
Outcome Assumptions and Duality Theory for Balancing Weights

David Bruns-Smith
UC Berkeley

Avi Feller
UC Berkeley

Abstract

We study balancing weight estimators, which reweight outcomes from a source population to estimate missing outcomes in a target population. These estimators minimize the worst-case error by making an assumption about the outcome model. In this paper, we show that this outcome assumption has two immediate implications. First, we can replace the minimax optimization problem for balancing weights with a simple convex loss over the assumed outcome function class. Second, we can replace the commonly-made overlap assumption with a more appropriate quantitative measure, the minimum worst-case bias. Finally, we show conditions under which the weights remain robust when our assumptions on the outcomes are wrong.

1 Introduction

Using covariates to transfer outcome information from one setting to another is a central task in domain adaptation, observational causal inference, and missing data imputation. These tasks share a common structure: we observe covariates and outcomes for a source data set and want to predict outcomes given covariates in a target data set, which might have a different covariate distribution than the source. One standard approach is to reweight the source distribution to have a similar covariate distribution to the target. When the source and target distributions have common support, using the density ratio for weights leads to unbiased estimation, known as *importance weighting* for domain adaptation under covariate shift (Sugiyama et al., 2007a) and *inverse propensity score weighting* (IPW) for observational causal inference (Rosenbaum and Rubin, 1983).

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

Importance weighting has several drawbacks. First, using the density ratio for weights can lead to extremely large variance and unstable estimation (Kang et al., 2007; Cortes et al., 2010). Second, the density ratio is notoriously difficult to estimate, and simple plug-in estimates do not guarantee covariate balance between the reweighted source distribution and target distributions (see Ben-Michael et al., 2021).

Due to these drawbacks, in practice we would like to use weights with smaller variance than the density ratio that directly target a specified level of covariate balance (e.g., Gretton et al., 2009; Imai and Ratkovic, 2014). In general, such a bias-variance trade-off only exists if we assume restrictions on the outcome model; without restrictions, only the density ratio can guarantee finite bias. This motivates the so-called *minimax balancing weights* estimators, which we study in this paper. These estimators find the minimum dispersion weights that constrain the worst-case bias between groups over an outcome function class (see Zubizarreta, 2015; Hirshberg et al., 2019; Zhao, 2019; Kallus, 2020b).

1.1 Summary of the Paper

We begin by reviewing existing balancing weights estimators, which achieve a smaller mean squared error than importance weighting by introducing an assumption on the outcome model. We argue that the outcome assumption implies two new results.

First, we use convex duality to show that the minimax optimization problem for balancing weights can be replaced with a simple convex loss over the outcome function class. Our dual formulation in Section 4 shows that the minimax weights are always a (rescaled and recentered) function from that class. For example, if the outcomes are bounded, the corresponding weights will be bounded. If the outcome function belongs to an RKHS with some kernel, the corresponding weights will belong to an RKHS with the same kernel. The outcome assumption pins down the shape of the balancing weights.

Second, we show that after making an outcome assumption, we do not need the density ratio to exist,

i.e., we do not need to make the additional “overlap” assumption that is common in causal inference. Instead, there is an explicit quantity, the minimum achievable bias, which depends on the outcome function class, and which acts as a quantitative measure of the degree of overlap violations. We show that this measure can be more appropriate than an overlap assumption in finite samples for quantifying the underlying difficulty of the reweighting problem.

Finally, given the central role of restrictions on the outcome model in both of the previous results, we briefly consider the setting in which this assumption is incorrect. In particular, we provide simple moment conditions under which we can retain a finite bound on the error when the true outcome model is not in the assumed class.

1.2 Related work

Estimators that target balance. Many reweighting estimators in causal inference explicitly target the discrepancy between source and target distribution, also known as *balance* (Hainmueller, 2012; Zubizarreta, 2015; Athey et al., 2018; Hirshberg et al., 2019; Zhao, 2019; Tan, 2020; Hazlett, 2020; Arbour et al., 2021). See Ben-Michael et al. (2021) for a summary. The literature on domain adaptation also uses worst-case discrepancy between distributions (Mansour et al., 2009; Gretton et al., 2009; Yu and Szepesvári, 2012; Courty et al., 2014). Some approaches learn representations that minimize these discrepancies (Ganin et al., 2016; Shen et al., 2018; Assaad et al., 2021). Closely related are estimators that target the density ratio between two groups through a surrogate loss (Sugiyama et al., 2007b; Nguyen et al., 2010; Sugiyama et al., 2012).

Overlap in causal inference and adversarial training. Many existing theoretical treatments of balancing weights require the density ratio to exist (called *overlap* in causal inference), which is typically used for proving asymptotic consistency (see Hirshberg et al., 2019; Kallus, 2020b). This assumption, however, can be highly restrictive especially in high-dimensions, as illustrated by D’Amour et al. (2021). See also Khan and Tamer (2010) for a discussion of the implications of overlap violations for causal inference. The same topic arises in adversarial training. See, for example, Dupuis and Mao (2019); Birrell et al. (2020a), who generalize ϕ -divergences to distributions that do not have common support. This idea is applied to GANs in Song and Ermon (2020); Glaser et al. (2021).

Domain adaptation and causal inference. We emphasize that domain adaptation and causal inference are both special cases of the same problem setup. Related work combines ideas from these two litera-

tures. For example, Shalit et al. (2017); Johansson et al. (2020) use integral probability metrics to estimate causal effects without the need for an overlap assumption. The same idea is used in Kallus (2020b) for matching estimators in causal inference. Other work has made the connection between causal inference and adversarial training (Yoon et al., 2018; Ozery-Flato et al., 2018).

2 Problem Setup

Let $X \in \mathcal{X}$ denote covariates, and $Y \in \mathbb{R}$ denote outcomes. We study the general class of problems with source and target populations, P and Q , with different joint distributions over X and Y . We observe X in both populations, but only observe the outcomes, Y , for the source population, P . The goal is to estimate the missing mean in the target population, $\mathbb{E}_Q[Y]$. Many important problems share this structure, including causal inference and domain adaptation.

Problem Setting 1 (Causal Inference). Consider the causal inference setting with a binary treatment status variable, T , and potential outcomes $Y(0)$ and $Y(1)$. For the control group, we observe covariates X and the potential outcome $Y(0)$. In the treated group, we still observe X but do not observe $Y(0)$. Therefore, finding the average treatment effect on the treated is equivalent to solving the problem setup described above, where P corresponds to the $T = 0$ population, Q corresponds to the $T = 1$ population, and Y corresponds to $Y(0)$.

Problem Setting 2 (Domain Adaptation). Consider a classification task with features X , labels Z , and loss function ℓ . In a source environment where we observe both X and Z , we train a model h for predicting Z given X . We would like to estimate the average risk of our classifier in a new environment where we observe X but not Z . This problem is equivalent to the setup above, where P corresponds to the source environment, Q corresponds to the target environment, and Y corresponds to the loss, $\ell(h(X), Z)$.

2.1 Ignorability and Overlap

To estimate the mean of Y in Q using the outcomes from P , we require some kind of regularity between the source and target populations. A common assumption is the *ignorability* assumption (also called the *covariate shift* assumption, or *selection on observables*), which requires the relationship between covariates and outcomes to be the same across the two groups:

Assumption 1 (Ignorability). For all $x \in \mathcal{X}$,

$$P(Y|X = x) = Q(Y|X = x).$$

Notice that in the causal inference setting described

above, Assumption 1 is equivalent to the standard conditional independence assumption, $Y(0) \perp\!\!\!\perp T|X$.

