of Biostatistics, Harvard T.H. Chan School of Public Health

Abstract

Imbalance in covariate distributions leads to biased estimates of causal effects. Weighting methods attempt to correct this imbalance but rely on specifying models for the treatment assignment mechanism, which is unknown in observational studies. This leaves researchers to choose the proper weighting method and the appropriate covariate functions for these models without knowing the correct combination to achieve distributional balance. In response to these difficulties, we propose a nonparametric generalization of several other weighting schemes found in the literature: Causal Optimal Transport. This new method directly targets distributional balance by minimizing optimal transport distances between treatment and control groups or, more generally, between any source and target population. Our approach is semiparametrically efficient and model-free but can also incorporate moments or any other important functions of covariates that a researcher desires to balance. Moreover, our method can provide nonparametric estimate the conditional mean outcome function and we give rates for the convergence of this estimator. Moreover, we show how this method can provide nonparametric imputations of the missing potential outcomes and give rates of convergence for this estimator. We find that Causal Optimal Transport outperforms competitor methods when both the propensity score and outcome models are misspecified, indicating it is a robust alternative to common weighting methods. Finally, we demonstrate the utility of our method in an external control trial examining the effect of misoprostol versus oxytocin for the treatment of post-partum hemorrhage.

Keywords: balancing weights, semiparametric efficiency, IPW, Sinkhorn divergence

---

[*]edunipace@mail.harvard.edu

1 Introduction

Inverse propensity score weighting (IPW) uses the predicted probabilities of treatment to estimate causal effects. Under the correct model, these weights will lead to distributional balance between treatment groups and thereby to unbiased effect estimates. Unfortunately, adequate distributional balance is a challenge to determine, so researchers often assess performance by measuring the balance of observed covariate functions (Li et al., 2018).

Rather than using weights that indirectly balance selected covariate functions, other authors construct weights that achieve such balance by design (Hainmueller, 2012; Imai and Ratkovic, 2014; Zubizarreta, 2015); however, the particular aspect of the covariates that will achieve distributional balance, especially in finite samples, is usually unknown. Weights built using reproducing kernel Hilbert spaces (RKHS) may get around this difficulty because they theoretically balance all covariate functions (Li and Li, 2021). Nonetheless, RKHS weights necessitate the tuning of unknown hyperparameters without an obvious metric to assess performance. Some authors tune the RKHS weights such that they balance predictions from an outcome model (Kallus, 2020; Kallus et al., 2018), meaning the performance of the weights in turn depends on the performance of these outcome models. However, all of these balancing methods fail to capture something fundamental about the covariates: the joint distribution.

Ultimately, if the distributions are the same between treatment groups, then all functions of the covariates will be balanced, leading to unbiased treatment effect estimates. But researchers may have certain covariate functions that they believe are important *a priori*. As such, researchers may not want to trade off balance on these quantities for better overall distributional balance. Fortunately, we do not have to choose between balancing distributions or covariate functions—we can do both.

Causal Optimal Transport (COT) is our proposed method that balances the joint distribution and any other selected covariate functions of interest in one framework. As we document, our method is semiparametrically efficient and performs well in simulation studies compared to competing methods—especially when the propensity score and outcome models are misspecified. We also demonstrate that several methods in the literature are actually special cases of COT, meaning our framework can be seen as an interpolation between several seemingly unconnected methods. Finally, we show how optimal transport methods can nonparametrically impute the missing potential outcomes by estimating the conditional mean outcome function at a $n^{-2/(2d'+9)}$ rate, where $d' = \lceil 5d/4 \rceil$ and d is the dimension of the covariate space. Code to implement the methods discussed in this work is available in the new R package `causalOT`, found at http://www.github.com/ericdunipace/causalOT.

The paper proceeds as follows: in Section 2, we describe the setting and assumptions necessary for identification. Then we introduce optimal transport and COT in Section 3. In Section 4, we demonstrate how COT unifies several existing methods, and in Section 5, we provide simulation results demonstrating the utility of COT. Section 6 presents a case study utilizing our method in a study of post-partum hemorrhage. Finally, we offer our concluding

remarks in Section 7.

2 Setup

2.1 The potential outcomes framework

We adopt the potential outcomes framework of Neyman and Rubin (Splawa-Neyman, 1923; Rubin, 1974). Assume that we have an independent, identically distributed (iid) sample of $n \in \mathbb{N}$ units from some population. Let Z be a binary variable that denotes receiving either a treatment ($Z = 1$) or control ($Z = 0$) condition. $Y(0)$ and $Y(1)$ are the potential outcomes, and $Y = Z \cdot Y(1) + (1 - Z) \cdot Y(0)$ is the observed outcome defined on a space $\mathcal{Y} \subseteq \mathbb{R}$. The confounders are $X \in \mathcal{X} \subseteq \mathbb{R}^d$. We will assume we have $n_0 \in \mathbb{N}$ control units and $n_1 \in \mathbb{N}$ treated units giving $n = n_0 + n_1$ total observations from this sample. Denote $\mathbf{a}_z = \sum_{i:Z_i=z}^{n} \delta_{x_i} a_i$ as the empirical distribution for treatment $Z = z$ and $\mathbf{a} = \sum_{i}^{n} \delta_{x_i} a_i$ as the empirical distribution for the full sample. Let α_z and α be the corresponding population distributions.

Finally, we assume the space $\mathcal{X}$ has a distance metric between observations, $d_\mathcal{X}(x_i, x_j) \in \mathbb{R}_+$. We will define a generic cost function as $c(x_i, x_j) = d_\mathcal{X}(x_i, x_j)^p$ with $p \geq 1$. As an example, if $d_\mathcal{X}$ is the Euclidean distance and $p = 2$, then c is the squared-Euclidean distance. From this function, we then construct a pairwise cost matrix $\mathbf{C} \in \mathbb{R}_+^{n \times m}$ between each unit i and j: $\mathbf{C}_{ij} = c(x_i, x_j)$.

2.2 Causal estimands

There are several potential causal contrasts of interest over these populations but we focus on the sample average treatment effect (ATE):

$$\tau = \mathbb{E}\left\{Y(1) - Y(0)\right\}, \tag{1}$$

Unfortunately, we cannot estimate Eq. (1) since one of the potential outcomes is missing for each individual.

Instead, we need to use the information in the source population to get valid treatment effect estimates. That is, we desire $\mathbb{E}\left\{Y(z)\mathbb{I}(Z = z) \cdot w\right\} = \mathbb{E}\left\{Y(z)\right\}$, for some function w. A common way to do this is to use an importance sampling weight $w = \frac{d\alpha}{d\alpha_z}$, which is also known as the Radon-Nikodym derivative of α with respect to α_z.

With a known w, we estimate Eq. (1) as

$$\hat{\tau} = \sum_i w_i Y_i Z_i - \sum_i w_i Y_i (1 - Z_i), \tag{2}$$

and such that the weights sum to one in the treatment and control groups: $\sum_i w_i Z_i = \sum_i w_i (1 - Z_i) = 1$.

2.3 Identifying assumptions

To identify these estimators, we need several assumptions, which we formalize below.

Assumption 1

Stable unit treatment value assumption, (Rubin, 1986):

$Y_i(Z_1, Z_2, ..., Z_i, ..., Z_n) = Y_i(Z_i)$ *and* $Y_i(Z_i) = Y_i(z)$ *if* $Z_i = z$.

Assumption 2

Strong ignorability of treatment assignment, (Rosenbaum and Rubin, 1983):

$Y(0), Y(1) \perp\!\!\!\perp Z \mid X, S = 1$ *and* $0 < P(Z = 1 \mid X, S = 1) < 1$.

These standard assumptions allow us to use the observed data to estimate the desired treatment effects in the target sample. Assumption 1 allows us to use the observed outcomes and not consider interference between units, while Assumption 2 gives common support between treatment populations. With these conditions, we now turn to optimal transport and COT.

3 Causal Optimal Transport

3.1 General properties of optimal transport

The popularity of optimal transport methods have exploded in recent years thanks to several recent theoretical and methodological advances (Cuturi, 2013; Peyré and Cuturi, 2019), but the field dates back centuries. We frame our discussion in terms of empirical samples $\mathbf{a}_z$ and $\mathbf{a}$ to align with the rest of the paper but these quantities can be arbitrary samples for general optimal transport problems.

The original optimal transport problem formulated by Monge (1781) involves finding optimal maps between distributions. Define such a map as a function $T : \mathcal{X} \mapsto \mathcal{X}$ and such that $\int_{\mathcal{X}} g(x) d\alpha_z = \int_{\mathcal{X}} g(T(x)) d\alpha$ for all measurable functions g. We denote the corresponding push-forward operator from α to α_z as $T_{\#}\alpha = \alpha_z$. The Monge formulation of the optimal transport problem is then

$$\inf_T \quad \sum_i c\{T(x_i), x_i\} a_j, \tag{3}$$

where $T_{\#}\alpha = \alpha_z$. Unfortunately, this problem can be intractable to solve in practice since the mapping must be injective.

To alleviate this issue, the Kantorovich (1942) formulation instead considers probabilistic assignments between

distributions $\mathbf{a}_z$ and $\mathbf{a}$ using a transport matrix $\mathbf{P}$:

$$\text{OT}(\mathbf{a}_z, \mathbf{a}) = \min_{\mathbf{P} \in \mathbf{U}(\mathbf{a}_z, \mathbf{a})} \sum_{i:Z_i=z, j} \mathbf{C}_{ij} \mathbf{P}_{ij} \tag{4}$$

where $\mathbf{U}(\mathbf{a}_z, \mathbf{a})$ is the set of joint distributions with margins $\mathbf{a}_z$ and $\mathbf{a}$. This metric is a proper distance that obeys the triangle inequality and metrizes the convergence in distribution, *i.e.* $\text{OT}(\mathbf{a}_z, \mathbf{a}) = 0 \iff \mathbf{a}_z = \mathbf{a}$ (Proposition 2.3, Peyré and Cuturi, 2019). When $c(x, x') = d_{\mathcal{X}}(x, x')^p$, as is the case for our setting, then Eq. (4) is also known as the p-Wasserstein distance. Unfortunately, this problem is known to have a decaying convergence with increasing dimension (Weed and Bach, 2019) and also to suffer from a large computational complexity (Cuturi, 2013).

Conveniently, regularized optimal transport offers improved rates of asymptotic convergence (Genevay et al., 2019; Mena and Weed, 2019) and computational speed (Altschuler et al., 2017) by adding a convex penalty to the objective function:

$$\text{OT}_\lambda(\mathbf{a}_z, \mathbf{a}) = \min_{\mathbf{P} \in \mathbf{U}(\mathbf{a}_z, \mathbf{a})} \sum_{i:Z_i=z, j} \mathbf{C}_{ij} \mathbf{P}_{ij} + H_\lambda(\mathbf{P}_{ij}). \tag{5}$$

Common penalties for H_λ include an entropy penalty, $\lambda \mathbf{P}_{ij} \log \mathbf{P}_{ij}$ (Cuturi, 2013), or an L_2 penalty, $\frac{\lambda}{2} \mathbf{P}_{ij}^2$ (Blondel et al., 2018). The solutions to this problem converge to the solutions from Eq. (4) as $\lambda \to 0$, while as $\lambda \to \infty$, the solutions put equal weight on every entry in $\mathbf{P}$.

To adjust for the fact that $\text{OT}_\lambda(\mathbf{a}_z, \mathbf{a}) \neq 0$, Genevay et al. (2018) introduced the Sinkhorn divergence for entropy penalized optimal transport:

$$S_\lambda(\mathbf{a}_z, \mathbf{a}) = \text{OT}_\lambda(\mathbf{a}_z, \mathbf{a}) - \frac{1}{2}\text{OT}_\lambda(\mathbf{a}_z, \mathbf{a}_z) - \frac{1}{2}\text{OT}_\lambda(\mathbf{a}, \mathbf{a}) \tag{6}$$

This has the advantage that $S_\lambda(\mathbf{a}_z, \mathbf{a}) = 0 \iff \mathbf{a}_z = \mathbf{a}$ (Feydy et al., 2019), while retaining the computational and theoretical advantages of regularized optimal transport.

Finally, we can still use Eqs. (4) or (5) to construct a map as in Eq. (3). In finite samples, this function can be estimated from the Kantorovich formulation as

$$T_{\mathbf{a} \mapsto \mathbf{a}_z}(j) = \underset{v}{\arg\min} \sum_{i:Z_i=z} c(X_i, v) \mathbf{P}_{ij}. \tag{7}$$

This mapping is alternatively known as the barycentric projection (Peyré and Cuturi, 2019). For the squared-Euclidean cost, this map equals $\frac{1}{a_j} \sum_{i:Z_i=z} \mathbf{P}_{ij} X_i$, or the weighted mean of the observations in the sample who received treatment $Z = z$. For an L_1 cost, $T_{\mathbf{a} \mapsto \mathbf{a}_z}$ is the weighted median of the corresponding X_i. Under an L_2 cost, this map will also converge to the optimal Monge map provided one of underlying measures is continuous (Ambrosio et al., 2005). With these general properties established, we now turn to our proposed method.

3.2 Problem formulation

We define the COT problem as

$$\text{COT}_\lambda\left(\mathbf{a}\right) = \min_{\mathbf{w} \in \Delta_n} \ S_\lambda(\mathbf{w}_1, \mathbf{a}) + S_\lambda(\mathbf{w}_0, \mathbf{a}), \tag{8}$$

where $\mathbf{w}_z$ is the empirical measure $\sum_{i:Z_i=z} \delta_{x_i} w_i$, Δ_n is the simplex with n vertices, and S_λ is defined in (6). The COT weights will then be the weights that minimize $\text{COT}_\lambda\left(\mathbf{a}\right)$. In a slight abuse of notation, we have the following marginal distribution:

$$\mathbf{w}_{\text{COT}} = \operatorname*{argmin}_{\mathbf{w} \in \Delta_n} \ \text{COT}_\lambda\left(\mathbf{a}\right). \tag{9}$$

In addition to seeking distributional balance, a researcher may also know a set of functions that he or she thinks are important to balance *a priori* for valid causal estimates. These functions may include a hypothesized outcome model or the moments of the covariates. Define $B_k(\cdot) : \mathcal{X} \mapsto \mathbb{R}$ for $k \in \{1, ..., K\}$ as these K functions of interest. We can then add an additional constraint to the problem in Eq. (8) to approximately balance these important functions between samples:

$$\left| \sum_{i:Z_i=z} B_k(X_i) w_i - \frac{1}{m} \sum_j B_k(X_j) \right| \leq \delta_k, \ \forall k \in \{1, ..., K\}. \tag{10}$$

Of course, other formulations of the problem are possible and we detail some of them in Appendix D of the Supplementary Materials. However, we find that in practice the formulation in Eq. (8) has the best performance in terms of bias and variance.

3.3 Convergence

We now discuss the convergence of our weights to the distribution of interest. First, we define the importance sampling weights as $\check{w}_i^\star = \frac{d\alpha(X_i)}{d\alpha_z(X_i)}$ and define the self-normalized importance sampling weights as $w_i^\star = \frac{1}{n} \check{w}_i^\star / \sum_i \frac{1}{n} \check{w}_i^\star$. In our setting, $\check{w}_i^\star = 1/\ \mathbb{P}(Z_i = z \,|\, X_i)$. Further, let δ_n be the smallest value of the balancing function constraints at which the importance sampling weights satisfy the condition in Eq. (10) for sample size n. We also rely on some additional assumptions to prove the convergence of the COT weights.

Assumption 3

$\exists x_0 \in \mathcal{X} : \int_{\mathcal{X}} c(x_0, x) d\alpha < \infty$ *and* $\mathbb{E}_\alpha |B(X)| < \infty$ *with* $\|\delta_n\|^2 = o_p(\|\delta\|^2)$.

Assumption 4

$c(\cdot, \cdot)$ *is in* $\mathcal{C}^\infty$ *and is L-Lipschitz and either* 1) $\frac{1}{\lambda^{\lceil 5d/4 \rceil + 2}} \frac{1}{\sqrt{n}} = o_p(1)$ *and* α_z *and* α *are σ^2-subgaussian with* $c = \|\cdot\|_2^2$ *or*

2) $\frac{\exp(\|\mathbf{C}\|_\infty/\lambda)}{\lambda^{\lceil d/2 \rceil}} \frac{1}{\sqrt{n}} = o_P(1)$ *and* $\mathcal{X} \subset \mathbb{R}^d$.