Typically, Assumption 1 is paired with a requirement that the density ratio dQ/dP exists, also known as *overlap* or *continuity* in different literatures:

Definition (Overlap). *We say that overlap holds if Q is absolutely continuous with respect to P .*

2.2 Importance Weighting

In the special case where Assumption 1 and overlap hold, we can estimate the mean of the missing outcomes by reweighting the observed outcomes with the density ratio. This estimator is called *importance weighting* or *inverse probability weighting* (IPW) and is unbiased:

$$\begin{aligned} \mathbb{E}_P \left[\frac{dQ}{dP}(X) Y \right] &= \mathbb{E}_P \left[\frac{dQ}{dP}(X) \mathbb{E}_P[Y|X] \right] \\ &= \mathbb{E}_Q[\mathbb{E}_P[Y|X]] = \mathbb{E}_Q[Y], \end{aligned}$$

where we use ignorability for the last equality.

Importance weighting has two main drawbacks. First, the overlap assumption is very strong, especially in high dimensions (D’Amour et al., 2021). But even if overlap holds in the super-population, in finite samples there are usually so-called practical overlap violations: regions of the covariate space that are well-represented in the target population, but very rare in the source population, leading to large importance weights.

2.3 Mean Squared Error

Large weights lead to large mean squared error. Consider arbitrary weights $w(X)$. We will expand the mean squared error (MSE) of $\mathbb{E}_P[w(X)Y]$ for estimating $\mathbb{E}_Q[Y]$ using the standard bias-variance decomposition. Define the *outcome function*, $f_0(x) := \mathbb{E}_P[Y|X = x] = \mathbb{E}_Q[Y|X = x]$ and likewise, let $\sigma_0^2(x)$ be the conditional variance of Y . Then,

$$\begin{aligned} \text{MSE}(w) &= \mathbb{E}_P[(w(X)Y - \mathbb{E}_Q[Y])^2] \\ &= (\mathbb{E}_P[w(X)Y] - \mathbb{E}_Q[Y])^2 + \text{Var}_P[w(X)Y] \\ &= (\mathbb{E}_P[w(X)f_0(X)] - \mathbb{E}_Q[f_0(X)])^2 \quad (1) \\ &\quad + \mathbb{E}_P[w(X)^2\sigma_0^2(X)]. \quad (2) \end{aligned}$$

The MSE depends on two quantities: (1) the imbalance of the mean of the outcome function f_0 between the reweighted source distribution and the target distribution; and (2) the variability of the weights under the source distribution, which amplifies the noise in the outcomes. With practical overlap violations in high dimensional problems, $w = dQ/dP$ can be enormous, and (2) will result in a large MSE (see, for example, Kang et al., 2007). Balancing weights, introduced in Section 3,

explicitly target the trade-off between bias and variance, as discussed extensively in Kallus (2020b).

2.4 Notation

We now introduce formal notation used for the remainder of the paper. Let $(\mathcal{X}, \mathcal{S})$ be a measurable space.¹ Let P and Q be given probability measures on $(\mathcal{X}, \mathcal{S})$. Let f_0 be a real-valued measurable function on \mathcal{X} . Denote $\mathcal{M}(\mathcal{X})$ the space of signed finite measures on $(\mathcal{X}, \mathcal{S})$ and $\mathcal{M}(P)$ those absolutely continuous with respect to P . Denote $\mathcal{P}(\mathcal{X})$ the space of probability measures on $(\mathcal{X}, \mathcal{S})$ and $\mathcal{P}(P)$ those absolutely continuous with respect to P .

With a slight abuse of notation, for measurable $f : \mathcal{X} \rightarrow \mathbb{R}$ and both $M \in \mathcal{M}(\mathcal{X})$ and $M \in \mathcal{P}(\mathcal{X})$, we will write $\mathbb{E}_M[f] := \int_{\mathcal{X}} f(x)dM(x)$. We assume $\mathbb{E}_P[|f_0|] < \infty$ and $\mathbb{E}_Q[|f_0|] < \infty$.

While our setting is quite general, it may be helpful for the reader to keep in mind the case where \mathcal{X} is finite and discrete with cardinality n . In this case, P and Q are probability vectors of length n and measurable functions are simply vectors in \mathbb{R}^n . Likewise, $\mathcal{M}(\mathcal{X})$ is just \mathbb{R}^n .

3 Balancing Weights

In this section, we briefly review the existing work on balancing weights estimators and then introduce our main contributions. Balancing weights estimators find weights $w(X)$ with minimum dispersion, subject to a balance constraint between the target covariate distribution and the reweighted source distribution. In general, we will consider weights such that $\mathbb{E}_P[w(X)] = 1$, i.e., we always end up with the same “size” population that we started with. The problem of choosing weights can be reformulated as finding a measure $R \in \mathcal{M}(P)$ such that $\int_{\mathcal{X}} dR(x) = 1$, with $w := dR/dP$. This corresponds to the intuition behind reweighting as creating a “pseudo-population” based on P intended to match Q . We will therefore often use w and R interchangeably.

A simple balancing weights estimator might constrain the mean of the covariates to match within tolerance δ , similarly to Zubizarreta (2015). For example, let $\mathcal{X} = \mathbb{R}^d$. We could find the minimum variance w such that

$$\|\mathbb{E}_R[X] - \mathbb{E}_Q[X]\|_2 \leq \delta, \quad (3)$$

where, as a reminder, $w = dR/dP$. If the outcome function $f_0(X)$ is linear with bounded coefficients, then this constraint will bound the bias term (2) and the

¹To side-step topological issues, we assume that \mathcal{X} is a separable Banach space.

tuning parameter δ lets us trade-off bias and variance to achieve a smaller MSE than the importance weights.

3.1 Assumptions on the Outcome Function

More generally, we may not want to assume that f_0 is linear. But this presents a difficulty: without any further restrictions, for any $w \neq dQ/dP$, there always exists an adversarial f_0 that can make the bias term (1) arbitrarily large. Only the density ratio guarantees bounded bias for *any* f_0 , but often at the cost of high variance.

Therefore, in practical settings, we instead restrict the outcome function in order to control the error. To make progress, we assume that f_0 belongs to some function class \mathcal{F} :

Assumption 2. *The outcome function f_0 belongs to \mathcal{F} where \mathcal{F} is a closed and convex set of measurable real-valued functions such that for all $f \in \mathcal{F}$, $\mathbb{E}_P[|f|] < \infty$, $\mathbb{E}_Q[|f|] < \infty$, and $-f \in \mathcal{F}$.*

For the causal inference problem setting, Assumption 2 requires making an assumption about the relation of the potential outcome $Y(0)$ to the covariates. For the domain adaptation problem setting, the assumption is about the relationship between the accuracy of a predictor, $\ell(h(X), Z)$, to its input features, X .

Many choices of \mathcal{F} in Assumption 2 are quite general and justifiable with domain knowledge. Some examples for $0 < B < \infty$ are:

- Bounded functions: $\mathcal{F}_\infty := \{f : \|f\|_\infty \leq B\}$
- Lipschitz functions: $\mathcal{F}_{\text{Lip}(c)} := \{f : \|f\|_{\text{Lip}(c)} \leq B\}$
- RKHS functions: $\mathcal{F}_{\mathcal{H}} := \{f : \|f\|_{\mathcal{H}} \leq B\}$,

where $\|\cdot\|_{\text{Lip}(c)}$ denotes the Lipschitz constant with respect to a metric c and $\|\cdot\|_{\mathcal{H}}$ denotes the norm in some Reproducing Kernel Hilbert Space (RKHS), \mathcal{H} .

Under Assumption 2, the bias is bounded by the worst-case discrepancy in means over \mathcal{F} . This quantity is called an *integral probability metric* (IPM), defined for any set of functions, \mathcal{G} , and any $M, N \in \mathcal{M}(\mathcal{X})$ as:²

$$\text{IPM}_{\mathcal{G}}(M, N) := \sup_{g \in \mathcal{G}} \{|\mathbb{E}_M[g] - \mathbb{E}_N[g]|\}.$$

The bias term (1) for a re-weighted population R under Assumption 2 is upper-bounded by:

$$|\mathbb{E}_Q[f_0] - \mathbb{E}_R[f_0]| \leq \text{IPM}_{\mathcal{F}}(Q, R). \quad (4)$$

This value is always finite by our assumptions on \mathcal{F} and we can trade it off against the variance of the weights.

²If $g \in \mathcal{G} \implies -g \in \mathcal{G}$ then the absolute value can be omitted.

Before introducing the general form of balancing weights in Section 3.2, we define two quantities that will be useful in our discussion below, the maximum and minimum bias.

Definition (Maximum and minimum bias). *The maximum bias, δ_{\max} , is the bias under uniform weights (when $R = P$). The minimum bias, δ_{\min} , is the smallest bias achievable by reweighting P .*

$$\delta_{\max} := \text{IPM}_{\mathcal{F}}(Q, P) \quad (5)$$

$$\delta_{\min} := \inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \{\text{IPM}_{\mathcal{F}}(Q, R)\}. \quad (6)$$

Since $R = P$ is feasible for (6), $\delta_{\min} \leq \delta_{\max}$. In the special case where overlap holds, $R = Q$ is also feasible, which implies $\delta_{\min} = 0$.

3.2 Minimax Balancing Weights

Assumption 2 and the resulting IPM bound on the bias (4) lead to a generalized balancing weights estimator as discussed in Kallus (2020b) and Ben-Michael et al. (2021). Define $\sigma^2 := \sup_{x \in \mathcal{X}} \sigma_0^2(x)$, where we assume $0 < \sigma^2 < \infty$. We can plug these bounds into the MSE to arrive at the following optimization problem:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \left\{ \text{IPM}_{\mathcal{F}}(Q, R)^2 + \sigma^2 \mathbb{E}_P \left[\left(\frac{dR}{dP} \right)^2 \right] \right\} \quad (7)$$

A solution always exists because the objective is finite for $R = P$, which is feasible. For $\sigma^2 > 0$, the problem is strongly convex in R and has a unique solution. Since the IPM term is itself a supremum, this estimator is sometimes referred to as minimax balancing weights.

Furthermore, $\exists \delta > 0$ such that (7) has the same minimizer as:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \mathbb{E}_P \left[\left(\frac{dR}{dP} \right)^2 \right] \quad (8)$$

such that $\text{IPM}_{\mathcal{F}}(Q, R) \leq \delta$.

We view σ^2 and δ as exchangeable tuning parameters: σ^2 represents the importance of reducing the variance of the weights; δ represents the level of acceptable bias. For $\sigma^2 \in (0, \infty)$, the corresponding δ lies in $(\delta_{\min}, \delta_{\max})$.

3.3 Our Contributions

In this paper, we start from the premise that Assumption 2 is necessary to achieve a reasonable MSE in high dimensions leading to estimators (7) and (8). Our main argument is that Assumption 2 immediately implies two additional results.

First, we derive a general duality result that lets us rewrite problems (7) and (8) as a single convex optimization problem over \mathcal{F} . Therefore, we can solve the minimax balancing weights problem by optimizing a simple convex loss over a function class. Furthermore, this reformulation shows that the optimal weights are always a rescaled and recentered member of \mathcal{F} .

Second, we no longer need an overlap assumption. Before restricting f_0 , we saw that only $w = dQ/dP$ could guarantee finite bias. Therefore, to bound the MSE we needed the density ratio to exist. But once we assume that $f_0 \in \mathcal{F}$, we no longer need the density ratio to exist, and we can simply minimize our bound on the MSE directly. Moreover, we argue that Assumption 2 provides us with a more appropriate *quantitative* measure of overlap — the minimum bias, δ_{\min} — that precisely characterizes the difficulty of translating results from one distribution to another.

4 Duality Theory for Balancing Weights

In this section, we derive a dual characterization of the solution, R^* , to problems (7) and (8) and the corresponding minimax weights $w^* = dR^*/dP$.

4.1 The Variance of the Weights

Our dual derivation uses the fact that the variance term can be written as a special case of a class of information-theoretic divergences called ϕ -divergences (Sriperumbudur et al., 2009). These have a variational representation that will allow us to simplify the minimax problems (7) and (8) into a single convex loss.

Definition (ϕ -Divergence). *For any convex function ϕ with $\phi(1) = 0$, the ϕ -divergence between $M \in \mathcal{M}(\mathcal{X})$ and $N \in \mathcal{P}(\mathcal{X})$ is:*

$$D_\phi(M||N) := \mathbb{E}_N[\phi(dM/dN)],$$

where $D_\phi(M||N) = \infty$ if M is not absolutely continuous with respect to N .

Notice that we can subtract off a constant to re-center our variance term in (8) without affecting the minimizer over R . We can then rewrite the objective as the divergence between R and P with $\phi(x) = x^2 - 1$. This is known as the χ^2 divergence, and we denote it $D_2(R||P)$:

$$\mathbb{E}_P \left[\left(\frac{dR}{dP} \right)^2 - 1 \right] = D_2(R||P).$$

Variational representations. It is possible to express ϕ -divergences in a dual form, called a *variational*

representation, as a supremum over measurable functions. Let $M \in \mathcal{M}(\mathcal{X})$ and let $N \in \mathcal{P}(\mathcal{X})$. Let ϕ^* denote the convex conjugate of ϕ . Then Keziou (2003) and Nguyen et al. (2005) show that:

$$D_\phi(M||N) = \sup_f \{ \mathbb{E}_M[f] - \mathbb{E}_N[\phi^*(f)] \}, \quad (9)$$

where the supremum is over all real-valued measurable functions on \mathcal{X} . If we additionally assume, as we do for R , that $\mathbb{E}_M[1] = 1$, then we have the tighter representation,

$$D_\phi(M||N) = \sup_f \{ \mathbb{E}_M[f] - \Lambda_N^\phi[f] \} \quad (10)$$

$$\text{where } \Lambda_N^\phi[f] := \inf_{\lambda \in \mathbb{R}} \{ \lambda + \mathbb{E}_N[\phi^*(f - \lambda)] \}.$$

This result, using the infimum over λ in the spirit of Ruderman et al. (2012), appears to have been independently proposed by Agrawal and Horel (2020) and Birrell et al. (2020b). Under minimal conditions on ϕ , the suprema in (9) and (10) are achieved by $\phi'(dM/dN)$.

4.2 Dual Formulation

We now present our main duality result under Assumption 2 where $f_0 \in \mathcal{F}$.

Theorem 4.1. *Under Assumptions 1 and 2, for $\delta > \delta_{\min}$, the optimization problem (8) has a unique solution,*

$$\frac{dR^*}{dP} = 1 + \left(\frac{\mathbb{E}_Q[f^*] - \mathbb{E}_P[f^*] - \delta}{\text{Var}_P[f^*]} \right) (f^* - \mathbb{E}_P[f^*]),$$

where, for a unique $\mu \geq 0$ corresponding to δ , f^* achieves the following supremum:

$$\sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_Q[f] - \mathbb{E}_P[f] - \frac{\mu}{4} \text{Var}_P[f] \right\}. \quad (11)$$

The resulting MSE is:

$$\text{MSE}(R^*) \leq \delta^2 + \sigma^2 \frac{(\mathbb{E}_Q[f^*] - \mathbb{E}_P[f^*] - \delta)^2}{\text{Var}_P[f^*]}. \quad (12)$$

Proof Sketch. The full proof is available in the Appendix. Here we provide a brief outline. In the first step, we show that problem (8) is equivalent to:

$$\sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_Q[f] + \inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \{ (1/\mu) D_2(R||P) - \mathbb{E}_R[f] \} \right\}$$

for some $\mu > 0$ corresponding to δ . In the second step, we apply (10) for the χ^2 divergence to show that the inner subproblem has an explicit solution:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \{ (1/\mu) D_2(R||P) - \mathbb{E}_R[f] \} = -\mathbb{E}_P[f] - \frac{\mu}{4} \text{Var}_P[f].$$

The theorem then follows from standard convex duality results.