These assumptions enforce some regularity on the constituent parts of COT. First, the cost function must exist and be continuously differentiable. Second, the measures are either subgaussian or defined on subsets of the real numbers. Third, the penalty term λ cannot go to zero too quickly if at all. And fourth, if using balancing constraints, there needs to be a value at which the importance sampling weights satisfy the constraints. With these assumptions, we have our first theorem.

Theorem 1

If Assumptions 2–4 hold, then as $n \to \infty$,

$$\mathbf{w}_{COT} \rightharpoonup \alpha,$$

where $\mathbf{w}_{COT}$ is defined in Eq. (9). Further,

$$\mathbb{E}\left\{\mathrm{OT}_\lambda\left(\mathbf{w}_{COT}, \mathbf{a}\right) - \mathrm{OT}_\lambda\left(\alpha, \alpha\right)\right\} = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right).$$

This theorem says that the COT weights converge to the distribution of the target sample at a $\sqrt{n}$-rate, which also has implications for the efficiency of estimators based on COT, as we will see in the next section. A proof of this theorem is provided in Appendix A.1 of the Supplementary Materials.

Finally, Theorem 1 also gives the following corollary.

Corollary 1

As $n \to \infty$,

$$\lim_{n \to \infty} \mathbf{w}_{COT} \overset{a.s.}{=} \lim_{n \to \infty} \mathbf{w}^\star.$$

The corollary follows as a consequence of the Radon-Nikodym Theorem and the fact that the Radon-Nikodym derivatives are almost surely unique.

3.4   Statistical Inference

For statistical inference, we turn our attention to the asymptotic distribution of Eq. (2) and its variance. We assume the following conditions also hold.

Assumption 5

$\mathbb{E}|Y - \mu_z(X)| < \infty$ *for* $\mu_z(X) \overset{def.}{=} \mathbb{E}\{Y(z) \mid X\}$, $\mathbb{E}(Y^2) < \infty$, *and either* $S_\lambda(\mathbf{w}_{COT}, \mathbf{a}) = o_p(n^{-1/2})$ *or, for basis function balancing,* $\|\delta\|_2^2 = o_p(n^{-1/2})$ *with* $\mu_z(X) \subseteq B(X)^\top \gamma$ *for* $\gamma \in \mathbb{R}^K$.

This assumption has several important parts. We assume that the second moment of the outcome is finite and that the residual is L_1-integrable, which are not strong assumptions for real data. Then we require one of two

additional conditions to hold. The first potential condition is that the convergence of $\mathbf{w}_{\text{COT}}$ to $\mathbf{a}$ occurs at a faster than $\sqrt{n}$-rate. We note that this is not actually that strong of an assumption in practice since the convergence to α happens at a $\sqrt{n}$-rate and COT is directly targeting the empirical distribution $\mathbf{a}$. Thus, we expect the convergence to the empirical distribution to be faster than $1/\sqrt{n}$, which is what we observe in practice. The second potential condition requires that the basis functions B encompass the true conditional mean and that the empirical means of B converge faster than $1/\sqrt{n}$. We observe in practice that the convergence of the basis functions is actually possible with relatively small sample sizes, making this assumption very plausible; however, outcome models are typically not known—though this can be ameliorated by using nonparametric models. With these conditions, we have our next theorem.

Theorem 2

If Assumptions 1–5 hold, then as $n \to \infty$,

$$\sqrt{n}\left(\hat{\tau} - \tau\right) \xrightarrow{\mathcal{L}} \mathcal{N}\left(0, V_{opt}\right),$$

where V_{opt} is the semiparametrically efficient variance as in Theorem 1 of Hahn (1998).

This result follows from the fact that the expansion of the bias $\hat{\tau} - \tau$ has the form of the semiparametrically efficient score function. We defer a proof to Appendix A.2 of the Supplementary Materials.

Theorem 2 also gives us the following corollary.

Corollary 2

Under Assumptions 1–5, then COT is doubly robust for a large enough n:

$$\hat{\tau}_{dr} = n^{-1} \sum_{i=1}^{n} nw_i Z_i \{Y_i - \mu_1(X_i)\} - nw_i(1 - Z_i)\{Y_i - \mu_0(X_i)\} + \{\mu_1(X_i) - \mu_0(X_i)\}$$

$$\approx n^{-1} \sum_{i=1}^{n} nw_i Z_i Y_i - nw_i(1 - Z_i)Y_i = \hat{\tau}$$

In practice, one can check if Corollary 2 holds by examining both the optimal transport distance between distributions $\mathbf{w}_{\text{COT}}$ and $\mathbf{a}$ as well as the balance of the hypothesized outcome models between samples of the covariate functions that determine the assumed outcome models . If the hypothesized outcome models are well-balanced, then there is little benefit to model augmentation. A manifestation of this phenomenon can be seen in the simulations in Section 5 where adding in model augmentation does not change the estimates from using COT even for sample sizes as low as 500.

3.5 Imputing the missing potential outcomes

One of the advantages of COT is that it provides a method to impute the missing potential outcomes, if so desired. Moreover, methods based on this estimator will also converge to the correct treatment effect, even when using weights from other methods.

We can construct a transportation matrix, $\mathbf{P}$, *a posteriori* for Eq. (8) by solving $\mathrm{OT}_{\lambda'}(\mathbf{w}_z, \mathbf{a})$, for any $\lambda' > 0$ using the appropriate weights from treatment group $Z = z$. Then the missing potential outcomes can be estimated by the barycentric projection in Eq. (7): $\hat{Y}_j(z) = \operatorname{argmin}_\nu \sum_i \mathbb{I}(Z_i = z) c(Y_i, \nu) \mathbf{P}_{ij}$. In practice, we do not have to use the same cost function used to estimate $\mathbf{P}$ but using a squared-L_2 cost gives us the following theorem

Theorem 3

Assume that α is compactly supported and admits a density with finite Fisher information I_0 and finite second moments, and $c = \|\cdot\|_2^2$. Further, assume μ_z is L-Lipschitz, $\mathrm{Var}(Y - \mu_z \mid X) < \xi^2 < \infty$ for all $X \in \mathcal{X}$, and that Assumptions 1–2 hold. Then if $\mathbf{w} \rightharpoonup \alpha$, $\mathbf{P}$ is estimated via Eq. (5) with an entropy penalty, and $\lambda \asymp n^{-\frac{1}{2d'+9}}$ for $d' = \lceil 5d/4 \rceil$,

$$\mathbb{E}_\alpha \|\hat{Y}(z) - \mu_z(X)\|^2 \lesssim (1 + I_0) n^{-\frac{2}{2d'+9}}.$$

There are several things to observe about this theorem and its assumptions. First, the assumptions on the cost and distribution allow us to connect the optimal transport solutions to the Monge maps of Eq. (3), but also allow us to give a rate for our theorem. Second, using penalized optimal transport ensures that we average out the errors $Y - \mu_z$. Third, in a related manner, the Lipschitz continuity of the outcome means that this averaging out of the errors will still achieve good estimates of μ_z. Fourth and finally, this theorem also suggests that μ_z is like a Monge map, T, between potential outcomes.

Unfortunately, these imputations are not necessarily useful by themselves.

Proposition 1

Under a squared-L_2 metric, an ATE estimator based solely on the barycentric projection, $n^{-1} \sum_i \hat{Y}_i(1) - \hat{Y}_i(0)$, is equivalent to Eq. (2).

Clearly, care must be taken when using these estimators. Proofs are located in Appendix A.3 and 1 of the Supplementary Materials.

3.6 Practical considerations

In this section, we turn to the practical considerations of optimizing the COT weights. Namely, we discuss the tuning of the hyperparameters and estimation of the weights.

Algorithm 1: Choosing hyperparameters for COT

Data: Grid of parameter values $\Delta = \{\{\lambda_1, \delta_1\}, \{\lambda_2, \delta_2\}, ...\}$, number of bootstrap samples K, empirical measure $\mathbf{a}$, empirical measure $\mathbf{a}_z$, treatment group of interest z, $\lambda' \geq 0$

Result: Value of hyperparameters, $\hat{\lambda}, \hat{\delta}$

foreach $\{\lambda, \delta\} \in \Delta$ **do**

    Estimate weights $\mathbf{w}_{\text{COT}}$ given parameters $\{\lambda, \delta\}$;

    for k *in* $1, ..., K$ **do**

        Bootstrap new target data $\mathbf{a}_k^\star \sim \mathbf{a}$;

        Bootstrap new source data $\mathbf{a}_{z,k}^\star \sim \mathbf{a}_z$;

        Set the unnormalized weights for empirical measure in treatment group z as

        $\tilde{\mathbf{w}}_k^\star = \mathbb{I}(Z = z) \odot \mathbf{w} \odot \mathbf{a}_{z,k}^\star$, where $\odot$ is the element-wise product;

        Renormalize the weights, $\mathbf{w}_k^\star = \tilde{\mathbf{w}}_k^\star / (\tilde{\mathbf{w}}_k^{\star\top} \mathbf{1}_n)$;

    Set $T_{\{\lambda, \delta\}} = K^{-1} \sum_k^K \text{OT}_{\lambda'}(\mathbf{w}_k^\star, \mathbf{a}_k^\star)$;

return $\hat{\lambda}, \hat{\delta} = \text{argmin}_{\{\lambda, \delta\}} T_{\{\lambda, \delta\}}$;

Hyperparameter tuning. Our goal is to select the hyperparameters λ and δ so that we achieve the best distributional balance without over-fitting the current data. To do so, we propose a bootstrap based tuning procedure detailed in Algorithm 1.

We justify this procedure in two ways. First, practitioners probably do not have an ideal weight penalty in mind based on subject matter knowledge. Second, because the COT weights target the Radon-Nikodym derivative, this tuning procedure will select the hyperparameters that lead to weights robust to sampling variation and better approximation of these population level quantities. We present an empirical examination of this tuning algorithm in Appendix B.3 of the Supplementary Materials that demonstrates its effectiveness at selecting the optimal λ.

Weight estimation. Given the known complexity of estimating optimal transport distances, Huling and Mak (2020) raise the concern that methodology like COT will not be feasible. Fortunately, these concerns are addressed by using regularized optimal transport.

In our simulations, we find that estimating the COT weights only takes a few seconds for a 1000 observations. Eq. (8) can be solved by alternating Sinkhorn divergence calculations in GeomLoss (Feydy et al., 2019) and optimization steps on the weights. With balancing constraints, we use the Frank-Wolfe algorithm to optimize the weights (Frank and Wolfe, 1956); without balancing constraints, we can use an LBFGS algorithm.

3.7 Target average treatment effects and multi-valued treatments

COT is also well-suited to the case where the target estimand is for a separate sample entirely. This is because the weights can be calibrated to target any arbitrary set of samples. The only additional assumption for convergence is that there is common support between distributions.

Finally, COT is easily extended to more than two treatments as long as the treatment values are discrete. This is because COT estimates weights separately for each treatment group.

4   Connections to Existing Methods

The COT framework is actually related to several other methods in the literature and can be seen as an interpolation between all of them, as we detail below.

Synthetic control method   Abadie and Gardeazabal ([2003](#)) first proposed the synthetic control method (SCM) as a way of performing counterfactual inference for a single treated unit, j. The objective function is

$$\min_{w:w^\top \mathbf{1}_n=1, w_i \geq 0} \left\| \sum_i X_i w_i - X_j \right\|_2^2,$$

which is the same objective as the Monge map in Eq. ([3](#)) when c is the squared-Euclidean distance and T has the corresponding form of the barycentric projection in Eq. ([7](#)). This means that SCM is actually estimating a version of the COT problem with $\lambda = 0$.

Proposition 2

If α admits a density, $c = \| \cdot \|_2^2$, and $\exists x_o$ s.t. $\mathbb{E}\{c(x_0, X)\} < \infty$, then SCM is asymptotically the same as COT with $\lambda = 0$.

For a proof, see Appendix [A.5](#) in the supplementary materials. One potential drawback of using SCM versus the formulation used for COT can be seen in the following simple example in Figure [1](#). COT favors the nearest point while the SCM method utilizes the points further away, which could be a problem if the response surface looks like Figure [1b](#). To avoid this, SCM could incorporate a modified objective that directly models both the barycentric projection and distance between units as in Perrot et al. ([2016](#)).

Nearest neighbor matching   When done with replacement, nearest neighbor matching (NNM) is also a reformulation of the COT problem with $\lambda = 0$: $\min_{\mathbf{P} \geq 0} \sum_{i:Z_i=z, j} \mathbf{C}_{ij} \mathbf{P}_{ij}$ subject to $\sum_{i,j} \mathbf{P}_{ij} \mathbb{I}(Z_i = z) = 1$ and $\mathbf{P}^\top \mathbf{1}_n = \mathbf{a}$. This will seek to find the unit i that is closest in terms of $\mathbf{C}$ for each unit j since this minimizes the total cost, which is the definition of NNM with replacement. Each observation will simply be weighted by the number of times it is matched, divided by the total sample size n.

NNM has some advantages and drawbacks relative to more general COT. Positive aspects of the method are that quick to estimate and also corresponds to an easily understood quantity of physical matching familiar to many researchers. However, the method can have poor convergence properties if $c \neq d_{\mathcal{X}}(\cdot, \cdot)^p$ where $p > d/2$ (Fournier
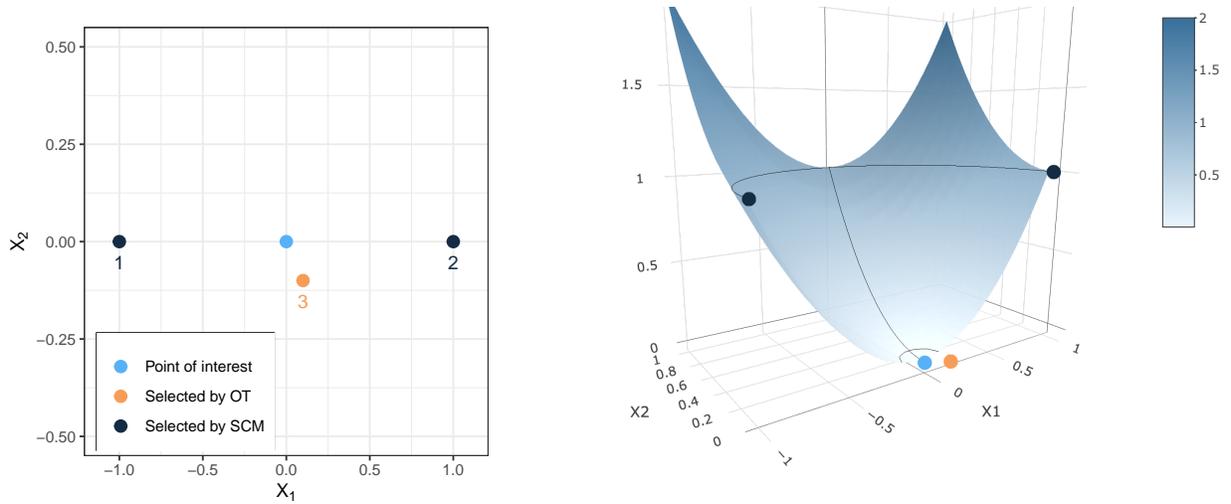
(a) Potential points for selection and the point of interest. **(b)** A response surface where synthetic controls would lead to poor estimates.

Figure 1: A simple example where synthetic controls (SCM) would perform worse than COT.

and Guillin, [2015]). This means that additional assumptions are necessary to ensure adequate convergence. Further, the weights will be given as rational numbers and as such, we would expect estimators based on them to have higher variance than weights without such constraints.

Optimal matching and MIP matching Optimal Matching (Rosenbaum, [1989]) corresponds to COT weights with integer solutions and $\lambda = 0$: $\min_{\mathbf{P} \in \{0,1\}} \sum_{i: Z_i = z, j} \mathbf{C}_{ij} \mathbf{P}_{ij}$ subject to $L \leq \sum_i \mathbf{P}_{ij} \mathbb{I}(Z_i = z) \leq U, \ \forall j$ and $1 \leq L \leq U \leq n$. We can turn this into Mixed Integer Program (MIP) matching by adding constraints:

$$\left| \sum_{ij} \frac{B_k(X_i)\mathbf{P}_{ij}\mathbb{I}(Z_i = z)}{nM_j} - \frac{1}{n} \sum_j B_k(X_j) \right| \leq \delta_k, \ \forall k \in \{1, ..., K\},$$

where $M_j = \sum_i \mathbf{P}_{ij}$ is the number of matches for unit j (Zubizarreta, [2012]). These methods also imply a re-weighting of treatment group $Z = z$ since the weights on element i will be $n^{-1} \sum_j \mathbf{P}_{ij}/M_j$.