Remark 4.1 (The Shape of the Weights). The weights dR^*/dP are equal to f^* multiplied by some scalar s_1 and then shifted by some scalar s_2 :

$$\frac{dR^*}{dP} = s_1 + s_2 f^*,$$

where s_1 and s_2 depend on both δ and f^* . Therefore, if we assume \mathcal{F} is the set of quadratic functions, then the balancing weights will also be quadratic, and if we assume \mathcal{F} is an RKHS with a certain kernel, then the balancing weights will belong to an RKHS with that same kernel.

Remark 4.2 (Other ϕ -Divergences). We can replace the χ^2 divergence in the balancing weight problems (7) and (8) with other ϕ -divergences. A duality result corresponding to Theorem 4.1 will hold for any convex function ϕ such that $\phi(1) = 0$ with convex conjugate ϕ^* such that $\{\phi^* < \infty\} = \mathbb{R}$. See the Appendix for details. We can use this general formulation to derive corresponding duality results for entropy balancing (Hainmueller, 2012) or other measures of dispersion (see Ben-Michael et al., 2021).

Remark 4.3 (Tuning Parameters). For every δ there is a unique corresponding μ . Therefore, we can treat μ as a tuning parameter instead of δ and solve (11) directly. In terms of μ , the solution to (8) is:

$$\frac{dR^*}{dP} = 1 + \frac{\mu}{2} (f^* - \mathbb{E}_P[f^*]) \quad (13)$$

and there is a closed form relationship between μ and δ given by:

$$\delta = \mathbb{E}_Q[f^*] - \mathbb{E}_P[f^*] - \frac{\mu}{2} \text{Var}_P[f^*].$$

Going forward, we will often use δ and μ interchangeably.

4.3 The Full Information Case

To help illustrate Theorem 4.1, consider the simplified setting where we know f_0 exactly. This corresponds to a special case of Assumption 2 where \mathcal{F} is the convex hull of $\{f_0, -f_0\}$. Assume without loss of generality that $\mathbb{E}_Q[f_0] \geq \mathbb{E}_P[f_0]$. Then, applying Theorem 4.1, we get $f^* = f_0$, and

$$\frac{dR^*}{dP} = 1 + \left(\frac{\mathbb{E}_Q[f_0] - \mathbb{E}_P[f_0] - \delta}{\text{Var}_P[f_0]} \right) (f_0 - \mathbb{E}_P[f_0]).$$

The optimal weights are always a rescaled and recentered version of f_0 . In this special case, the dual optimal f^* does not depend on δ ; only the scaling factor does. Therefore, the MSE bound (12) becomes a quadratic in δ and we can solve for the optimal bias:

$$\delta^* = \left(\frac{\sigma^2}{\text{Var}_P[f_0] + \sigma^2} \right) |\mathbb{E}_Q[f_0] - \mathbb{E}_P[f_0]|,$$

which gives

$$\text{MSE}(w^*) \leq \left(\frac{\sigma^2}{\text{Var}_P[f_0] + \sigma^2} \right) (\mathbb{E}_Q[f_0] - \mathbb{E}_P[f_0])^2.$$

This is an independently interesting result. With complete information, we can analytically find the optimal bias-variance trade-off. Under homoskedasticity, these weights have the smallest possible MSE over all w such that $\mathbb{E}_P[w] = 1$.

4.4 The Linear Case

For a second simple example, we return to the linear problem in (3). In this case, duality shows that balancing weights are equivalent to fitting a linear model. In fact, for a certain choice of linear \mathcal{F} , problem (11) is identical to linear regression.

Let $g : \mathcal{X} \rightarrow \mathbb{R}^d$ be some feature map. Assume that our balance constraint is:

$$\|\mathbb{E}_R[g(X)] - \mathbb{E}_Q[g(X)]\|_2 \leq \delta.$$

This is equivalent to problem (8) using the following linear function class:

$$f_0 \in \mathcal{F}_{\text{lin}} = \{\beta^T g(X) : \|\beta\|_2 \leq 1\}.$$

Applying Theorem 4.1, we know $f^* = (\beta^*)^T g(X) \in \mathcal{F}_{\text{lin}}$ and therefore the optimal weights will be linear. Solving (11) via calculus, we get:

$$\beta^* = c_1 (\text{Cov}_P[g(X)] + c_2 I)^{-1} (\mathbb{E}_Q[g(X)] - \mathbb{E}_P[g(X)])$$

for some scalars c_1 and c_2 that depend on μ . Notice the dependence on the inverse of the covariance of the features plus a regularization term. This is another way of deriving a well-known result: for \mathcal{F}_{lin} , problem (8) is identical to estimating $\mathbb{E}_Q[Y]$ by fitting a ridge regression in the P population and then applying it to the Q population. See Kallus (2020b) for a direct proof. If we replace the ℓ_2 -norm with the ℓ_1 -norm, we obtain a similar equivalence for Lasso. When the regularization term is 0, we obtain linear regression as a special case.

Several other papers have recognized the duality between linear regression and balancing weights estimators for a linear function class; see Zhao and Percival (2017); Zhao (2019); Wang and Zubizarreta (2020); Tan (2020); Ben-Michael et al. (2021). Theorem 4.1 generalizes these existing duality results for linear function classes to general function classes \mathcal{F} .

5 Outcome Assumptions and Overlap

In this section, we discuss the implications of the outcome assumption for overlap. First, we show that if

overlap holds, then under conditions on \mathcal{F} , as $\delta \rightarrow 0$, the balancing weights converge to the importance weights dQ/dP .

However, when overlap is violated, the only impact on the balancing weights estimator is that the minimum bias (which depends on our function class \mathcal{F}) is greater than zero. Due to Assumption 2, for any $\delta > \delta_{\min}$ there still exists a solution to (8) that bounds the MSE. If the variance of the outcomes is large, then we naturally want to choose δ larger than δ_{\min} and the failure of overlap has no impact on our estimator.

Instead, we argue that we should use the minimum bias, δ_{\min} , directly as a measure of practical overlap violations. We illustrate that in finite samples, δ_{\min} can be large even when overlap holds in the super population, and likewise that δ_{\min} can be small even when overlap is violated in the super population. Therefore, under Assumption 2, δ_{\min} is a more precise summary of the underlying difficulty of the reweighting problem.

5.1 Convergence to Importance Weights

We begin with an example. Let $\mathcal{X} = \mathbb{R}$. Let P be Gaussian with mean 1 and variance 1, let Q be Gaussian with mean 2 and variance 1, and let p and q denote their densities. Let $\mathcal{F} = \{f : \|f\|_{\infty} \leq 1\}$ so that the outcome function is bounded between -1 and 1 . The solution to the dual problem, f^* , and the corresponding optimal weights are illustrated in Figure 1.

The weights have a distinctive form. When $\delta = \delta_{\max}$, the optimal weights are uniform. As the allowed bias δ decreases, the optimal weights trace out the density ratio dQ/dP but truncated above and below. This is the form of a well-known estimator in the causal inference literature, IPW with a trimmed propensity score (Yang and Ding, 2018): under Assumption 2 with bounded functions, the balancing weights formulation provides formal justification for using the truncated density ratio for weights. As $\delta \rightarrow 0$, the optimal weights converge to dQ/dP .

In general, convergence to the importance weights will *always* occur as $\delta \rightarrow 0$ under certain conditions on \mathcal{F} .

Definition (Distribution-defining). \mathcal{F} is distribution-defining if $\forall M, N \in \mathcal{P}(\mathcal{X})$, $IPM_{\mathcal{F}}(M, N) = 0$ if and only if $M = N$.