These methods have a similar flavor to NNM but with additional linear constraints. As such, they would likely share some of its benefits and drawbacks. As an advantage, these methods again yields matches which have an easy interpretation; however, as a disadvantage, the additional linear constraints will slow down the problem estimation. Moreover, the problem has additional tuning parameters L, U, and δ_k not present in NNM. Similar to NNM, we would again expect weights based on rational numbers to have higher variance.

Energy distance. The Energy Distance (ED) is defined as $\mathcal{E}(\mathbf{a}_z, \mathbf{a}) = \frac{2}{n^2} \sum_{i: Z_i = z, j} d_{\mathcal{X}}(x_i, x_j)^p - \frac{1}{n^2} d_{\mathcal{X}}(x_i, x_i)^p - \frac{1}{m^2} d_{\mathcal{X}}(x_j, x_j)^p$, where $p \geq 1$. Then $S_\lambda(\mathbf{a}_z, \mathbf{a}) \to \frac{1}{2}\mathcal{E}(\mathbf{a}_z, \mathbf{a})$ as $\lambda \to \infty$ (Feydy et al., [2019]). Thus, Energy Balancing

Weights (Huling and Mak, 2020) are a special case of COT.

One advantage of this method is that there is no tuning parameter necessary to estimate the weights. However, we might assume that there would be an advantage to interpolating between all of these various methods. Indeed, in our experiments we find that the larger values of λ do not approximate the true inverse propensity score as well as intermediate values. See Appendix B.3 in the Supplementary Materials for an empirical evaluation.

Mean Maximum Discrepancy. Optimal transport is also related to the mean maximum discrepancy (MMD) through the ED. The MMD is equal to $\mathcal{M} = 0.5 \int_{\mathcal{X} \times \mathcal{X}} k(x, x') d\phi(x)$, for $\phi = \alpha - \beta$. For some reproducing kernel Hilbert space k and for a distance $d_{\mathcal{X}}$ defined as $d_{\mathcal{X}}(x, x') = \frac{1}{2} k(x, x') + \frac{1}{2} k(x, x') - k(x, x')$, MMD is equivalent to the ED (Feydy et al., 2019) and, therefore, to COT. To our knowledge, there has not been a proposed weighting method based on the MMD but we would expect it to have performance similar to that of the ED and COT with large values of λ.

5   Simulation Study

To evaluate the finite sample performance of the proposed weighting methodology, we use the simulation study originally presented in Hainmueller (2012). For each setting, we run 1000 experiments with a sample size of $n = 512$. The estimand of interest is the ATE.

For additional experiments examining the convergence of COT, its confidence interval coverage, and the efficacy of Algorithm 1, see Appendix B of the Supplementary Materials.

5.1   Setup

Study design. We generate six covariates $X_1, ..., X_6$ from the following distributions

$$
\begin{bmatrix} X_1 \\ X_2 \\ X_3 \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 2 & 1 & -1 \\ 1 & 1 & -0.5 \\ -1 & -0.5 & 1 \end{bmatrix} \right)
$$

$$X_4 \sim \text{Unif}(-3, 3)$$

$$X_5 \sim \chi_1^2$$

$$X_6 \sim \text{Bern}(0.5).$$

In this study, the last three covariates are mutually independent of each other and also of the first three covariates.

The treatment indicator is generated as $Z = \mathbb{I}(X_1 + 2X_2 - 2X_3 - X_4 - 0.5X_5 + X_6 + \nu > 0)$, where ν is drawn from one of three distributions leading to different degrees of overlap: $\nu \sim \mathcal{N}(0, 100)$ for high overlap, $\nu \sim g(\chi_5^2)$ for medium overlap, and $\mathcal{N}(0, 30)$ for low overlap. The function g in the medium-overlap setting gives the χ_5^2 draws expectation 0.5 and variance 67.6. We expect the scenarios that will lead to the highest bias to be in the low-overlap setting where there is a strong separation between treatment groups and in the medium-overlap setting where the errors are leptokurtic.

Given the treatment indicator Z and the covariates $X_1, ..., X_6$, we draw the outcome Y from $Y(0) = Y(1) = (X_1 + X_2 + X_5)^2 + \eta$, with $\eta \sim \mathcal{N}(0, 1)$. There are two things to note about this outcome model. The first is that there is no effect of the treatment at the unit level and hence the ATE is 0. The second is that a linear outcome model should be biased.

Methods under examination. We compare our methodology to several other weighting methods commonly used in the literature. The first method we consider is a logistic regression (GLM) using only first order terms. We also consider balancing methods such as the covariate balancing propensity score (CBPS) of Imai and Ratkovic (2014) and the stable balancing weights (SBW) of Zubizarreta (2015) both targeting mean balance. Finally, we also utilize SCM and NNM.

For the COT weights, we include two variations both using an L_2 metric with standardized covariates. The first only balances the joint distribution (no constraints or "none") and the second demonstrates basis function balancing by targeting the joint distribution as well as mean balance ("means").

5.2 Estimators

In our simulations, we consider three estimators to target the ATE. The first is known as the Hájek estimator (Hajek, 1971) and is simply a weighted mean with sum to one weights as in Eq. (2). The second is an augmented or doubly robust estimator of Robins et al. (1994), and the third is a weighted least squares estimator—both only including linear terms of the covariates. We do not include a barycentric projection estimator like in Eq. (7) since under an L_2 metric it gives the same result as the Hájek estimator (see Proposition 1).

5.3 Results

We now turn our attention to the results. Due to its nonparametric to semiparametric formulation, we expect COT to do better than other methods when the true propensity score model diverges from a logistic regression—*e.g.*, in the medium-overlap scenario.

Indeed, the COT methods have the lowest RMSE across all overlap scenarios, as we can see in Table 1. Further,

across the medium and low overlap settings, COT has the lowest bias as well; in the high overlap scenario, there is a negligible difference between COT, SBW, and GLM.

We should also note that COT gives estimates that do not vary between estimators. This is because the COT already balances the basis functions used in the augmented estimator and weighted least squares. Therefore, there is no difference between running an outcome regression utilizing linear terms of the covariates and the Hájek estimator. This is a function of the fact that by balancing distributions of the covariates, COT will also balance functions of the covariates. Similar phenomenon can also be observed with SBW.

overlap	method	constraint	Bias			RMSE		
			Hajek	DR	WOLS	Hajek	DR	WOLS
high	GLM	none	**-0.01**	**-0.01**	-0.02	1.18	1.14	1.14
	CBPS	means	0.24	**-0.01**	-0.02	1.12	1.11	1.09
	SBW	means	**-0.01**	**-0.01**	**-0.01**	1.00	1.00	1.00
	SCM	none	0.36	0.27	0.28	1.63	1.57	1.55
	NNM	none	0.43	0.32	0.28	0.69	0.65	0.56
	COT	none	**0.01**	**0.01**	**0.01**	0.61	0.61	0.61
		means	**0.01**	**0.01**	**0.01**	**0.42**	**0.42**	**0.42**
medium	GLM	none	1.12	1.10	1.04	1.72	1.69	1.70
	CBPS	means	1.20	1.06	0.95	1.72	1.64	1.56
	SBW	means	0.63	0.63	0.63	1.20	1.20	1.20
	SCM	none	1.19	1.12	1.10	2.05	1.97	1.95
	NNM	none	0.73	0.65	0.58	0.94	0.91	0.79
	COT	none	0.23	0.23	0.23	0.74	0.74	0.74
		means	**-0.03**	**-0.03**	**-0.03**	**0.43**	**0.43**	**0.43**
low	GLM	none	0.19	0.06	0.02	1.72	1.49	1.51
	CBPS	means	0.45	0.06	0.01	1.42	1.46	1.42
	SBW	means	0.03	0.03	0.03	1.03	1.03	1.03
	SCM	none	0.64	0.42	0.43	1.75	1.69	1.65
	NNM	none	0.81	0.56	0.49	1.02	0.89	0.77
	COT	none	0.05	0.05	0.05	0.85	0.85	0.85
		means	**0.00**	**0.00**	**0.00**	**0.41**	**0.41**	**0.41**

Table 1: Performance of various weighting methods under the simulation settings of Hainmueller, 2012. Bold values are the values with the lowest bias or root mean-squared error (RMSE) of the methods under the same conditions. GLM refers to weighting by the inverse of the propensity score as calculated from a logistic regression model, CBPS is the covariate balancing propensity score, SBW is the stable balancing weights, SCM is the synthetic control method, and COT is the optimal transport formulation proposed in this paper. The estimators are Hajek weights (Hajek), doubly-robust augmented IPW (DR), and weighted least squares (WOLS). All weights are normalized to sum to 1. Constraints refer to balancing constraints and are one of "none" for no constraints or "mean" for mean constraints.

6   Case Study

In this section, we apply our methodology to a real data set. There is growing interest in the literature to utilize libraries of randomized control trials (RCTs) to evaluate new interventions, the idea being that running new RCTs to evaluate every new intervention is expensive and time consuming (Schmidli et al., 2020). These studies, alternatively called externally controlled trials or synthetic control group trials, compare a set of study subjects receiving a treatment to a group of individuals external to the trial at hand who did not receive the intervention of interest. The participants used for the control group can be taken from a variety of sources such as an observational study,

electronic medical records, or from historical clinical trial data (Davi et al., 2020). To demonstrate this in practice, we present an analysis utilizing data originally from a multi-site RCT discussed by Blum et al. (2010).

6.1 Misoprostol for the Treatment of Postpartum Hemorrhage

The original study was a double-blind, non-inferiority trial that exposed 31,055 women to prophylactic oxytocin during labor at five hospitals across Burkina Faso, Egypt, Turkey, and Vietnam. The 807 women in this group with uncontrolled blood-loss after delivery—a condition known as post-partum hemorrhage or PPH—were then randomized to receive either 800 milligrams misoprostol (treatment condition) or 40 international units oxytocin (control condition). There were 407 and 402 women in each treatment group, respectively. The primary outcome for the study was whether blood loss was controlled within 20 minutes after PPH diagnosis. The authors measure several important confounders like maternal age, blood loss at treatment, whether cord traction was maintained, maternal hemoglobin, whether the mother is currently married, whether the cord was clamped early, fetal gestational age, whether labor was augmented, whether labor was induced, maternal education, number of previous live births, whether the placenta was delivered prior to hemorrhage, and whether a uterine massage was given.

6.2 Modifications and methods

We modify the study in a couple of ways to make it more similar to an externally controlled trial. For each site, we separate the paired treatment groups and attempt to estimate effects using the units from other sites—*e.g.*, for the misoprostol group in Egypt we remove the oxytocin group from Egypt and attempt to estimate a causal effect using the oxytocin groups from the other sites. In this manner, we generate effect estimates and confidence intervals for each treatment group at each site.

To estimate the treatment effects, we use COT with hyperparameter tuning as in Algorithm 1 and the squared-Euclidean distance as the cost function. We also compare COT to GLM, CBPS, SBW, SCM, and NNM. For GLM, CBPS, and SBW we utilize all first and second covariate moments. Finally, our estimator is that of Eq. (2).

For our estimates to be valid, we require that there be no unmeasured confounding but also that the estimates are "transportable," *i.e.* we are able to take estimates from the other hospitals in the external group and "transport" them to site of interest. This requires that conditional on the observed covariates there are no other variables that can effect the outcome and treatment indicator (d-separation holds, Pearl and Bareinboim, 2013).

6.3 Case study results

Amazingly, the three optimal transport flavored methods are able to achieve estimates close to the original treatment effects on average. In Figure 2, we can see that COT, NNM, and SCM do the best job in terms of average
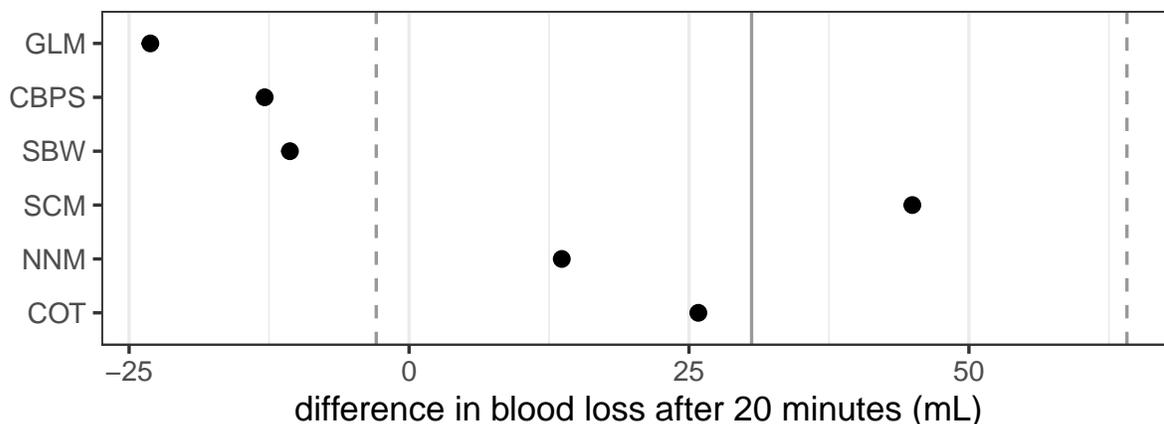
Figure 2: Results for treatment effect estimation averaged across treatment groups and study sites. The gray vertical line is the original treatment effect estimate for the entire study while the dotted vertical lines are the original confidence interval. The weighting methods under examination are logistic regression (GLM), Covariate Balancing Propensity Score (CBPS), Stable Balancing Weights (SBW), Synthetic Control Method (SCM), Nearest Neighbor Matching (NNM), and Causal Optimal Transport (COT).

bias; moreover, these are the only methods that have estimates inside the original confidence interval. Of these, we note that COT has the least overall bias across all treatment groups and sites.

In terms of inference, COT also performs best across all sites and treatment groups. Table 2 displays how well the calculated confidence intervals cover the original treatment effects and also if the calculated estimates are inside the original confidence interval. In both cases, COT has the highest percentage of confidence intervals covering the original treatment effect and estimates inside the original confidence interval at 60% each, respectively. And while NNM had a good overall bias, only 20% of its confidence intervals covered the true treatment effect and its estimates were inside of only 20% of the original confidence intervals.

Method	% C.I. covering original effect	% of estimates in original C.I.
GLM	20	20
CBPS	30	30
SBW	30	40
SCM	50	50
NNM	20	20
COT	60	60

Table 2: For each method, the table displays the percentage of times that the calculated 95% confidence interval (C.I.) covered the true treatment effect and whether the estimated treatment effect was inside the original C.I. from the study. The weighting methods under examination are logistic regression (GLM), Covariate Balancing Propensity Score (CBPS), Stable Balancing Weights (SBW), Synthetic Control Method (SCM), Nearest Neighbor Matching (NNM), and Causal Optimal Transport (COT).

7 Summary and Remarks

We have described a new tool for the estimation of causal effects in observational studies: Causal Optimal Transport. This method allows for checks of distributional overlap and model-free weight estimation that is semipara-

metrically efficient. We also showed how several other methods for causal inference are closely related to COT and that COT can be seen as an interpolation between these methods. In our simulation study, we demonstrated that this methodology performs well even when both the outcome and propensity score models are misspecified. Compared to other common weighting methods, COT generally has lower bias and lower root mean-squared error under model misspecification.

There are several areas for future research. First, the sensitivity of COT to the choice of cost function remains to be elucidated. Second, selecting covariates through typical model selection frameworks such as an L_1 penalized regression is not obvious given that COT does not generate clear predictive models, though this may not matter given the nonparametric nature of the weights. Third, more work needs to be done to extend this framework to time series data, but further connections to SCM may offer a way forward.

Acknowledgments

Supplementary Materials

The Supplementary Materials contain proofs of the theorems in the paper, further empirical studies, an additional case study, and proofs for other formulations of COT. The last two s Proofs are found in Appendix A, and empirical studies of the convergence, confidence interval coverage, and efficacy of the tuning algorithm are located in Appendix B. Appendix C is the additional case study and Appendix D are the proofs for other formulations of COT.

References

Abadie, Alberto and Javier Gardeazabal (2003). "The Economic Costs of Conflict : A Case Study of the Basque Country". In: *The American Economic Review* 93.1, pp. 113–132 (see p. 11).

Altschuler, Jason, Jonathan Weed, and Philippe Rigollet (2017). "Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration". In: *Advances in Neural Information Processing Systems* 2017, pp. 1965–1975 (see p. 5).

Ambrosio, Luigi, Nicola Gigli, and Giuseppe Savare (2005). *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Series Title: Lectures in Mathematics ETH Zürich. Basel: Birkhäuser-Verlag (see pp. 5, 26, 29).