For example, \mathcal{F}_{∞} and $\mathcal{F}_{\text{Lip}(c)}$ are distribution-defining, as is $\mathcal{F}_{\mathcal{H}}$ for a universal kernel. When \mathcal{F} is distribution-defining then only dQ/dP can achieve worst-case bias zero. Therefore, when overlap holds and \mathcal{F} is distribution-defining, the optimal weights, $w^* \rightarrow dQ/dP$ as $\delta \rightarrow 0$.

This connection between balancing weights and the




Figure 1: The optimal weights and corresponding dual optimal function for the Gaussian example, with δ starting at δ_{\max} and shrinking towards zero.

density ratio is not new: among others, Zhao (2019) makes a similar point. We discuss the connections to our duality result in the Appendix.

5.2 Balancing Weights Without Overlap

What if overlap does not hold? Then if \mathcal{F} is distribution-defining, by definition, $\delta_{\min} > 0$. In this case, problem (8) still has a solution that bounds the MSE for any $\delta \geq \delta_{\min}$, but the failure of overlap precludes us from using $\delta = 0$. However, the motivation behind balancing weights is to avoid using an unbiased estimator: if the variance of the outcomes is sufficiently high, we might still prefer to use $\delta > \delta_{\min}$.

Consider a simple example in which we reweight $P = \text{Uniform}(1, 2)$ to target $Q = \text{Uniform}(1.01, 2.01)$. While Q is not absolutely continuous with respect to P , intuitively, we should be able to find w that achieves small error because the distributions are close to each other. The function class \mathcal{F} provides a formal definition of “close to each other” for the purposes of reweighting.

For these uniform P and Q , there is an irreducible bias,

δ_{\min} , for any possible weights:

$$\delta_{\min} = \sup_{f \in \mathcal{F}} \int_1^{1.01} f(x) dx + \sup_{f \in \mathcal{F}} \int_2^{2.01} f(x) dx$$

If \mathcal{F} is unrestricted, then f could take on arbitrarily large values on the intervals $[1, 1.01]$ and $[2, 2.01]$. Therefore, the bias is unbounded without Assumption 2, which typically justifies imposing an overlap assumption. However, if we assume $\mathcal{F} \in \{f : \|f\|_{\infty} \leq B\}$, for example, then we have $\delta_{\min} = 0.02B$ which may be quite small.

The parameter $\delta > \delta_{\min}$ in problem (8) is a tuning parameter that trades off bias and variance. If the variance of the outcomes is very large, then we may prefer to use a value of δ larger than δ_{\min} . In this case, the overlap violation would not have any impact on our estimator at all. On the other hand, if the variance of the outcomes is small relative to δ_{\min} , we may prefer to use a value of δ close to 0. The best we could do would be to set $\delta = \delta_{\min}$; without overlap, the best achievable lower bound for the MSE is δ_{\min}^2 .

5.3 Quantitative Overlap

In finite samples, we argue that δ_{\min} will often be a more useful measure of overlap than the existence of the density ratio in a super-population. For example, let $P_{\text{super}} = \mathcal{N}(100, 1)$ and $Q_{\text{super}} = \mathcal{N}(-100, 1)$. Technically, overlap holds and the density ratio exists over all of \mathbb{R} . For concreteness, let $\mathcal{F} = \mathcal{F}_{\mathcal{H}}$ be an RKHS with a Gaussian kernel. Then for P_{super} and Q_{super} , $\delta_{\min} = 0$, because $w = dQ/dP$ will perfectly balance the RKHS. However, any finite dataset will have severe practical overlap violations. Let P be a sample of n data points from P_{super} and likewise for Q . With high probability, the points in P and the points in Q will be far apart, and as a result, δ_{\min} over the RKHS will be large.

On the other hand, if we let $P_{\text{super}} = \text{Uniform}(1, 2)$ and $Q_{\text{super}} = \text{Uniform}(1.01, 2.01)$, overlap does not hold and δ_{\min} will be non-zero for the super-population. But, for corresponding finite samples P and Q , δ_{\min} is still likely to be very small. In these examples, super-population overlap is misleading, whereas δ_{\min} is a precise quantitative summary of the difficulty of the reweighting problem for function class \mathcal{F} .

6 IHDP Example

In this section, we walk through an example on a real dataset to make the previous two sections more concrete. We apply balancing weights to the Infant Health and Development Program (IHDP) using an RKHS function class.

6.1 The IHDP Dataset and Setup

The Infant Health and Development Program (IHDP) data set is a standard observational causal inference benchmark from Hill (2011), based on data from a randomized control trial of an intensive home visiting and childcare intervention for low birth weight infants born in 1985. We consider a non-experimental subset of the original data with $n_0 = 608$ children assigned to control, $n_1 = 139$ children assigned to treatment, and $n = 747$ total children. For all children, we have a range of baseline covariates, including both categorical covariates, like the mother’s educational attainment, and continuous covariates, like the child’s birth weight. Our goal is to estimate the average outcome (a standardized test score) in the absence of the intensive intervention. We observe this outcome for the 608 control children, and want to re-weight these observations to estimate the missing mean for the 139 treated children.

To do so, we use an RKHS as a flexible but tractable functional form for f_0 . In particular, we assume that $\mathcal{F} = \mathcal{F}_{\mathcal{H}}^B := \{f : \|f\|_{\mathcal{H}} \leq B\}$ for $B < \infty$, where \mathcal{H} is the RKHS induced by the Gaussian kernel,

$$\mathcal{K}(x_1, x_2) = \exp\left(-\frac{1}{2}\|x_1 - x_2\|_2^2\right).$$

Define $K \in \mathbb{R}^{n \times n}$ with $K_{ij} = \mathcal{K}(X_i, X_j)$. Then for any $f \in \mathcal{F}$, there exists an $\alpha \in \mathbb{R}^n$ such that $\alpha^T K \alpha \leq B$ and $f(X_j) = \sum_{i=1}^n \alpha_i K_{ij}, \forall j$.

6.2 Solving the Dual Problem

We compute the minimax balancing weights by solving the dual problem (11) directly for many values of the tuning parameter $\mu > 0$. The dual problem can be written as a quadratic optimization problem over the vectors α that characterize the $f \in \mathcal{F}$. See the Appendix for details. We obtain the corresponding optimal weights by plugging the resulting f^* into (13).

The balancing weights interpolate between two extremes. See Figure 2 for an illustration. At one extreme are the weights with maximum bias and minimum variance. This is achieved at $\mu = 0$, which results in uniform weights and corresponding bias $\delta = \delta_{\max}$.

At the other extreme are the weights with maximum variance and minimum bias. Since some of the covariates are continuous, the data points for the control and treated groups have disjoint support. Therefore, there are no weights that achieve zero bias. Instead, we find weights that achieve the smallest possible bias over $\mathcal{F}_{\mathcal{H}}^B$, δ_{\min} , which will correspond to some $\mu = \mu_{\max} < \infty$. We find μ_{\max} by increasing μ until the bias stops decreasing. The corresponding weights are shown in black in Figure 2.

The function class $\mathcal{F}_{\mathcal{H}}^B$ has the nice property that the worst-case bias scales with the norm bound B , $\text{IPM}_{\mathcal{F}_{\mathcal{H}}^B}(R, Q) = B \cdot \text{IPM}_{\mathcal{F}_{\mathcal{H}}^1}(R, Q)$. Furthermore, regardless of B , the optimal weights remain identical. Therefore, we can report the bias as a fraction of the size of functions in \mathcal{F} . For the IHDP data, $\delta_{\max} = 0.102B$ and $\delta_{\min} = 0.089B$ with corresponding variances σ^2 (with uniform weights) and $1.7\sigma^2$. For this particular problem, we achieve most of the bias reduction with smaller weights: the intermediate weights in Figure 2 have $\delta = 0.090B$ with variance $1.35\sigma^2$, highlighting the relevance of the bias-variance trade-off.