Blondel, Mathieu, Vivien Seguy, and Antoine Rolet (2018). "Smooth and sparse optimal transport". In: *International Conference on Artificial Intelligence and Statistics, AISTATS 2018*, pp. 880–889 (see pp. 5, 37–39).

Blum, Jennifer et al. (2010). "Treatment of post-partum haemorrhage with sublingual misoprostol versus oxytocin in women receiving prophylactic oxytocin : a double-blind , randomised , non-inferiority trial". In: *The Lancet* 375.9710. Publisher: Elsevier Ltd, pp. 217–223 (see p. 16).

Boyd, Stephen and Lieven Vandenberghe (2004). *Convex Optimization*. ISSN: 0018-9286. New York: Cambridge University Press (see p. 41).

Brenier, Yann (1987). "Décomposition polaire et réarrangement monotone des champs de vecteurs". In: *C. R. Acad. Sci. Paris Sér. I Math* 305.19, pp. 805–808 (see pp. 25, 29).

Chizat, Lénaïc et al. (2020). "Faster Wasserstein Distance Estimation with the Sinkhorn Divergence". In: *NeurIPS* 2 (NeurIPS), pp. 1–29 (see p. 28).

Cuturi, Marco (June 4, 2013). "Sinkhorn Distances: Lightspeed Computation of Optimal Transport". In: *Advances in Neural Information Processing Systems*. Ed. by C J C Burges et al. Vol. 26. Curran Associates, Inc., pp. 1–9 (see pp. 4, 5).

Davi, Ruthie et al. (2020). "Informing single-arm clinical trials with external controls". In: *Nature reviews. Drug discovery* 19.12. Publisher: Springer US, pp. 821–822 (see p. 16).

Feydy, Jean et al. (2019). "Interpolating between Optimal Transport and MMD using Sinkhorn Divergences". In: *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2681–2690 (see pp. 5, 10, 12, 13, 22).

Fournier, Nicolas and Arnaud Guillin (2015). "On the rate of convergence in Wasserstein distance of the empirical measure". In: *Probability Theory and Related Fields* 162.3, pp. 707–738 (see pp. 11, 31, 40).

Frank, Marguerite and Philip Wolfe (Mar. 1956). "An algorithm for quadratic programming". In: *Naval Research Logistics Quarterly* 3.1, pp. 95–110 (see p. 10).

Genevay, Aude, Gabriel Peyré, and Marco Cuturi (2018). "Learning Generative Models with Sinkhorn Divergences". In: *AISTATS*, p. 10 (see p. 5).

Genevay, Aude et al. (2019). "Sample complexity of sinkhorn divergences". In: *AISTATS* 89 (see pp. 5, 23, 40).

Hahn, Jinyong (1998). "On the Role of the Propensity Score in Efficient Semiparametric Estimation of Average Treatment Effects". In: *Econometrica* 66.2, pp. 315–331 (see p. 8).

Hainmueller, Jens (2012). "Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies". In: *Political Analysis* 20.1, pp. 25–46 (see pp. 2, 13, 15).

Hajek, J. (1971). "Comment on "Comment on An essay on the logical foundations of survey sampling"". In: *Foundations of Statistical Inference*. Ed. by V.P. Godambe and D.A. Sprott. Holt, Rinehart and Winston, pp. 201–233 (see p. 14).

Huling, Jared D and Simon Mak (2020). "Energy Balancing of Covariate Distributions". In: *Unpublished* (see pp. 10, 13).

Imai, Losuke and Marc Ratkovic (2014). "CBPS: Covariate Balancing Propensity Score". In: pp. 243–263 (see pp. 2, 14).

Janati, Hicham, Marco Cuturi, and Alexandre Gramfort (2020). "Debiased Sinkhorn barycenters". In: *Proceedings of the 37 th International Conference on Machine Learning*, PMLR. Vol. 119. Online, p. 10 (see p. 22).

Kallus, Nathan (2020). "Generalized Optimal Matching Methods for Causal Inference". In: *Journal of Machine Learning Research* 21, pp. 1–54 (see p. 2).

Kallus, Nathan, Brenton Pennicooke, and Michele Santacatterina (2018). "More robust estimation of sample average treatment effects using Kernel Optimal Matching in an observational study of spine surgical interventions". In: *Unpublished*, pp. 1–37 (see p. 2).

Kantorovich, L. (1942). "On the transfer of masses (in Russian)". In: *Doklady Akademii Nauk* 37.2, pp. 227–229 (see p. 4).

LaLonde, Robert J. (1986). "American Economic Association Evaluating the Econometric Evaluations of Training Programs with Experimental Data Author (s): Robert J . LaLonde Source : The American Economic Review , Vol . 76 , No . 4 (Sep ., 1986), pp . 604-620 Published by : Americ". In: *The American Economic Review* 76.4, pp. 604–620 (see p. 34).

Li, Fan, Kari Lock Morgan, and Alan M. Zaslavsky (Jan. 2, 2018). "Balancing Covariates via Propensity Score Weighting". In: *Journal of the American Statistical Association* 113.521, pp. 390–400 (see p. 2).

Li, Yan and Liang Li (Feb. 1, 2021). "Propensity score analysis methods with balancing constraints: A Monte Carlo study". In: *Statistical Methods in Medical Research* (see p. 2).

Mena, Gonzalo and Jonathan Weed (2019). "Statistical bounds for entropic optimal transport: sample complexity and the central limit theorem". In: *NeurIPS*, pp. 1–23 (see pp. 5, 23, 28, 40).

Monge, Gaspard (1781). *Mémoire sur la théorie des déblais et des remblais*. De l'Imprimerie Royale (see p. 4).

Owen, Art B. (2013). "9: Importance Sampling". In: *Monte Carlo theory, methods and examples*. Unpublished (see p. 22).

Pearl, Judea and Elias Bareinboim (2013). "Transportability across studies: A formal approach introduction". In: *Technical Report No. R372* (April), pp. 1–18 (see p. 16).

Perrot, Michael et al. (2016). "Mapping estimation for discrete optimal transport". In: *Advances in Neural Information Processing Systems* (Nips), pp. 4204–4212 (see p. 11).

Peyré, Gabriel and Marco Cuturi (2019). "Computational Optimal Transport". In: *Foundations and Trends in Machine Learning* 11.5. ISBN: 2200000073, pp. 355–607 (see pp. 4, 5, 25, 27).

Pooladian, Aram-Alexandre and Jonathan Niles-Weed (Sept. 24, 2021). "Entropic estimation of optimal transport maps". In: *arXiv:2109.12004 [math, stat]* (see p. 28).

Robins, James M., Andrea Rotnitzky, and Lue Ping Zhao (1994). "Estimation of regression coefficients when some regressors are not always observed". In: *Journal of the American Statistical Association* 89.427, pp. 846–866 (see p. 14).

Rosenbaum, Paul R. (1989). "Optimal matching for observational studies". In: *Journal of the American Statistical Association* 84.408, pp. 1024–1032 (see p. 12).

Rosenbaum, Paul R. and Donald B. Rubin (1983). "The Central Role of the Propensity Score in Observational Studies for Causal Effects". In: *Biometrika* 70.1, pp. 41–55 (see p. 4).

Rubin, Donald B (1986). "Comment: Which Ifs Have Causal Answers". In: *Journal of the American Statistical Association* 81.396, pp. 961–962 (see p. 4).

— (1974). "Estimating causal effects of treatments in randomized and nonrandomized studies." In: *Journal of Educational Psychology* 66.5, pp. 668–701 (see p. 3).

Schmidli, Heinz et al. (2020). "Beyond Randomized Clinical Trials: Use of External Controls". In: *Clinical Pharmacology and Therapeutics* 107.4, pp. 806–816 (see p. 15).

Splawa-Neyman, Jerzy (1923). "On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9". In: *Roczniki Nauk Rolniczych Tom X*, pp. 1–51 (see p. 3).

Villani, C\{{\textbackslash}'e\}dric (2008). *Optimal transport, old and new*. Springer (see pp. 26, 29, 38, 39).

Villani, Cedric (2006). *Optimal transport, old and new*. Springer.

Weed, Jonathan and Francis Bach (2019). "Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance". In: *Bernoulli* 25.4, pp. 2620–2648 (see p. 5).

Zubizarreta, José R. (2012). "Using mixed integer programming for matching in an observational study of kidney failure after surgery". In: *Journal of the American Statistical Association* 107.500, pp. 1360–1371 (see p. 12).

— (2015). "Stable Weights that Balance Covariates for Estimation With Incomplete Outcome Data". In: *Journal of the American Statistical Association* 110.511, pp. 910–922 (see pp. 2, 14).

A Proofs

In this section we offer our proofs of the theorems and propositions stated in the paper.

A.1 Proof of Theorem 1

In this section, we prove Theorem 1 which establishes the convergence of the COT weights. First, we need the following lemma.

Lemma 1 (The importance sampling weights converge to α)

Let Assumption 2 holds. Define the importance sampling weights as $\breve{w}_i^\star = \frac{d\alpha_z(X_i)}{d\alpha(X_i)}$ and define the self-normalized importance sampling weights as $w_i^\star = \frac{1}{n}\breve{w}_i^\star / \sum_i \frac{1}{n}\breve{w}_i^\star$. Then,

$$\mathbf{w}^\star \rightharpoonup \alpha.$$

Proof. By Assumption 2, $\breve{w}_i^\star$ exists for all i. Then define $\mathbb{P}_n(X \in E) = \sum_i \mathbb{I}(X_i \in E) w_i^\star$, for some $E \subset \mathcal{X}$. Take $f(X) = \mathbb{I}(X \in E)$. By Theorem 9.2 in Owen (2013), $\mathbb{P}_n(X \in E) = \sum_i f(X_i) w_i^\star \xrightarrow{\text{a.s.}} \mathbb{E}_\alpha(f(X)) = \int_{\mathcal{X}} f \, d\alpha = \mathbb{P}_\alpha(X \in E)$, where $\mathbb{P}_\alpha(\cdot)$ denotes the probability of $X \in E$ when $X \sim \alpha$. Thus we have $\lim_{n\to\infty} \mathbb{P}_n(X \in E) = \mathbb{P}_\alpha(X \in E)$ and the result follows. $\qquad\square$

Now we are ready to proceed.

Proof. Under Assumption 2, $\mathbf{w}^\star$ exists and by Lemma 1, $\mathbf{w}^\star \rightharpoonup \alpha$. By Assumptions 3 and 4, S_λ is convex in its entries and metrizes the convergence in measure (Feydy et al., 2019; Janati et al., 2020). Thus, $0 \leq S_\lambda(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}) \leq S_\lambda(\mathbf{c}, \mathbf{a})$ for all $\mathbf{c} \in \Delta_n$ that meet the chosen balancing constraints, δ, which includes $\mathbf{w}^\star$ for a large enough n and m. Hence, $0 \leq S_\lambda(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}) \leq S_\lambda(\mathbf{w}^\star, \mathbf{a})$ and since $\mathbf{w}^\star \rightharpoonup \alpha$, $S_\lambda(\mathbf{w}^\star, \mathbf{a}) \to 0 \Rightarrow S_\lambda(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}) \to 0 \Rightarrow \mathbf{w}_{\mathsf{COT}} \rightharpoonup \alpha$.

We then take the same inequality and modify it slightly: $0 = S_\lambda(\mathbf{a}, \mathbf{a}) \leq S_\lambda(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}) \leq S_\lambda(\mathbf{w}^\star, \mathbf{a})$. If we add $\mathsf{OT}_\lambda(\alpha, \alpha) - \mathsf{OT}_\lambda(\alpha, \alpha)$ to each term and rearrange, we get

$$\frac{1}{2}\{\mathsf{OT}_\lambda(\mathbf{a}, \mathbf{a}) - \mathsf{OT}_\lambda(\alpha, \alpha)\} \leq \mathsf{OT}_\lambda(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}) - \frac{1}{2}\mathsf{OT}_\lambda(\mathbf{w}_{\mathsf{COT}}, \mathbf{w}_{\mathsf{COT}}) - \frac{1}{2}\mathsf{OT}_\lambda(\alpha, \alpha)$$

$$\leq \mathsf{OT}_\lambda(\mathbf{w}^\star, \mathbf{a}) - \frac{1}{2}\mathsf{OT}_\lambda(\mathbf{w}^\star, \mathbf{w}^\star) - \frac{1}{2}\mathsf{OT}_\lambda(\alpha, \alpha).$$

For n large enough, the terms in the last two inequalities will be approximately equal to $\frac{1}{2}\{\mathsf{OT}_\lambda(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}) -$

$\mathsf{OT}_\lambda(\alpha,\alpha)\}$ and $\frac{1}{2}\{\mathsf{OT}_\lambda(\mathbf{w}^\star,\mathbf{a})-\mathsf{OT}_\lambda(\alpha,\alpha)\}$, respectively. Therefore,

$$\frac{1}{2}\{\mathsf{OT}_\lambda(\mathbf{a},\mathbf{a})-\mathsf{OT}_\lambda(\alpha,\alpha)\}\leq\frac{1}{2}\{\mathsf{OT}_\lambda(\mathbf{w}_{\mathsf{COT}},\mathbf{a})-\mathsf{OT}_\lambda(\alpha,\alpha)\}$$

$$\leq\frac{1}{2}\{\mathsf{OT}_\lambda(\mathbf{w}^\star,\mathbf{a})-\mathsf{OT}_\lambda(\alpha,\alpha)\}$$

Under Assumptions 3–4, Theorem 1 of Genevay et al. (2019) or Corollary 1 of Mena and Weed (2019) hold and $\mathbb{E}\{\mathsf{OT}_\lambda(\mathbf{a},\mathbf{a})-\mathsf{OT}_\lambda(\alpha,\alpha)\}=\mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$ and $\mathbb{E}\{\mathsf{OT}_\lambda(\mathbf{w}^\star,\mathbf{a})-\mathsf{OT}_\lambda(\alpha,\alpha)\}=\mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$. Thus, $\mathbb{E}\{\mathsf{OT}_\lambda(\mathbf{w}_{\mathsf{COT}},\mathbf{a})-\mathsf{OT}_\lambda(\alpha,\alpha)\}=\mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$.

$\square$

A.2 Proof of Theorem 2

Proof. Define $e_{z,i}=\mathbb{P}(Z_i=z\,|\,X_i,S_i=1)\,\mathbb{P}(S_i=1\,|\,X_i)/\mathbb{P}(S_i=0\,|\,X_i)$, or the inverse weight targeting the TATE. This is also the inverse of the Radon-Nikodym derivative. As a reminder, $\mu_z(X)=\mathbb{E}\{Y(z)\,|\,X\}$. We first decompose $\hat{\tau}-\tau$ into several residual terms:

$$\hat{\tau}-\tau=\sum_i w_i Z_i Y_i-\sum_i w_i(1-Z_i)Y_i-\tau$$

$$=\sum_i w_i Z_i\{Y_i-\mu_1(X_i)\}-\sum_i w_i(1-Z_i)\{Y_i-\mu_0(X_i)\}$$

$$+\sum_i w_i Z_i\mu_1(X_i)-\sum_i w_i(1-Z_i)\mu_0(X_i)-\tau$$

$$=\sum_i w_i Z_i\{Y_i-\mu_1(X_i)\}-\sum_i w_i(1-Z_i)\{Y_i-\mu_0(X_i)\}$$

$$+\sum_i w_i Z_i\mu_1(X_i)-\sum_i w_i(1-Z_i)\mu_0(X_i)-\tau$$

$$=\frac{1}{n}\sum_i\frac{Z_i}{e_{1,i}}\{Y_i-\mu_1(X_i)\}-\frac{1}{n}\sum_i\frac{1-Z_i}{e_{0,i}}\{Y_i-\mu_0(X_i)\}$$

$$+\sum_i\left(w_i-\frac{1}{n\cdot e_{1,i}}\right)Z_i\{Y_i-\mu_1(X_i)\}$$

$$-\sum_i\left(w_i-\frac{1}{n\cdot e_{0,i}}\right)(1-Z_i)\{Y_i-\mu_0(X_i)\}$$

$$+\sum_i w_i Z_i\mu_1(X_i)-\sum_i w_i(1-Z_i)\mu_0(X_i)$$

$$-\frac{1}{m}\left\{\sum_j\mu_1(X_j)-\sum_j\mu_0(X_j)\right\}$$

$$+\frac{1}{m}\left\{\sum_j\mu_1(X_j)-\sum_j\mu_0(X_j)\right\}-\tau$$

$$= A + B + C,$$

where

$$A = \frac{1}{n} \sum_i \frac{Z_i}{e_{1,i}} \{Y_i - \mu_1(X_i)\} - \frac{1}{n} \sum_i \frac{1 - Z_i}{e_{0,i}} \{Y_i - \mu_0(X_i)\}$$

$$+ \frac{1}{m} \left\{ \sum_j \mu_1(X_j) - \sum_j \mu_0(X_j) \right\} - \tau$$

$$B = \frac{1}{n} \sum_i \left(n \cdot w_i - \frac{1}{e_{1,i}} \right) Z_i \{Y_i - \mu_1(X_i)\}$$

$$- \frac{1}{n} \sum_i \left(n \cdot w_i - \frac{1}{e_{0,i}} \right) (1 - Z_i) \{Y_i - \mu_0(X_i)\}$$

$$C = \sum_i w_i Z_i \mu_1(X_i) - \sum_i w_i (1 - Z_i) \mu_0(X_i)$$

$$- \frac{1}{m} \left\{ \sum_j \mu_1(X_j) - \sum_j \mu_0(X_j) \right\}.$$

The goal is to show that both $n^{1/2}B$ and $n^{1/2}C$ are $o_p(1)$. Then, since A has the form of the semiparametrically efficient score function, the result follows.

First, for B, we have that $\lim_{n \to \infty} \mathbf{w} = \lim_{n \to \infty} \mathbf{w}^\star$ by Corollary 1. This also implies that $\lim_{n \to \infty} n \cdot \mathbf{w}_i = 1/e_i$ since by the Radon-Nikodym theorem, the Radon-Nikodym derivative is unique almost surely. To prove that $n^{-1/2}B$ goes to 0, it will be sufficient to prove that $\frac{\sqrt{n}}{n} \sum_i \left| \left(n \cdot w_i - \frac{1}{e_{z,i}} \right) \mathbb{I}(Z_i = z) \{Y_i - \mu_z(X_i)\} \right| \to 0$. We then have

$$\frac{\sqrt{n}}{n} \sum_i \left| \left(n \cdot w_i - \frac{1}{e_{z,i}} \right) \mathbb{I}(Z_i = z) \{Y_i - \mu_z(X_i)\} \right|$$

$$\leq \frac{\sqrt{n}}{n} \sum_i \left| \left(n \cdot w_i - \frac{1}{e_{z,i}} \right) \mathbb{I}(Z_i = z) \right| |Y_i - \mu_z(X_i)|$$

$$\leq \operatorname{ess.\,sup}_i \left| \left(n \cdot w_i - \frac{1}{e_{z,i}} \right) \mathbb{I}(Z_i = z) \right| \frac{1}{\sqrt{n}} \sum_i |Y_i - \mu_z(X_i)|,$$

where ess. sup is the essential supremum. The essential supremum quantity goes to 0 as a consequence of Corollary 1 while the residual quantity has finite expectation and variance by assumption. This implies $\frac{1}{\sqrt{n}} \sum_i |Y_i - \mu_z(X_i)| \xrightarrow{\mathcal{L}} L$ for some random variable L and that the desired quantity goes to 0 by Slutsky's theorem.

Second, for C, we have by assumption that $S_\lambda(\mathbf{w}, \mathbf{a}) = o_p(n^{-1/2})$. This implies that the empirical expectations also converge at a faster than $\sqrt{n}$-rate. Alternatively, for the basis function constraints, we have

$$|\mathbb{E}_{\mathbf{w}}\{\mathbb{I}(Z = z)\mu_z(X)\} - \mathbb{E}_{\mathbf{a}}\{\mu_z(X)\}| \leq \sum_k \delta_k |\gamma_k| \leq \|\delta\|_2^2 \|\gamma\|_2^2.$$

Thus, for each value z of Z, $\sqrt{n}\,|\,\mathbb{E}_{\mathbf{w}}\{\mathbb{I}(Z=z)\mu_z(X)\} - \mathbb{E}_{\mathbf{a}}\{\mu_z(X)\}| = o_p(1)$.

Finally, with Assumption 1, we can replace Y_i with $Y_i(Z_i)$. This means that by Assumption 1 and Assumption 5, $\mathbb{E}(A) = 0$ and $\lim_{n\to\infty} \mathbb{E}(\hat{\tau}) = \tau$. Thus, under Assumptions 1–5 and by the fact that A has the form of the semiparametrically efficient score function, A converges to τ at a $\sqrt{n}$-rate and $\sqrt{n}(\hat{\tau} - \tau)$ has the desired asymptotic distribution. $\qquad\square$

A.3   Proof of Theorem 3

Before we proceed to our proof, a slight digression is necessary to discuss how the barycentric projection in the setting of this paper will differ slightly from the usual formulation in optimal transport. Typically in optimal transport problems, we would include all of the available data in our distance metric; however, the missing potential outcomes make this inadvisable. Simply throwing in the observed outcomes into the optimal transport problem could lead to weights that bias treatment effects towards zero. Thus, we estimate the optimal transport plan only on the covariate data X and then incorporate the outcomes after estimation of the transport plan.

Further, we require the following definition.

Definition 1

The primal form of (5) with an entropy penalty for generic measures α and β on $\mathcal{X}$ is equivalent to

$$\text{OT}_\lambda(\alpha,\beta) = \inf_{\pi\in\mathbf{U}(\alpha,\beta)} \int_{\mathcal{X}\times\mathcal{X}} c(x,y)d\pi(x,y) + \lambda\log\left(\frac{\pi(x,y)}{\alpha(x)\otimes\beta(y)}\right)d\pi(x,y). \tag{11}$$

Eq. (11) has the dual form

$$\sup_{f,g} \int_{\mathcal{X}} f(x)d\alpha(x) + \int_{\mathcal{X}} g(y)d\beta(y) - \lambda\int_{\mathcal{X}\times\mathcal{X}} e^{(f(x)+g(y)-c(x,y))/\lambda}d\alpha(x)d\beta(y) + \lambda, \tag{12}$$

with a primal solution equal to $e^{(f(x)+g(y)-c(x,y))/\lambda}d\alpha(x)d\beta(y)$.

We note that the dual form is justified by Fenchel-Rockafellar duality but defer a proof to sources such as Peyré and Cuturi (2019).

Proof. Denote $\lim_{n\to\infty}\mathbf{w} = \omega$ and $\lim_{n\to\infty}\mathbf{a} = \alpha$. By assumption, $\omega = \alpha$. Without loss of essential generality, assume the weights only adjust one treatment group towards the full sample. We also assume both groups have equal sample sizes since it will make some of the notation easier to follow.

We will denote the optimal transportation plan between ω and α as π. The barycentric projection from α into ω is then $\int_{\mathcal{X}} x\,d\pi(x\,|\,x')$, where $d\pi(x\,|\,x') = d\pi(x,x')/d\alpha(x')$. Under the assumptions of the theorem, this transport plan, π is unique and is supported on the graph of a Monge map (Brenier, 1987). Moreover, the barycentric

projection will be the optimal map, T (Ambrosio et al., 2005, Lemma 12.2.3). Finally, the optimal transport plans in finite samples will converge to the limiting value: $\mathbf{P}_n \to \pi$ (Villani, 2008, Theorem 5.20). We also note that the Monge map T will be the identity function since the distributions are the same.

μ_z is like a Monge map for the outcomes. For ease of exposition, we can think of units from ω as having their outcome $Y(z)$ observed, while the outcomes are completely missing for units in α. Thus, we need some way of projecting an individual from α to ω and generating their hypothetical outcome. For an individual with covariate values x', this will be

$$\int \begin{bmatrix} y(z) \\ x \end{bmatrix} dp(y(z)\,|\,x) d\pi(x\,|\,x') = \int \begin{bmatrix} y \\ x \end{bmatrix} dp(y\,|\,x) d\pi(x\,|\,x') = \int \begin{bmatrix} \mu_z(x) \\ x \end{bmatrix} d\pi(x\,|\,x') = \begin{bmatrix} \mu_z(x') \\ x' \end{bmatrix},$$

where under Assumptions 1–2, the observed outcomes can be used for the potential outcomes and there is no interference.

This means the Monge map is $\begin{bmatrix} \mu_z(x') \\ x' \end{bmatrix}$ in our particular case when we map from a space without outcomes to a space with outcomes. The failure of the values of x' to change aligns with our assumption that the covariates precede treatment causally so their value does not change depending on the treatment group. Also, we note that since the outcomes are unknown, the best projection we can hope for is the conditional mean outcome function.

Given this fact and that we do not care about projecting the covariates, we will focus on the barycentric projection of just the outcome:

$$\bar{\pi}(x') = \int y\, dp(y\,|\,x) d\pi(x\,|\,x') = \int \mu_z(x) d\pi(x\,|\,x') = \mu_z(x') = T(x').$$

The estimated barycentric projection converges to μ_z. Denote $\bar{\mathbf{P}}^\lambda$ as the distribution implied by the barycentric projection from the entropy-penalized optimal transport problem and $\bar{\pi} = \pi(\cdot\,|\,x')$ as the optimal barycentric projection distribution. First, we rewrite our equation to make the barycentric projections explicit. Starting with the L_2 metric, we can represent Y as the convolution of its conditional mean, μ_z, with another random variable with mean zero, ϵ:

$$\left\| \int y\, d\bar{\mathbf{P}}^\lambda - \int \mu_z\, d\bar{\pi} \right\| = \left\| \int (\mu_z + \epsilon)\, d\bar{\mathbf{P}}^\lambda - \int \mu_z\, d\bar{\pi} \right\|$$

$$\leq \left\| \int \mu_z(\, d\bar{\mathbf{P}}^\lambda - d\bar{\pi}) \right\| + \left\| \int \epsilon\, d\bar{\mathbf{P}}^\lambda \right\|$$

$$\leq L \left\| \int x(\, d\bar{\mathbf{P}}^\lambda - d\bar{\pi}) \right\| + \left\| \int \epsilon\, d\bar{\mathbf{P}}^\lambda \right\|.$$

Then taking the square and expectation:

$$\mathbb{E}_\alpha \left\| \int y \, d\bar{\mathbf{P}}^\lambda - \int \mu_z \, d\bar{\pi} \right\|^2 \leq L^2 \, \mathbb{E}_\alpha \left\| \int x(d\bar{\mathbf{P}}^\lambda - d\bar{\pi}) \right\|^2$$

$$+ 2L \, \mathbb{E}_\alpha \left\| \int x(d\bar{\mathbf{P}}^\lambda - d\bar{\pi}) \right\| \left\| \int \epsilon \, d\bar{\mathbf{P}}^\lambda \right\|$$

$$+ \mathbb{E}_\alpha \left\| \int \epsilon \, d\bar{\mathbf{P}}^\lambda \right\|^2$$

$$\leq L^2 \, \underbrace{\mathbb{E}_\alpha \left\| \int \mu_z(d\bar{\mathbf{P}}^\lambda - d\bar{\pi}) \right\|^2}_{E}$$

$$+ 2L \, \underbrace{\sqrt{\mathbb{E}_\alpha \left\| \int \mu_z(d\bar{\mathbf{P}}^\lambda - d\bar{\pi}) \right\|^2} \sqrt{\mathbb{E}_\alpha \left\| \int \epsilon \, d\bar{\mathbf{P}}^\lambda \right\|^2}}_{F}$$

$$+ \underbrace{\mathbb{E}_\alpha \left\| \int \epsilon \, d\bar{\mathbf{P}}^\lambda \right\|^2}_{G}$$

For the term in G, we have

$$G = \mathbb{E}_\alpha \left(\sum_i \epsilon_i \bar{\mathbf{P}}_i^\lambda \right)^2$$

$$= \mathbb{E}_\alpha \left(\sum_i \sum_{i'} \epsilon_i \epsilon_{i'} \bar{\mathbf{P}}_i^\lambda \bar{\mathbf{P}}_{i'}^\lambda \right)$$

$$= \mathbb{E}_\alpha \, \mathbb{E} \left(\sum_i \sum_{i'} \epsilon_i \epsilon_{i'} \bar{\mathbf{P}}_i^\lambda \bar{\mathbf{P}}_{i'}^\lambda \mid X, X' \right)$$

$$\leq \mathbb{E}_\alpha \left(\sum_i \xi^2 (\bar{\mathbf{P}}_i^\lambda)^2 \right).$$

We then argue that this term goes to zero since $\bar{\mathbf{P}}_i^\lambda$ converges to $\bar{\pi}^\lambda$, which is without an atom. By the primal solution in Definition 1 and the structure of the penalty in of Eq. (5), we know $\bar{\mathbf{P}}_i^\lambda = e^{[f(x_i)+g(x')-c(x,x')]/\lambda} w_i < 1, \forall i$. Further, because $(\bar{\mathbf{P}}_i^\lambda)^2 < \bar{\mathbf{P}}_i^\lambda$, each value of the square is bounded away from one as well. Additionally, $\bar{\mathbf{P}}_i^\lambda$ will converge to zero for all i since π^λ has a density with respect to ω and α (Peyré and Cuturi, 2019). This also implies that the empirical mean and variance of $\bar{\mathbf{P}}_i^\lambda$ will also go to zero and, thus, $\sum_i (\bar{\mathbf{P}}_i^\lambda)^2 \to 0$.

The rate of this convergence will be determined by the rate of convergence of the regularized optimal transport problem,

$$K_d \lambda \left(1 + \frac{\sigma^{\lceil 5d/2 \rceil + 6}}{\lambda^{d'+3}} \right) n^{-1/2}$$

for a constant K_d depending only on the dimension d, σ a constant determined by the subgaussian tail of the

random variable, and $d' = \lceil 5d/4 \rceil$ (Mena and Weed, 2019, Corollary 1). As long as λ goes to zero slowly enough, the sum will converge to zero. This will be the case if we take

$$\lambda \asymp n^{-\frac{1}{2d'+9}}. \tag{13}$$

Lastly,

$$\mathbb{E}_\alpha \left(\sum_i \xi^2 (\overline{\mathbf{P}}_i^\lambda)^2 \right) \to 0$$

by Lebesgue's Dominated Convergence Theorem.

Turning next to the term in E, we have by Theorems 4 and 5 in Pooladian and Niles-Weed (2021) that

$$E \lesssim \lambda^{1-d''/2} \log(n) n^{-1/2} + \lambda^{(t+1)/2} + \lambda^2 I_0 + \lambda^{-1} (1 + \lambda^{1-d''/2}) \log(n) n^{-1/2},$$

where I_0 is the Fisher Information of α, $d'' = 2\lceil d/2 \rceil$, and $t \in [2,3]$ is related to the number of derivatives that exist for the optimal Kantorovich potential (Chizat et al., 2020; Pooladian and Niles-Weed, 2021). These theorems also require several assumptions which hold for our particular case. First, in our particular case where Brenier's Theorem holds, the derivatives of the dual potentials of the unpenalized problem equal the Monge map. These maps are then infinitely differentiable since in our case they are equal to the identity function. This also means $t = 3$. Also since we assume I_0 exists, we do not need the bounded densities required in Pooladian and Niles-Weed (2021). Then using the value of λ from (13), we get

$$E \lesssim (1 + I_0) n^{-\frac{2}{2d'+9}}.$$

Finally, to establish the rate is determined by E, we need for the term in F to go to zero faster than the term in E. To see this, we have that

$$\frac{2LF}{E} = \frac{2L\sqrt{E}\sqrt{G}}{E} = \frac{2L\sqrt{G}}{\sqrt{E}}.$$

Since the numerator goes to zero faster than the denominator, the overall rate is determined by E.

$\square$

A.4   Proof of Proposition 1

Proof. It will be sufficient to show the proof for the estimator of $\mathbb{E}\,Y(z)$ for one such group $Z = z$. The barycentric projection estimator of the mean in this group will be $n^{-1}\sum_i \hat{Y}_i(z)$. Expanding the term,

$$n^{-1}\sum_i \hat{Y}_i(z) = n^{-1}\sum_i n\sum_j Y_j\mathbb{I}(Z_j = z)\mathbf{P}_{ji}$$

$$= \sum_j Y_j\mathbb{I}(Z_j = z)\sum_i \mathbf{P}_{ji}$$

$$= \sum_j Y_j\mathbb{I}(Z_j = z)w_j,$$

where the first equality is by definition and the last equality is by the constraints of the optimal transport problem.

$\square$

A.5   Proof of Proposition 2

Proof. Under the assumptions of the proposition, the limiting optimal transport plan is unique and supported on the graph of a Monge map, the Monge map will be equal to the barycentric projection, and the empirical transport plans converge to the limiting value (Brenier, 1987; Ambrosio et al., 2005; Villani, 2008).