In any real data set with continuous covariates, two finite samples will typically have disjoint support like we have here. The standard approach in causal inference is to assume that overlap holds in the super-populations from which the samples were drawn. In this case, we could approximate the density ratio asymptotically. However, as we emphasized in Section 5.3, the implications of overlap for balancing weights are entirely summarized by δ_{\min} so we do not need to make such an assumption.

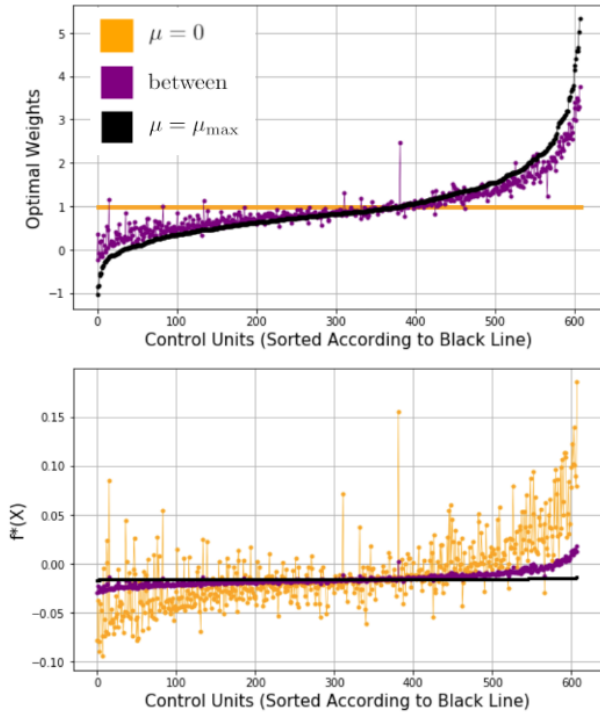


Figure 2: The optimal weights and corresponding dual optimal function for the IHDP example for the extreme values of μ and one intermediate value.

Remark 6.1 (Computational Advantages of the Dual). For an RKHS, there is a closed form of the IPM, which makes the primal and dual problems equally easy to solve. But in some situations, it is computationally easier to solve the dual problem (11) directly instead

of the primal problem (8). Consider a class of neural networks parameterized by bounded network weights θ . Then handling the IPM constraint in the primal problem requires adversarial training, as in Kallus (2020a), which can be quite computationally challenging. On the other hand, (11) requires training a neural network once with a convex loss function which can be accomplished with off-the-shelf SGD.

7 Robustness

Balancing weights rely heavily on the function class in Assumption 2. In this section, we show that with minimal moment conditions we can still retain a bound on the bias even if we have misspecified the function class \mathcal{F} . We consider two function classes. First, a *misspecified* \mathcal{F} for which we solve (8) to find R^* such that $\text{IPM}_{\mathcal{F}}(Q, R^*) \leq \delta$. Second, the *true* function class, \mathcal{G} such that $f_0 \in \mathcal{G}$ and $f_0 \notin \mathcal{F}$. To bound the bias, we need to show that

$$\text{IPM}_{\mathcal{F}}(Q, R^*) \leq \delta \implies \text{IPM}_{\mathcal{G}}(Q, R^*) \leq \rho(\delta) \quad (14)$$

for some $\rho < \infty$ which has good scaling with δ . Without further assumptions, (14) will *not* hold for any \mathcal{G} .

IPMs correspond to common perturbations in the robust statistics literature. For example, $\text{IPM}_{\mathcal{F}_{\infty}}$ and $\text{IPM}_{\mathcal{F}_{\text{Lip}(c)}}$ are equivalent to the total variation (TV) distance and Wasserstein distance respectively. For $\mathcal{F} = \mathcal{F}_{\infty}$, we can apply Lemma E.2 from Zhu et al. (2019) to achieve (14) for any \mathcal{G} . We require an Orlicz norm bound under Q and R^* on $g(X)$ for all $g \in \mathcal{G}$. For a simple example, let \mathcal{G} be linear. Then we get the following result:

Proposition 7.1. *Let $\text{TV}(Q, R^*) \leq \delta$ and let $f_0 \in \{\beta^T x : \|\beta\| \leq 1\}$. If R^* and Q have bounded covariance, then we have the following upper bound on the bias:*

$$|\mathbb{E}_Q[f_0] - \mathbb{E}_{R^*}[f_0]| \leq \rho_1(\delta),$$

where $\rho_1(\delta) = O(\sqrt{\delta})$.

If instead R^* and Q are sub-Gaussian, then we have the following upper bound on the bias:

$$|\mathbb{E}_Q[f_0] - \mathbb{E}_{R^*}[f_0]| \leq \rho_2(\delta),$$

where $\rho_2(\delta) = O(\delta\sqrt{\log(1/\delta)})$.

For general \mathcal{G} , the rate of ρ in terms of δ is similar, but the moment conditions on X become stronger. In practice, these robust statistics results mean that we can make a best guess about \mathcal{F} and as long as Q is sufficiently “nice”, the true bias will not be much larger than δ .

References

- R. Agrawal and T. Horel. Optimal bounds between f -divergences and integral probability metrics. In *International Conference on Machine Learning*, pages 115–124. PMLR, 2020.
- D. Arbour, D. Dimmery, and A. Sondhi. Permutation weighting. In *International Conference on Machine Learning*, pages 331–341. PMLR, 2021.
- S. Assaad, S. Zeng, C. Tao, S. Datta, N. Mehta, R. Henao, F. Li, and L. Carin Duke. Counterfactual representation learning with balancing weights. In A. Banerjee and K. Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 1972–1980. PMLR, 13–15 Apr 2021. URL <https://proceedings.mlr.press/v130/assaad21a.html>.
- S. Athey, G. W. Imbens, and S. Wager. Approximate residual balancing: debiased inference of average treatment effects in high dimensions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(4):597–623, 2018.
- E. Ben-Michael, D. Hirshberg, A. Feller, and J. Zubizarreta. The balancing act in causal inference. 2021.
- J. Birrell, P. Dupuis, M. A. Katsoulakis, Y. Pantazis, and L. Rey-Bellet. (f, γ) -divergences: Interpolating between f -divergences and integral probability metrics. *arXiv preprint arXiv:2011.05953*, 2020a.
- J. Birrell, M. A. Katsoulakis, and Y. Pantazis. Optimizing variational representations of divergences and accelerating their statistical estimation. *arXiv preprint arXiv:2006.08781*, 2020b.
- M. Broniatowski and A. Keziou. Minimization of φ -divergences on sets of signed measures. *Studia Scientiarum Mathematicarum Hungarica*, 43(4):403–442, 2006.
- C. Cortes, Y. Mansour, and M. Mohri. Learning bounds for importance weighting. In *Nips*, volume 10, pages 442–450. Citeseer, 2010.
- N. Courty, R. Flamary, and D. Tuia. Domain adaptation with regularized optimal transport. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 274–289. Springer, 2014.
- P. Dupuis and Y. Mao. Formulation and properties of a divergence used to compare probability measures without absolute continuity. *arXiv preprint arXiv:1911.07422*, 2019.
- A. D’Amour, P. Ding, A. Feller, L. Lei, and J. Sekhon. Overlap in observational studies with high-dimensional covariates. *Journal of Econometrics*, 221(2):644–654, 2021.
- Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- P. Glaser, M. Arbel, and A. Gretton. Kale flow: A relaxed kl gradient flow for probabilities with disjoint support. *arXiv preprint arXiv:2106.08929*, 2021.
- A. Gretton, A. Smola, J. Huang, M. Schmittfull, K. Borgwardt, and B. Schölkopf. Covariate shift by kernel mean matching. *Dataset shift in machine learning*, 3(4):5, 2009.
- C. Ha. A noncompact minimax theorem. *Pacific Journal of Mathematics*, 97(1):115–117, 1981.
- J. Hainmueller. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political analysis*, pages 25–46, 2012.
- C. Hazlett. Kernel balancing. *Statistica Sinica*, 30(3):1155–1189, 2020.
- J. L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- D. A. Hirshberg, A. Maleki, and J. R. Zubizarreta. Minimax linear estimation of the retargeted mean. *arXiv preprint arXiv:1901.10296*, 2019.
- K. Imai and M. Ratkovic. Covariate balancing propensity score. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):243–263, 2014.
- F. D. Johansson, U. Shalit, N. Kallus, and D. Sontag. Generalization bounds and representation learning for estimation of potential outcomes and causal effects. *arXiv preprint arXiv:2001.07426*, 2020.
- N. Kallus. Deepmatch: Balancing deep covariate representations for causal inference using adversarial training. In *International Conference on Machine Learning*, pages 5067–5077. PMLR, 2020a.
- N. Kallus. Generalized optimal matching methods for causal inference. *Journal of Machine Learning Research*, 21(62):1–54, 2020b.
- J. D. Kang, J. L. Schafer, et al. Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical science*, 22(4):523–539, 2007.
- A. Keziou. Dual representation of φ -divergences and applications. *Comptes rendus mathématique*, 336(10):857–862, 2003.

- S. Khan and E. Tamer. Irregular identification, support conditions, and inverse weight estimation. *Econometrica*, 78(6):2021–2042, 2010.
- Y. Mansour, M. Mohri, and A. Rostamizadeh. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*, 2009.
- X. Nguyen, M. J. Wainwright, and M. I. Jordan. On divergences, surrogate loss functions, and decentralized detection. *arXiv preprint math.ST/0510521*, 2005.
- X. Nguyen, M. J. Wainwright, and M. I. Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010.
- M. Ozery-Flato, P. Thodoroff, M. Ninio, M. Rosen-Zvi, and T. El-Hay. Adversarial balancing for causal inference. *arXiv preprint arXiv:1810.07406*, 2018.
- P. R. Rosenbaum and D. B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- A. Ruderman, M. Reid, D. García-García, and J. Petersen. Tighter variational representations of f-divergences via restriction to probability measures. *arXiv preprint arXiv:1206.4664*, 2012.
- U. Shalit, F. D. Johansson, and D. Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*, pages 3076–3085. PMLR, 2017.
- J. Shen, Y. Qu, W. Zhang, and Y. Yu. Wasserstein distance guided representation learning for domain adaptation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- M. Sion. On general minimax theorems. *Pacific Journal of mathematics*, 8(1):171–176, 1958.
- J. Song and S. Ermon. Bridging the gap between f-gans and wasserstein gans. In *International Conference on Machine Learning*, pages 9078–9087. PMLR, 2020.
- B. K. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf, and G. R. Lanckriet. On integral probability metrics, ϕ -divergences and binary classification. *arXiv preprint arXiv:0901.2698*, 2009.
- M. Sugiyama, M. Krauledat, and K.-R. Müller. Covariate shift adaptation by importance weighted cross validation. *Journal of Machine Learning Research*, 8(5), 2007a.
- M. Sugiyama, S. Nakajima, H. Kashima, P. Von Buehnan, and M. Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In *NIPS*, volume 7, pages 1433–1440. Citeseer, 2007b.
- M. Sugiyama, T. Suzuki, and T. Kanamori. Density-ratio matching under the bregman divergence: a unified framework of density-ratio estimation. *Annals of the Institute of Statistical Mathematics*, 64(5):1009–1044, 2012.
- Z. Tan. Regularized calibrated estimation of propensity scores with model misspecification and high-dimensional data. *Biometrika*, 107(1):137–158, 2020.
- Y. Wang and J. R. Zubizarreta. Minimal dispersion approximately balancing weights: asymptotic properties and practical considerations. *Biometrika*, 107(1):93–105, 2020.
- S. Yang and P. Ding. Asymptotic inference of causal effects with observational studies trimmed by the estimated propensity scores. *Biometrika*, 105(2):487–493, 2018.
- J. Yoon, J. Jordon, and M. Van Der Schaar. Ganite: Estimation of individualized treatment effects using generative adversarial nets. In *International Conference on Learning Representations*, 2018.
- Y. Yu and C. Szepesvári. Analysis of kernel mean matching under covariate shift. *arXiv preprint arXiv:1206.4650*, 2012.
- Q. Zhao. Covariate balancing propensity score by tailored loss functions. *The Annals of Statistics*, 47(2):965–993, 2019.
- Q. Zhao and D. Percival. Entropy balancing is doubly robust. *Journal of Causal Inference*, 5(1), 2017.
- B. Zhu, J. Jiao, and J. Steinhardt. Generalized resilience and robust statistics. *arXiv preprint arXiv:1909.08755*, 2019.
- J. R. Zubizarreta. Stable weights that balance covariates for estimation with incomplete outcome data. *Journal of the American Statistical Association*, 110(511):910–922, 2015.

Supplementary Material: Outcome Assumptions and Duality Theory for Balancing Weights

CONTENTS

- A Proof of Theorem 4.1
- B General statement and proof for ϕ -divergences for Remark 4.2
- C Extension with non-negative weights
- D Connection to surrogate loss for density ratio estimation
- E Details of RKHS optimization problem in Section 6

A Proof of Theorem 4.1

We derive a dual formulation of the optimization problem:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} D_2(R||P)$$

such that $\text{IPM}_{\mathcal{F}}(Q, R) \leq \delta$,

where $\delta > \delta_{\min}$. As a reminder, $D_2(R||P) := \mathbb{E}_P[(dR/dP)^2 - 1]$ is the χ^2 divergence and $\text{IPM}_{\mathcal{F}}(Q, R) := \sup_{f \in \mathcal{F}} \{\mathbb{E}_Q[f] - \mathbb{E}_R[f]\}$. Note that this problem takes the form of a projection in D_2 of P onto an IPM ball around Q .

By the definition of δ_{\min} , the constraint set is non-empty and convex. D_2 is strictly convex in R and $0 \leq D_2(R||P) < \infty$ so there is a unique solution.

When P already satisfies the IPM constraint, then $R = P$ has objective 0 and we're done (i.e. we don't need to do a projection, P is already on or in the ball). Otherwise, by standard use of the Lagrangian, we claim (details in Section A.5 below) that for some $\mu > 0$ corresponding to δ , this problem is equivalent to:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \left\{ (1/\mu)D_2(R||P) + \sup_{f \in \mathcal{F}} \{\mathbb{E}_Q[f] - \mathbb{E}_R[f]\} \right\}.$$

A.1 Exchanging hard subproblem for an easy subproblem

The inner supremum is hard to solve for arbitrary \mathcal{F} . However, we can make a series of transformations to get an easier subproblem with a closed-form solution:

$$\begin{aligned} & \inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \left\{ (1/\mu)D_2(R||P) + \sup_{f \in \mathcal{F}} \{\mathbb{E}_Q[f] - \mathbb{E}_R[f]\} \right\} \\ &= \inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \sup_{f \in \mathcal{F}} \left\{ (1/\mu)D_2(R||P) + \mathbb{E}_Q[f] - \mathbb{E}_R[f] \right\} \\ &= \sup_{f \in \mathcal{F}} \inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \left\{ (1/\mu)D_2(R||P) + \mathbb{E}_Q[f] - \mathbb{E}_R[f] \right\} \\ &= \sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_Q[f] + \inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \left\{ (1/\mu)D_2(R||P) - \mathbb{E}_R[f] \right\} \right\} \end{aligned}$$

The only non-trivial step is the interchange of the inf and the sup. This follows by Sion's Minimax Theorem (Sion, 1958). We assumed that \mathcal{X} was a separable Banach space so we have the necessary topological properties. The objective on the second line is continuous and strictly convex in R and is linear in f . The set \mathcal{F} is convex and closed and R is in a linear subspace. Furthermore, we know there is a unique solution R^* , and so we can always find the necessary compact subset of the linear subspace for R to apply the theorem (e.g. ala Ha, 1981).