Let $\hat{T}(X_j) = \sum_{i:Z_i=z} X_i w_i$. We also take our data to be univariate for notational simplicity but the results are easily extended to data in $\mathbb{R}^d$. Then the objective for SCM can be written as

$$\text{SCM} = \inf_{\hat{T}} n^{-1}\sum_j \left\| \hat{T}(X_j) - X_j \right\|_2^2$$

$$= n^{-1}\sum_j \hat{T}(X_j)^2 - 2\hat{T}(X_j)X_j + X_j^2.$$

We can represent $\hat{T}(X_j)$ as a matrix of weights with certain constraints instead of the sum over the weights for each j:

$$\text{SCM} = \inf_{\tilde{\mathbf{P}}\in\mathbb{R}_+^{n_z\times n}:\tilde{\mathbf{P}}^\top\mathbf{1}=\mathbf{1}} n^{-1}\sum_j \left(\sum_{i:Z_i=z} X_i\tilde{\mathbf{P}}_{ij}\right)^2 - 2X_j\sum_{i:Z_i=z} X_i\tilde{\mathbf{P}}_{ij} + X_j^2$$

$$= \inf_{\tilde{\mathbf{P}}\in\mathbb{R}_+^{n_z\times n}:\tilde{\mathbf{P}}^\top\mathbf{1}=\mathbf{1}} n^{-1}\sum_j \sum_{i:Z_i=z} X_i\tilde{\mathbf{P}}_{ij}\tilde{X}_j - 2\sum_{i:Z_i=z} X_jX_i\tilde{\mathbf{P}}_{ij} + X_j^2$$

$$= \inf_{\mathbf{P}\in\mathbb{R}_+^{n_z\times n}:\mathbf{P}^\top\mathbf{1}=\mathbf{a}} \sum_{i:Z_i=z,j} X_i\tilde{X}_j\mathbf{P}_{ij} - 2\sum_{i:Z_i=z,j} X_jX_i\mathbf{P}_{ij} + \sum_{i:Z_i=z,j} X_j^2\mathbf{P}_{ij} \qquad (14)$$

The terms in $\tilde{X}_j$ are equivalent to a barycentric projection from $\mathbf{a}$ into $\mathbf{a}_z$. And we can see that the solution $\mathbf{P}_{ij}$ is

equivalent to an optimal transport plan with one of the margins allowed to vary. The corresponding problem under a Kantorovich relaxation is

$$\inf_{\mathbf{P}\in\mathbb{R}_+^{n_z\times n}:\mathbf{P}^\top\mathbf{1}=\mathbf{a}} \sum_{i:Z_i=z,j} X_i^2\mathbf{P}_{ij} - 2\sum_{i:Z_i=z,j} X_jX_i\mathbf{P}_{ij} + \sum_{i:Z_i=z,j} X_j^2\mathbf{P}_{ij}. \tag{15}$$

The marginal distributions $\mathbf{w}$ from both of these problems can be found as $\mathbf{w} = \mathbf{P1}$. Eqs. (14) and (15) will be equal in finite samples if $\sum_j \tilde{X}_j\mathbf{P}_{ij} = X_i \sum_j \mathbf{P}_{ij}$, which is not guaranteed. However, in the limit they will be the same, as we show next.

Under a similar argument to the proof of Theorem 1, OT $(\mathbf{w}_{NNM}, \mathbf{a}) \leq$ OT $(\mathbf{w}^\star, \mathbf{a}) \to 0$. Then, denoting the Monge problem using barycentric projections as M, M$(\mathbf{w}_{SCM}, \mathbf{a}) \leq$ M$(\mathbf{w}_{NNM}, \mathbf{a}) \to 0$. The last term holds since $\mathbf{w}_{NNM} \to \alpha$ and $\int x d\pi(x\,|\,x') = x'$.

Finally, to see the equivalence of the Monge and Kantorovich formulations for fixed margins $\alpha,\ \omega$ under the given assumptions:

$$\begin{aligned}
\inf_T \int |T(x') - x'|^2 d\alpha &= \inf_T \int T(x')^2 d\alpha + \int (x')^2 d\alpha - 2\int x'T(x')d\alpha \\
&= \inf_T \int x^2 dT_\#\alpha + \int (x')^2 d\alpha - 2\int xT(x)d\alpha \\
&= \int x^2 d\omega + \int (x')^2 d\alpha - \inf_\pi \int x' \left[\int x d\pi(x\,|\,x')\right] d\alpha \\
&= \int x^2 d\omega + \int (x')^2 d\alpha - \inf_\pi \int xx' d\pi(x,x') \\
&= \inf_{\pi\in\mathbf{U}(\omega,\alpha)} \int |x - x'|^2 d\pi(x,x')
\end{aligned}$$

$\square$

B   Further empirical studies

B.1   Empirical convergence

The data generating model in this section comes from the setting in Section 5 with high-overlap between covariate distributions and the estimand of interest is the ATE. We examine three measures of performance of the estimated weights, $\mathbf{w}$, at approximating 1) the target empirical distribution, $\mathbf{a}$, in terms of 2-Sinkhorn divergence, 2) the distribution of the self-normalized propensity score, $\mathbf{w}^\star$, in terms of the 2-Sinkhorn divergence, and 3) the difference between $\mathbf{w}$ and $\mathbf{w}^\star$ under an L_2 norm. We present averages across 1000 replications.

The comparator methods we use are: a Probit Generalized Linear Model (GLM), the true data-generating model;
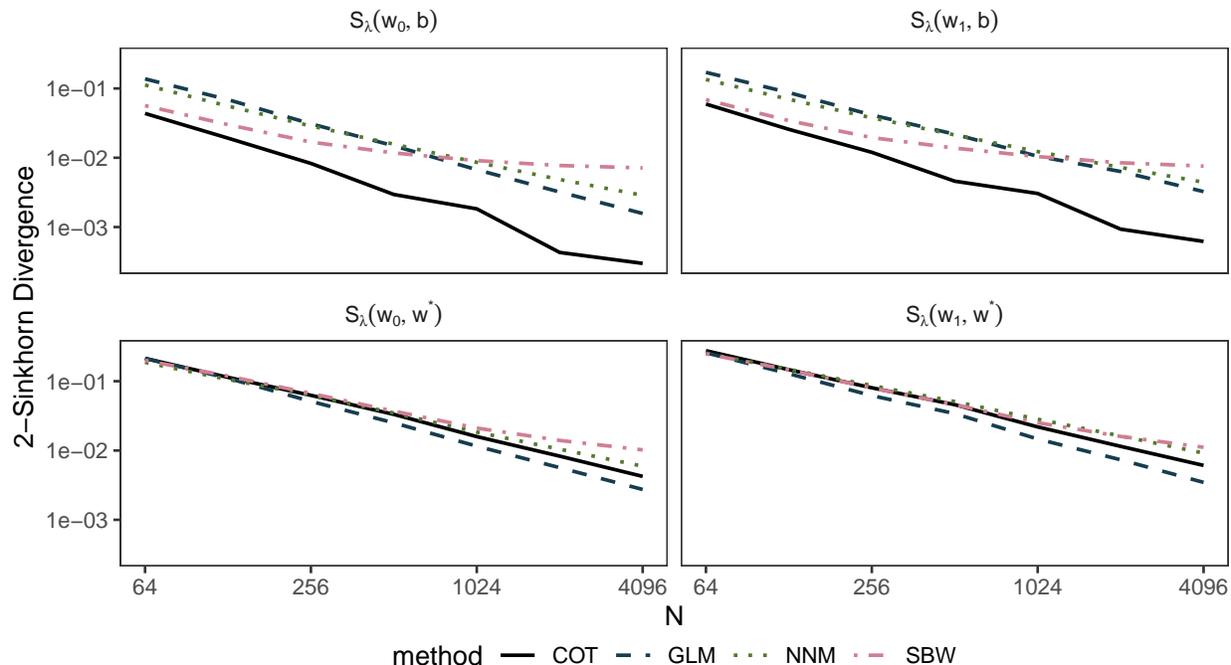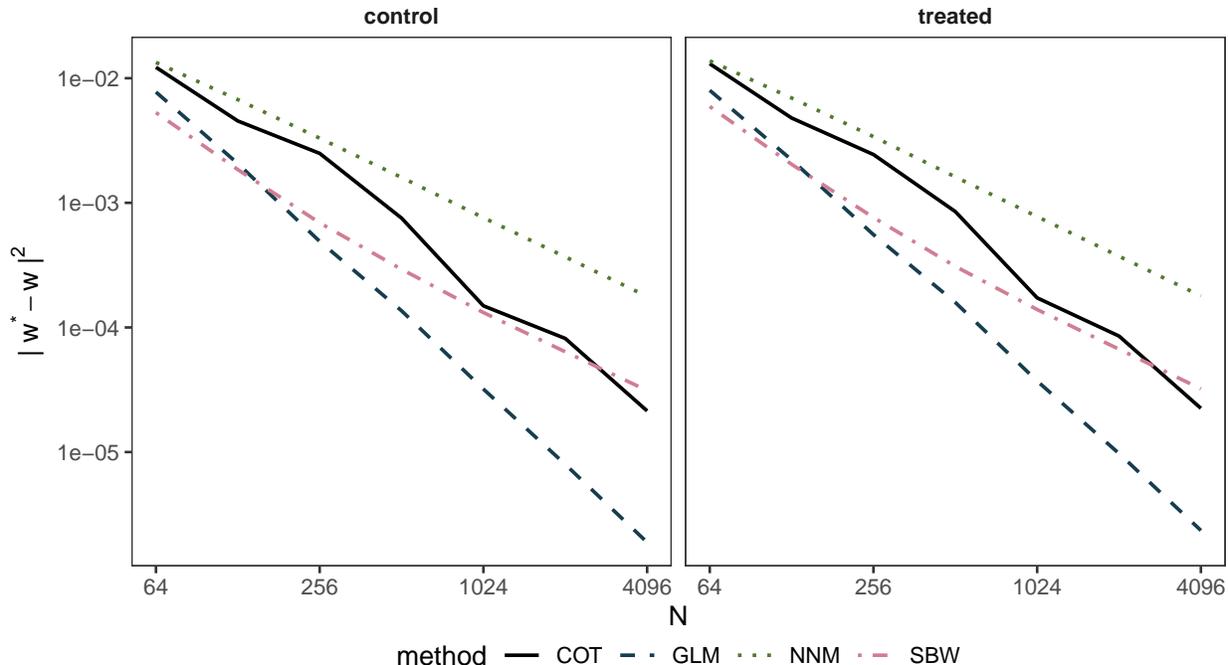
Figure 3: Convergence of the weights to the distributions specified by the empirical distributions (top) and the distributions specified by the true propensity score/Radon-Nikodym derivatives (bottom). Weights are a Causal Optimal Transport (COT), Nearest Neighbor Matching (NNM), a Probit model (GLM), and Stable Balancing Weights (SBW). Lines denote means across 1000 simulations. Both axes are on the log scale.

Stable Balancing Weights (SBW) using the correct propensity score covariate functions; and Nearest Neighbor Matching (NNM) with replacement, which is of course equivalent to an unpenalized COT (see Section 4). For NNM, we use a cost function that is equal to $\| \cdot \|_p^p$ with $p = d/2 + 1$ to meet the conditions of Theorem 1 in Fournier and Guillin (2015).

In Figure 3, we see that the COT weights do a better job of approximating the target distribution under the 2-Sinkhorn Divergence but that the GLM model does better at targeting the distribution implied by the true inverse propensity score under the same metric. Of note, SBW displays decaying rates convergence as the sample size increases and even performs worse than NNM for large sample sizes.

Figure 4 displays the convergence in L_2-norm for the various methods. As we would expect, the GLM model converges fastest to the values of the true inverse propensity score. The COT weights using the Sinkhorn divergence display slightly worse rates of convergence, on average, followed by NNM. SBW again displays a rate that decays with the sample size, though it does perform better than other methods when sample sizes are small.

Figure 4: Convergence of the estimated weights to the values of the true inverse propensity score in terms of the L_2 norm. Weights are a Causal Optimal Transport (COT), Nearest Neighbor Matching (NNM), a Probit model (GLM), and Stable Balancing Weights (SBW). Lines denote means across 1000 simulations. Both axes are on the log scale.
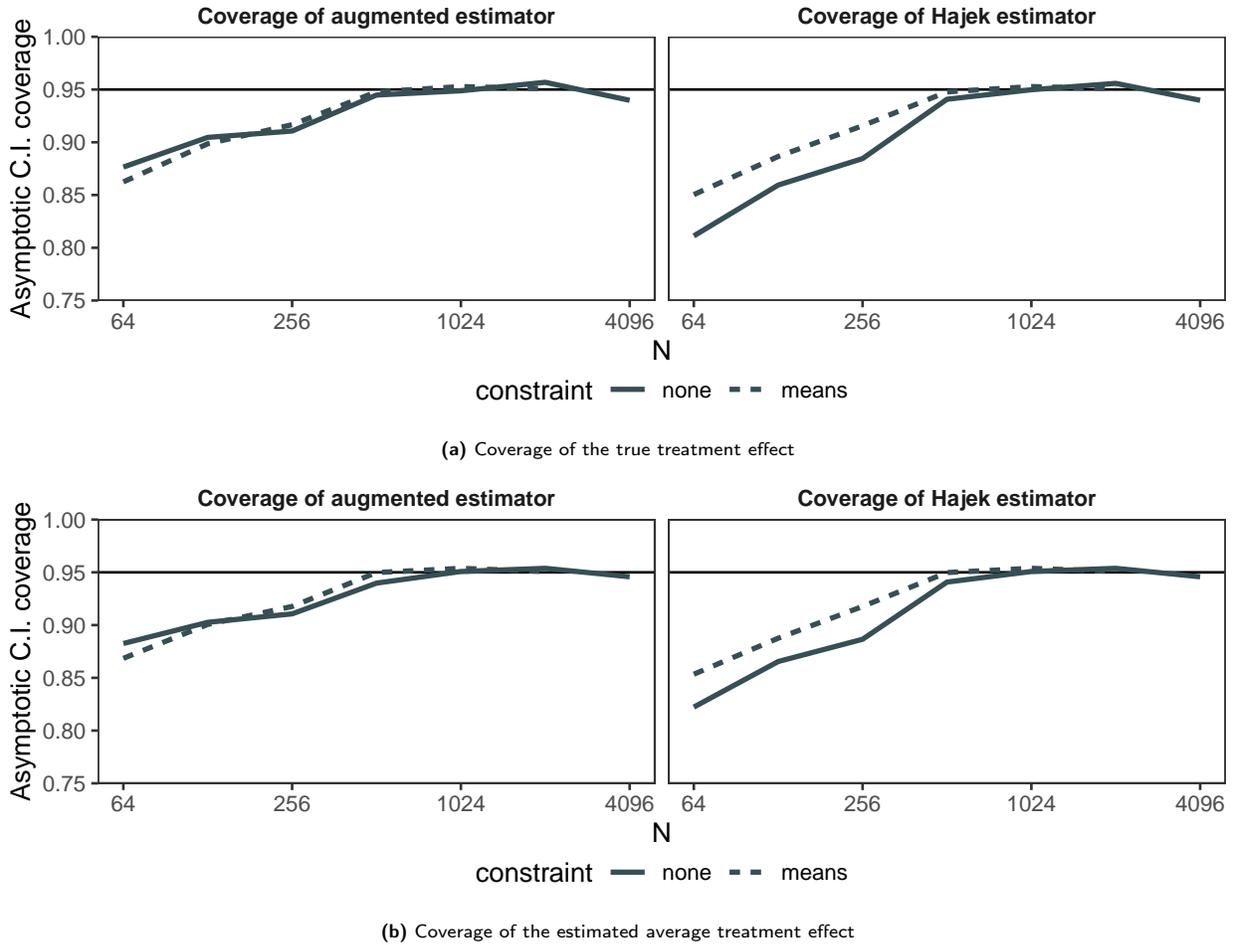
B.2 Empirical coverage of asymptotic confidence interval

The empirical coverage of the confidence interval is the focus of this subsection. We utilize the generating model from the setting in Section 5 with high-overlap between covariate distributions and a linear outcome model linear outcome model

$$Y(0) = Y(1) = X_1 + X_2 + X_3 - X_4 + X_5 + X_6 + \eta$$

with $\eta \sim \mathcal{N}(0, 1)$. The target distribution, **a**, is the full sample, making the estimand of interest the ATE. We run 1000 replications of our experiment.

Figure 5 displays results for increasing sample sizes. In the top part, Figure 5a, we examine the coverage of the true estimate of zero in a variety of settings. Amazingly, the COT method achieves well-calibrated confidence intervals without using an augmented estimator or mean constraints. Similar results are observed for the empirical expectations in Figure 5b. In both cases, the non-augmented balancing constraint method converges a bit faster than the non-augmented method without balancing constraints.

(a) Coverage of the true treatment effect

(b) Coverage of the estimated average treatment effect

Figure 5: Coverage of the asymptotic confidence interval for both the true effect (a) and the average estimated effect (b). Solid lines denote no basis function balancing ("none") and the balancing of the means of covariates ("means").

B.3   Tuning algorithm

In this section, we examine the performance of the tuning algorithm presented in Algorithm 1. We again use the setting of Section 5 with high-overlap between covariate distributions, use COT with an L_2 metric and no balancing functions, and for a variety sample sizes from 32 to 4096. The target distribution in this case is the full sample making the estimand the ATE. Performance is measured in terms of an Anderson-Darling statistic between the estimated weights, $\mathbf{w}$, and the self-normalized inverse propensity score, $\mathbf{w}^\star$:

$$\frac{1}{n} \sum_{i}^{n} \frac{(w_i - w_i^\star)^2}{w_i^\star (1 - w_i^\star)}.$$

We use this term rather than a simple L_2 norm because it will appropriately adjust for the discrepancy between weight vectors as the values become small. Finally, we run this experiment 1000 times.

Figure 6 displays the results for the tuning algorithm as the sample size increases. We can see that initially the

algorithm avoids the highest discrepancy area for intermediate values of the penalty parameter λ. As the sample size increases, the algorithm concentrates on λ values that minimize the difference between the true and estimated propensity scores. This holds true for both the treated and control observations.

C   Additional case study: the LaLonde Data

We also validate our method on the LaLonde data set (LaLonde, 1986).

## C.1   The National Supported Work Demonstration program

The original data come from a job training program called the National Supported Work Demonstration program (NSW) in which people were randomized to receive or not receive training from the program in the year 1976. The outcome of interest was then to look at the difference in incomes between the treatment and control groups in 1978. The original experimental estimate was a difference of $1,794 with a confidence interval of ($551, $3038). The variables available in the original study include 10 pre-intervention characteristics: earnings and employment in 1974 and 1975, years of education, whether the person received a high school degree, marital status, and indicators for black or Hispanic ethnicity.

C.2   LaLonde's modification

LaLonde then proceeded to modify the original study data by removing the control group and seeing if he could recover the original treatment effect by utilizing an observational data sample taken from the Current Population Survey (CPS) with the same variables measured. This gives 185 participants from the NSM in the treated group and 15,992 non-participants from the CPS in the control group.

C.3   Methods

From the Causal Optimal Transport weighting methods, we include no constraints ("none") and mean constraints ("means"). Hyperparameters were tuned with the algorithm detailed in Algorithm 1. The distance metric is an L_2 metric on the binary covariates and a Mahalanobis L_2 metric on the continuous covariates. We consider the Háyek estimator in (2), a doubly robust/augmented estimator using linear regression with linear terms of the covariates, a weighted least squares estimator, and the barycentric projection estimator of Eq. (7) utilizing an assignment matrix $\mathbf{P}$ constructed utilizing an L_1 cost.
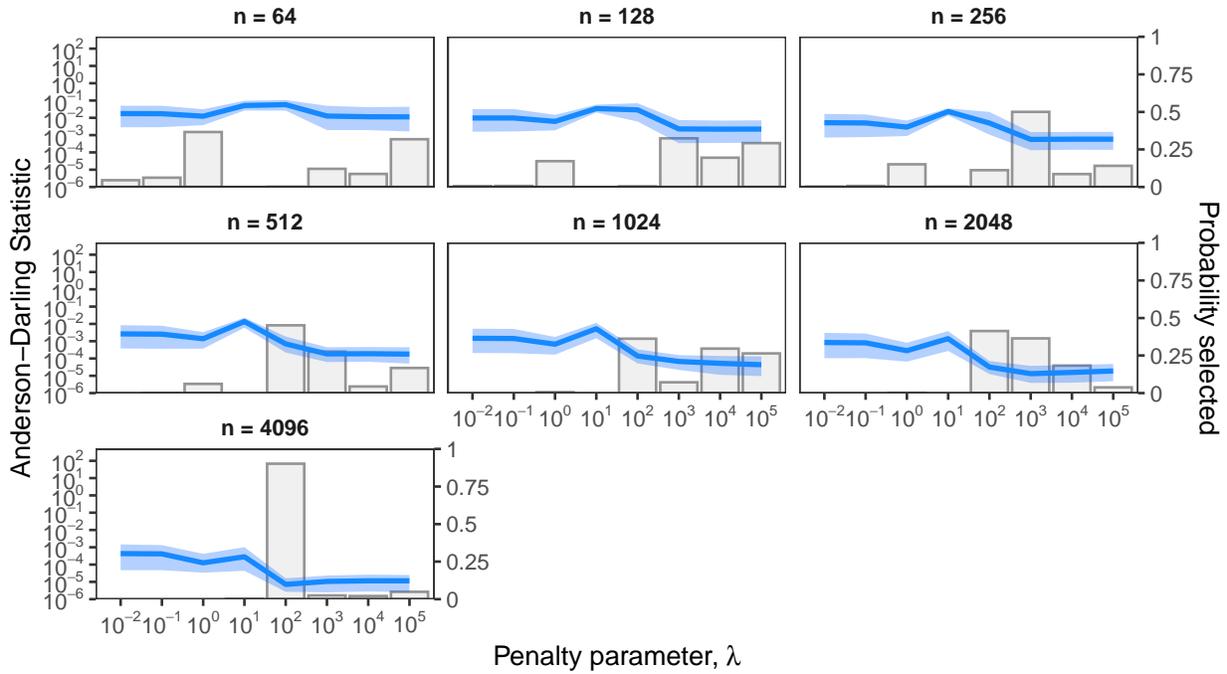
(a) Selection of penalty parameter, λ, for the treated
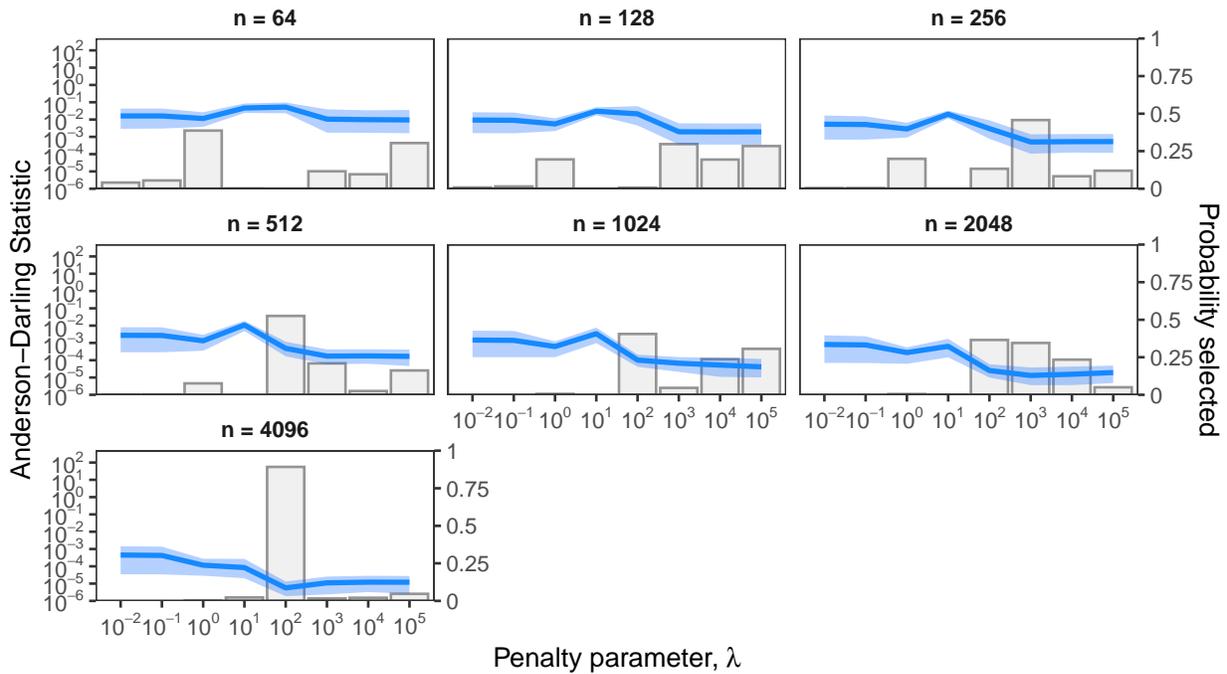
(b) Selection of the penalty parameter, λ, for the control

Figure 6: Performance of the tuning algorithm in Algorithm 1. Discrepancy between the weights estimated by COT and the true self-normalized inverse propensity score in terms of an Anderson Darling statistic in blue. The probability that a λ value was selected is given by the histogram.

C.4 Design diagnostics

We now display the before and after weighting balance in variable means and 2-Sinkhorn divergence to give a sense of distributional balance. We can see that for all weighting methods both means and distributions are much more similar after weighting than before (Figure 7).

C.5 Results

In Table 3 and Figure 8, we see that we are able to get very close to the original effects for the Hájek, Augmented, and weighted least squares approaches. The barycentric projection estimators have a notable upward bias but still have confidence intervals covering the true effect.

	Hajek	Augmented	Weighted OLS	Barycentric Projection
COT	1791 (649, 2932)	1791 (650, 2932)	1791 (418, 3164)	2435 (1287, 3583)
COT, means	1816 (675, 2957)	1816 (675, 2957)	1816 (532, 3100)	2390 (1243, 3538)

Table 3: Results for treatment effect estimation for the National Work Support demonstration treated group and the weighted set of controls from the Current Population Survey. The estimate is the difference in 1978 earnings in dollars between the two groups. Values are estimates with asymptotic 95% confidence intervals.

D   Other versions of Causal Optimal Transport

We can also represent COT using Eq. (5):

$$\mathrm{COT}_\lambda\left(\mathbf{a}\right) = \min_{\mathbf{w}\in\Delta_n} \mathrm{OT}_\lambda\left(\mathbf{w}_1, \mathbf{a}\right) + \mathrm{OT}_\lambda\left(\mathbf{w}_0, \mathbf{a}\right). \tag{16}$$

In these equations, we can either use an entropy or an L_2 penalty and can also incorporate balancing constraints. We can, of course, show that Theorems 1 and 2 hold.

D.1 Other versions of COT also converge

Starting with the proof of convergence, we need the following additional assumption:

Assumption 6

For $\mathrm{COT}_\lambda\left(\mathbf{a}\right)$ *in Eq* (16) *with an entropy penalty:* $\lambda \rightarrow 0$ *as* $n \rightarrow \infty$

Proof. We begin by proving the L_2 regularized weights converge, then the entropically regularized weights, and finally, the Sinkhorn divergence. We also have that under Assumption 2, $\mathbf{w}^\star$ exists. Then by Lemma 1, $\mathbf{w}^\star \rightharpoonup \alpha$.
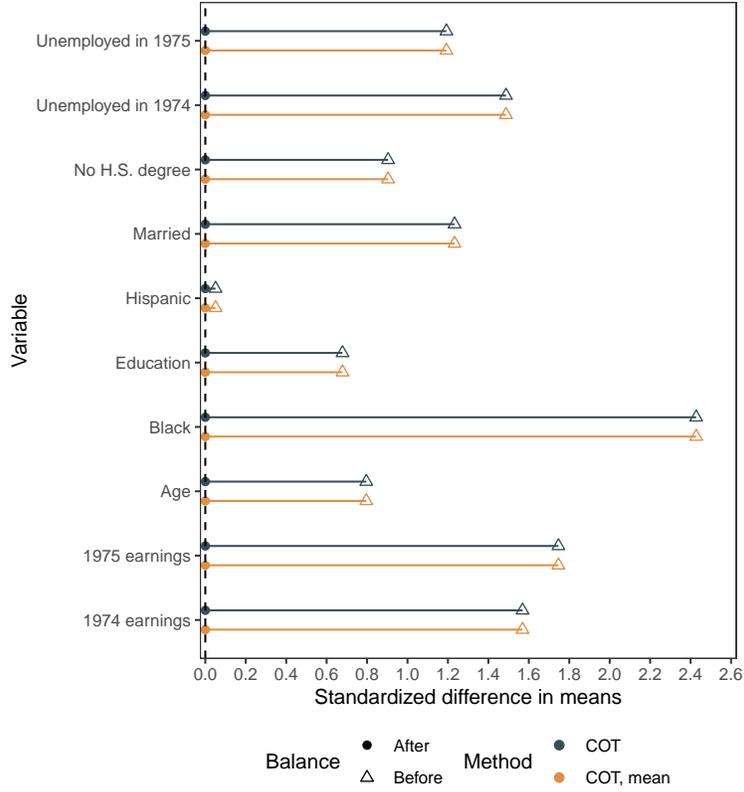
(a) Change in the standardized difference in means between the two groups before and after weighting
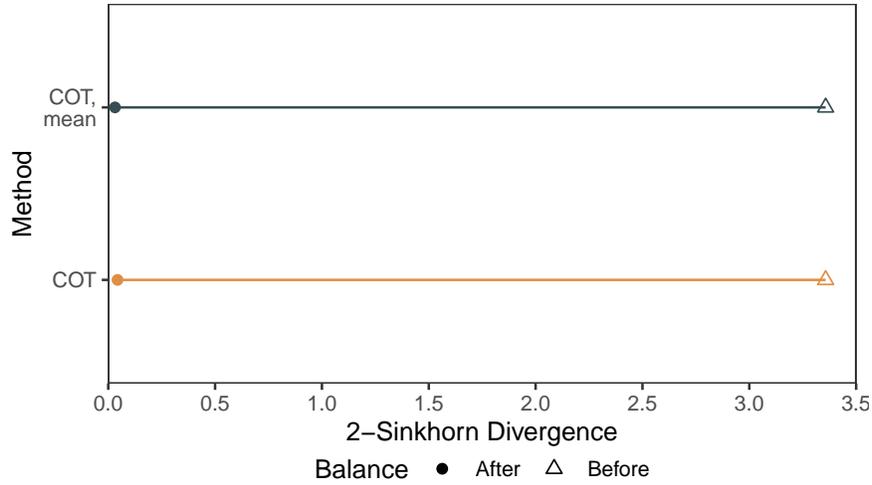
(b) Change in the 2-Sinkhorn divergence between the two groups before and after weighting

Figure 7: An examination in the change in balance before and after utilizing the optimal transport methods with the listed constraints for the LaLonde data. "COT" corresponds to no constraints and "COT, mean" corresponds to constraints on the mean balance between distributions.

L$_2$ penalization. Theorem 1 of Blondel et al. (2018) give bounds on $\mathrm{OT}_\lambda\left(\mathbf{a}_z, \mathbf{a}\right)$:

$$\frac{\lambda}{2} \sum_{i:Z_i=z,j} \left(\frac{\mathbf{a}_{z,i}}{n} + \frac{\mathbf{a}_j}{n} - \frac{1}{n^2}\right)^2 \leq \mathrm{OT}_\lambda\left(\mathbf{w}^\star, \mathbf{a}\right) - \mathrm{OT}\left(\mathbf{a}_z, \mathbf{a}\right) \leq \frac{\lambda}{2} \min\{\|\mathbf{a}_z\|^2, \|\mathbf{a}\|^2\}.. \tag{17}$$
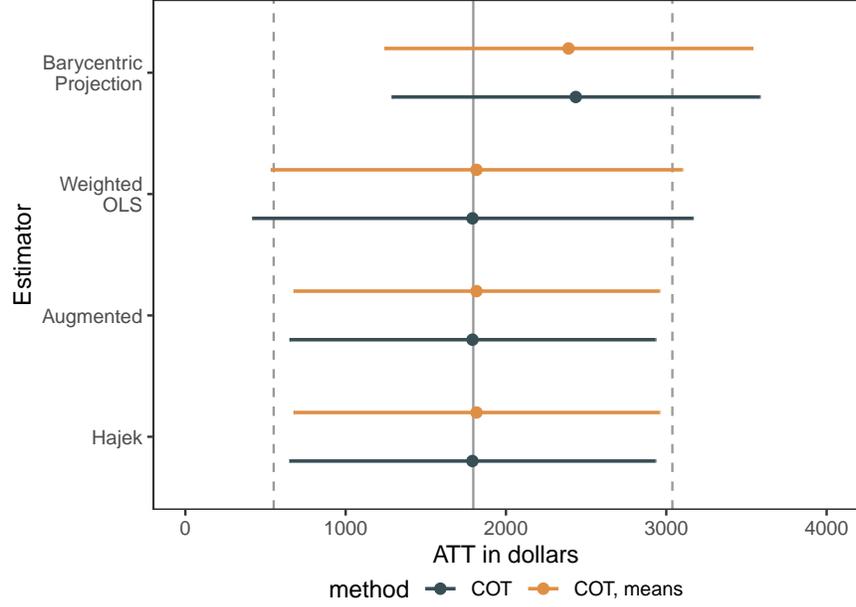
Figure 8: Results for treatment effect estimation for the National Work Support demonstration treated group and the weighted set of controls from the Current Population Survey. The estimate is the difference in 1978 earnings in dollars between the two groups targeting the average treatment effect of the treated (ATT). We see that all optimal transport methods and estimators displayed are able to get close to the original treatment effect. Note that "COT" corresponds to no constraints and "COT, mean" corresponds to constraints on the mean balance between distributions.

Then the upper bounds on the L_2 regularized problem for the importance sampling weights are

$$\mathsf{OT}_\lambda\left(\mathbf{a}_z, \mathbf{a}\right) \leq \mathsf{OT}\left(\mathbf{w}^\star, \mathbf{a}\right) + \frac{\lambda}{2}\|\mathbf{a}\|,$$

where the first inequality follows from rearrangement of Eq. (17) and the fact that $\min\{\|\mathbf{w}^\star\|^2, \|\mathbf{a}\|^2\}$ is minimized by the measure where all the observations have the same weight. Also,

$$\|\mathbf{a}\|^2 = \sum_j \mathbf{a}_j^2 = \frac{1}{n} \to 0.$$

Thus, $\mathsf{OT}_\lambda\left(\mathbf{w}^\star, \mathbf{a}\right) \to \mathsf{OT}\left(\mathbf{w}^\star, \mathbf{a}\right)$ and by Corollary 6.9 of Villani (2008), $\mathsf{OT}\left(\mathbf{w}^\star, \mathbf{a}\right) \to 0$.

Now we turn directly to the Causal Optimal Transport weights. The problem is convex (Blondel et al., 2018), which means that

$$\mathsf{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) \leq \mathsf{OT}_\lambda\left(\mathbf{c}, \mathbf{a}\right)$$

for all $\mathbf{c} \in \Delta_n$ that satisfy the constraints of the problem. Further, by assumption $\exists n > 0$ such that the importance sampling weights $\mathbf{w}^\star$ also satisfy the balancing constraints. This means that

$$\mathsf{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) \leq \mathsf{OT}_\lambda\left(\mathbf{w}^\star, \mathbf{a}\right)$$

and both quantities also satisfy the problem constraints for some n.

Finally, if $\mathrm{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right)$ goes to 0, this will mean $\mathbf{w}_{\mathsf{COT}} \rightharpoonup \alpha$ since $\mathrm{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) \rightarrow \mathrm{OT}\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right)$. Thus, since $\mathrm{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) \rightarrow 0$ because $\mathrm{OT}_\lambda\left(\mathbf{w}^\star, \mathbf{a}\right)$ goes to 0, by Corollary 6.9 in Villani (2008)

$$\mathbf{w}_{\mathsf{COT}} \rightharpoonup \alpha.$$

Entropy penalization. The entropy penalized Causal Optimal Transport problem is also a convex problem, which allows us to conclude

$$\mathrm{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) \leq \mathrm{OT}_\lambda\left(\mathbf{c}, \mathbf{a}\right)$$

for $\forall \mathbf{c} \in \Delta_n$ since $\mathbf{w}_{\mathsf{COT}}$ minimizes this loss. This gives us the bound

$$0 \leq \mathrm{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) \leq \mathrm{OT}_\lambda\left(\mathbf{w}^\star, \mathbf{a}\right).$$

Since the entropy penalized optimal transport problem does not metrize weak convergence, we require that $\lambda \rightarrow 0$.

As $\lambda \rightarrow 0$ (by assumption) and $n \rightarrow \infty$,

$$\mathrm{OT}_\lambda\left(\mathbf{w}^\star, \mathbf{a}\right) \rightarrow 0$$

since $\mathbf{w}^\star \rightharpoonup \alpha$. This implies that

$$\mathrm{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) \rightarrow 0,$$

which implies that $\mathbf{w}_{\mathsf{COT}} \rightharpoonup \alpha$ by Corollary 6.9 in Villani (2008) since at $\lambda = 0$, $\mathrm{OT}_\lambda = \mathrm{OT}$.

$\square$

D.2   Convergence happens at a $\sqrt{n}$-rate

Then semiparametric efficiency also holds with the following additional assumption:

Assumption 7

For L_2 penalized weights, $c(\cdot, \cdot) = d_{\mathcal{X}}(\cdot, \cdot)^p$, with $p > d/2$ and $\mathbb{E}\left|X\right|^q < \infty$ for $q > 2p$.

Proof. First, $\mathbf{w}^\star$ exist under Assumption 2 and $\mathbf{w}^\star \rightharpoonup \alpha$ by Lemma 1. Also, under Assumption 3, Theorem 1 holds and $\mathbf{w}_{\mathsf{COT}} \rightharpoonup \alpha$.

L_2 regularization. Theorem 1 of Blondel et al. (2018) give bounds on $\mathrm{OT}_\lambda\left(\mathbf{a}_z, \mathbf{a}\right)$:

$$\frac{\lambda}{2} \sum_{i: Z_i = z, j} \left(\frac{\mathbf{a}_{z,i}}{n} + \frac{\mathbf{a}_j}{n} - \frac{1}{n^2}\right)^2 \leq \mathrm{OT}_\lambda\left(\mathbf{a}_z, \mathbf{a}\right) - \mathrm{OT}\left(\mathbf{a}_z, \mathbf{a}\right) \leq \frac{\lambda}{2} \min\{\|\mathbf{a}_z\|^2, \|\mathbf{a}\|^2\}.$$

This implies that regularized problem converges at a linear rate to the unregularized problem because $\min\{\|\mathbf{a}_z\|^2, \|\mathbf{a}\|^2\} = \|\mathbf{a}\|^2$ because under an iid assumption $a_j = 1/n, \forall j$. Therefore,

$$\mathrm{OT}_\lambda\left(\mathbf{a}_z, \mathbf{a}\right) \leq \mathrm{OT}\left(\mathbf{a}_z, \mathbf{a}\right) + \frac{\lambda}{2}\|\mathbf{a}\|^2.$$

Then

$$\|\mathbf{a}\|_2^2 = \sum_j a_j^2 = \frac{1}{n} \to 0.$$

This also implies that

$$\lim_{n\to\infty} \mathrm{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) = \mathrm{OT}_\lambda\left(\alpha, \alpha\right) = \mathrm{OT}\left(\alpha, \alpha\right)$$

since $\mathbf{w}_{\mathsf{COT}} \rightharpoonup \alpha$.

Further, using Theorem 1 in Fournier and Guillin (2015) we have that under Assumption 7,

$$\mathbb{E}\{\mathrm{OT}\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right)\} \leq \mathbb{E}\{\mathrm{OT}\left(\mathbf{w}_{\mathsf{COT}}, \alpha\right) + \mathrm{OT}\left(\mathbf{a}, \alpha\right)\} = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right).$$

Then,

$$\mathbb{E}\left\{\mathrm{OT}\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) - \mathrm{OT}_\lambda\left(\alpha, \alpha\right)\right\} = \mathbb{E}\left\{\mathrm{OT}\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right)\right\} = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right).$$

Entropy regularization. First, by convexity

$$0 \leq \mathrm{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) \leq \mathrm{OT}_\lambda\left(\mathbf{w}^\star, \mathbf{a}\right)$$

and

$$0 \leq \mathrm{OT}_\lambda\left(\mathbf{a}, \mathbf{a}\right) \leq \mathrm{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right).$$

Also,

$$\mathrm{OT}_\lambda\left(\mathbf{a}, \mathbf{a}\right) - \mathrm{OT}_\lambda\left(\alpha, \alpha\right) \leq \mathrm{OT}_\lambda\left(\mathbf{w}_{\mathsf{COT}}, \mathbf{a}\right) - \mathrm{OT}_\lambda\left(\alpha, \alpha\right) \leq \mathrm{OT}_\lambda\left(\mathbf{w}^\star, \mathbf{a}\right) - \mathrm{OT}_\lambda\left(\alpha, \alpha\right).$$

Then with Assumptions 2–4, the conditions of either Theorem 3 of Genevay et al. (2019) or Corollary 1 of Mena and Weed (2019) hold. This means that

$$\mathbb{E}\left\{\mathrm{OT}_\lambda\left(\mathbf{w}^\star, \mathbf{a}\right) - \mathrm{OT}_\lambda\left(\alpha, \alpha\right)\right\} = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right)$$

and

$$\mathbb{E}\left\{\mathrm{OT}_\lambda\left(\mathbf{a}, \mathbf{a}\right) - \mathrm{OT}_\lambda\left(\alpha, \alpha\right)\right\} = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right).$$

Thus,

$$\mathbb{E}\left\{\mathrm{OT}_\lambda\left(\mathbf{w}_{\mathrm{COT}}, \mathbf{a}\right) - \mathrm{OT}_\lambda\left(\alpha, \alpha\right)\right\} = \mathcal{O}\left(\frac{1}{\sqrt{n}}\right).$$

□

D.3 Dual formulation

Finally, we also have the following dual form for this problem:

Theorem 4

The dual of each term in Eq. (16) *is*

$$\max_{g,\xi} \; g^\top \mathbf{a} - \sum_k \delta_k |\xi_k| - \frac{1}{m} \sum_j \xi^\top B(X_j) - \sum_{i,j'} \mathbb{I}(Z_i = z) H_\lambda^* \left(g_{j'} - \xi^\top B(X_{j'}) - \mathbf{C}_{ij'}\right), \tag{18}$$

where $B(X) = \left(\begin{matrix} B_1(X) & \cdots & B_K(X) \end{matrix}\right)^\top$ *and* H_λ^* *is the convex conjugate of the penalty function* H_λ.

Note that the convex conjugate of $H_\lambda(x) = \lambda x \log(x)$ is $H_\lambda^*(y) = \exp\{(y - 1)/\lambda\}$ and the convex conjugate of $H_\lambda(x) = \lambda x^2/2$ is $H_\lambda^*(y) = y^2/(2\lambda)$.

Proof. We present the proof of the dual form provided in Theorem 4. First, some tools from convex analysis (Boyd and Vandenberghe, 2004).

Strong duality. If strong duality holds then the value of the primal objective at the optimal primal solution is equal to the dual objective at the optimal dual solution.

Slater's conditions. Slater's conditions are that the objective function is convex and only has equality and inequality constraints .

Slater's theorem. If Slater's condition's hold, then strong duality holds.

We are now ready to proceed. As a reminder, the primal optimization problem is

$$\text{COT}_{\lambda,z}(\mathbf{a}) = \underset{\mathbf{P} \geq 0}{\text{argmin}} \quad \sum_{i,j} C_{i,j} P_{i,j} + \lambda \frac{1}{2} P_{i,j}^2$$

$$\text{subject to } \sum_{i,j} P_{i,j} \mathbb{I}(Z_i = z) = 1$$

$$\mathbf{P}^\top \mathbf{1}_n = \mathbf{a}$$

$$\left| \sum_{i,j} B_k(X_i) P_{i,j} - \frac{1}{n} \sum_{j'} B_k(X_{j'}) \right| \leq \delta_k, \quad \forall k \in \{1, ..., K\}.$$

We first note that we can separate the basis function constraint into the following two inequality constraints

$$\sum_{i,j} B_k(X_i) P_{i,j} - \frac{1}{n} \sum_{j'} B_k(X_{j'}) < \delta_k,$$

$$-\sum_{i,j} B_k(X_i) P_{i,j} + \frac{1}{n} \sum_{j'} B_k(X_{j'}) < \delta_k.$$

Further we combine the k basis function upper bounds into one vector δ and similarly denote $B(X)$ as a $n \times k$ matrix of the basis function constraints and $\overline{B}$ as the average of the basis functions in the target population: $\frac{1}{m} \sum_{j'} B_k(X_{j'})$. Then we re-write the primal problem in its Lagrangian form, defining $\langle x, y \rangle = \text{tr}(x^\top y)$.

$$\mathcal{L} = \underset{\mathbf{P} \geq 0}{\min} \underset{g, \xi_L, \xi_U}{\max} \langle \mathbf{C}, \mathbf{P} \rangle + \frac{\lambda}{2} \langle \mathbf{P}, \mathbf{P} \rangle - \langle g, \mathbf{P}^\top \mathbf{1}_n - \mathbf{a} \rangle +$$

$$\langle \xi_U, B(X)^\top \mathbf{P} \mathbf{1}_n - \delta - \overline{B} \rangle + \langle \xi_L, -B(X)^\top \mathbf{P} \mathbf{1}_n - \delta + \overline{B} \rangle.$$

Because the primal problem contains only equality and inequality constraints and the primal objective is a convex function, then strong duality holds,

$$= \underset{g, \xi_L, \xi_U}{\max} \underset{\mathbf{P} \geq 0}{\min} \langle \mathbf{C}, \mathbf{P} \rangle + \frac{\lambda}{2} \langle \mathbf{P}, \mathbf{P} \rangle - \langle g, \mathbf{P}^\top \mathbf{1}_n - \mathbf{a} \rangle +$$

$$\langle \xi_U, B(X)^\top \mathbf{P} \mathbf{1}_n - \delta - \overline{B} \rangle + \langle \xi_L, -B(X)^\top \mathbf{P} \mathbf{1}_n - \delta + \overline{B} \rangle$$

$$= \underset{g, \xi_L, \xi_U}{\max} g^\top \mathbf{a} - (\xi_U + \xi_L)^\top \delta - (\xi_U - \xi_L)^\top \overline{B} +$$

$$\underset{\mathbf{P} \geq 0}{\min} \langle \mathbf{C}, \mathbf{P} \rangle + \frac{\lambda}{2} \langle \mathbf{P}, \mathbf{P} \rangle - g^\top \mathbf{P}^\top \mathbf{1}_n + (\xi_U - \xi_L)^\top B(X)^\top \mathbf{P} \mathbf{1}_n$$

$$= \underset{g, \xi}{\max} g^\top \mathbf{a} - \delta \|\xi\|_1 - \xi^\top \overline{B} +$$

$$\underset{\mathbf{P} \geq 0}{\min} \langle \mathbf{C}, \mathbf{P} \rangle + \frac{\lambda}{2} \langle \mathbf{P}, \mathbf{P} \rangle - g^\top \mathbf{P}^\top \mathbf{1}_n + \xi^\top B(X)^\top \mathbf{P} \mathbf{1}_n$$

where we have combined the two positive dual variables ξ_U and ξ_L into one unconstrained variable. Then we can rearrange to get

$$
\begin{aligned}
= \max_{g,\xi} \; & g^\top \mathbf{a} - \delta \|\xi\|_1 - \xi^\top \overline{B} + \\
& \min_{\mathbf{P} \geq 0} \langle \mathbf{C} - \mathbf{1}_n g^\top + B(X)\xi \mathbf{1}_n^\top, \mathbf{P} \rangle + \frac{\lambda}{2} \langle \mathbf{P}, \mathbf{P} \rangle.
\end{aligned}
$$

Then taking the derivative with respect to $\mathbf{P}$,

$$
\nabla_\mathbf{P} \mathcal{L} = \mathbf{C} - \mathbf{1}_n g^\top + B(X)\xi \mathbf{1}_n^\top + \lambda \mathbf{P}
$$

and finding the critical point of the gradient gives

$$
\mathbf{P} = \frac{1}{\lambda} \left(\mathbf{1}_n g^\top - B(X)\xi \mathbf{1}_n^\top - \mathbf{C} \right)_+,
$$

where the function $(x)_+ = \max(0, x)$ applied element-wise ensures that the weights are constrained to be positive. Plugging this back in, the objective is now

$$
\mathcal{L} = \max_{g,\xi} g^\top \mathbf{a} - \delta \|\xi\|_1 - \xi^\top \overline{B} - \frac{1}{2\lambda} \left(\mathbf{1}_n g^\top - B(X)\xi \mathbf{1}_n^\top - \mathbf{C} \right)_+^2,
$$

as desired. □