A.2 Solving the easy subproblem with the variational representation

Next, we apply the variational representation of ϕ -divergences to get a dual formulation of the inner sub-problem over R . Define: $\phi(x) = (1/\mu)(x^2 - 1)$ which has convex conjugate $\phi^*(y) = (\mu/4)y^2 + (1/\mu)$. The Lagrangian for the infimum for a fixed f is:

$$\mathcal{L}_f(R, \lambda) = D_\phi(R||P) - \mathbb{E}_R[f - \lambda] - \lambda$$

We get the first-order condition:

$$\phi'(dR^*/dP) = f - \lambda^* \implies \frac{dR^*}{dP} = \frac{\mu}{2}(f - \lambda^*)$$

where λ^* solves the supremum $\sup_{\lambda \geq 0} g(\lambda)$ over the dual function:

$$\begin{aligned} g(\lambda) &:= -\lambda + \inf_{R \in \mathcal{M}(P)} \{D_\phi(R||P) - \mathbb{E}_R[f - \lambda]\} \\ &= -\lambda - \sup_{R \in \mathcal{M}(P)} \{\mathbb{E}_R[f - \lambda] - D_\phi(R||P)\} \\ &= -\lambda - D_\phi^*(f - \lambda), \end{aligned}$$

where D_ϕ^* is the convex conjugate of the ϕ -divergence as a function of R for a fixed P . We can then use the standard result (Proposition 4.2 Broniatowski and Keziou, 2006), $D_\phi^*(f) = \mathbb{E}_P[\phi^*(f)]$.

Using this form of the dual function, we can write our subproblem over R as:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \{(1/\mu)D_2(R||P) - \mathbb{E}_R[f]\} = \sup_{\lambda \geq 0} \{-\lambda - \mathbb{E}_P[\phi^*(f - \lambda)]\}$$

Plugging in ϕ^* , we can solve for λ^* by straightforward calculus:

$$\lambda^* = \mathbb{E}_P[f] - \frac{2}{\mu}.$$

Now using the first-order conditions, we can find the optimal R^* :

$$\frac{dR^*}{dP} = \frac{\mu}{2}(f - \mathbb{E}_P[f]) + 1$$

and after some algebra, a closed form of the subproblem:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \{(1/\mu)D_2(R||P) - \mathbb{E}_R[f]\} = -\mathbb{E}_P[f] - \frac{\mu}{4}\text{Var}_P[f].$$

A.3 Writing the original problem as a single optimization problem over \mathcal{F}

Finally, we substitute this form of the sub-problem into our original optimization problem:

$$\begin{aligned} &\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \left\{ (1/\mu)D_2(R||P) + \sup_{f \in \mathcal{F}} \{\mathbb{E}_Q[f] - \mathbb{E}_R[f]\} \right\} \\ &= \sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_Q[f] + \inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \{(1/\mu)D_2(R||P) - \mathbb{E}_R[f]\} \right\} \\ &= \sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_Q[f] - \mathbb{E}_P[f] - \frac{\mu}{4}\text{Var}_P[f] \right\} \end{aligned}$$

and therefore by duality:

$$\frac{dR^*}{dP} = \frac{\mu}{2}(f^* - \mathbb{E}_P[f^*]) + 1$$

where f^* achieves this supremum.

A.4 Recovering δ in terms of μ

Most of the proof of the theorem is complete. We just need to rewrite μ in terms of the original tuning parameter δ . Remember from the projection perspective, that $\mu > 0$ corresponds to P outside of the IPM ball. As a result:

$$\delta = \sup_{f \in \mathcal{F}} \{\mathbb{E}_Q[f] - \mathbb{E}_{R^*}[f]\}$$

We just proved that R^* achieves the infimum of the objective which equals the supremum of the dual:

$$\begin{aligned} & \text{IPM}_{\mathcal{F}}(Q, R^*) + (1/\mu)D_2(R^*||P) \\ &= \mathbb{E}_Q[f^*] - \mathbb{E}_P[f^*] - \frac{\mu}{4}\text{Var}_P[f^*] \\ &= \left(\mathbb{E}_Q[f^*] - \mathbb{E}_{R^*}[f^*]\right) + \left(\mathbb{E}_{R^*}[f^*] - \mathbb{E}_P[f^*] - \frac{\mu}{4}\text{Var}_P[f^*]\right) \end{aligned}$$

But then, using the variational representation of the ϕ -divergence and the definition of the IPM we have:

$$\begin{aligned} & \text{IPM}_{\mathcal{F}}(Q, R^*) + (1/\mu)D_2(R^*||P) \\ &= \sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_Q[f] - \mathbb{E}_{R^*}[f] \right\} + \sup_f \left\{ \mathbb{E}_{R^*}[f] - \mathbb{E}_P[f] - \frac{\mu}{4}\text{Var}_P[f^*] \right\} \\ &= \left(\mathbb{E}_Q[f^*] - \mathbb{E}_{R^*}[f^*]\right) + \left(\mathbb{E}_{R^*}[f^*] - \mathbb{E}_P[f^*] - \frac{\mu}{4}\text{Var}_P[f^*]\right) \end{aligned}$$

which implies

$$\delta = \sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_Q[f] - \mathbb{E}_{R^*}[f] \right\} = \mathbb{E}_Q[f^*] - \mathbb{E}_{R^*}[f^*]$$

Finally, substituting the form of R^* in terms of f^* we get

$$\begin{aligned} \delta &= \mathbb{E}_Q[f^*] - \mathbb{E}_P[f^*] - \frac{\mu}{2}\text{Var}_P[f^*] \\ \implies \mu &= 2 \left(\frac{\mathbb{E}_Q[f^*] - \mathbb{E}_P[f^*] - \delta}{\text{Var}_P[f^*]} \right), \end{aligned}$$

which concludes the proof.

A.5 Transformation from δ to μ via Lagrangian

Here we provide the details for our earlier claim that we can rewrite the problem over δ as a problem over μ . The dual function corresponding to the original δ problem is:

$$\begin{aligned} g(\mu) &= \inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \mathcal{L}(R, \mu) \\ &= \inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \left\{ D_2(R||P) + \mu \left(\sup_{f \in \mathcal{F}} \{ \mathbb{E}_Q[f] - \mathbb{E}_R[f] \} - \delta \right) \right\} \end{aligned}$$

where \mathcal{L} is the Lagrangian. Notice that the original optimization problem has a strictly convex objective. Furthermore, since the function class \mathcal{F} is convex and closed, $\delta > \delta_{\min}$, and the individual constraints for $f \in \mathcal{F}$ are all linear, the feasible set is convex with a non-empty interior. Then by standard convex duality there exists R^* and $\mu^* \geq 0$ such that R^* solves the original optimization problem, μ^* achieves $\sup_{\mu \geq 0} g(\mu)$, and R^* achieves the infimum inside $g(\mu^*)$.

By complementary slackness, $\mu^* = 0$ only when the worst-bias constraint doesn't bind which only occurs when $R = P$ already satisfies the IPM constraint. Then $R = P$ has minimum variance and we're done. So we only need to consider the case where $\mu^* > 0$.

At $\mu = \mu^*$, the solution to:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \left\{ D_2(R||P) + \mu \sup_{f \in \mathcal{F}} \{ \mathbb{E}_Q[f] - \mathbb{E}_R[f] \} \right\}$$

has the same solution as the original problem. Furthermore, since $\mu^* > 0$, we can apply one more transformation without affecting the infimum to get:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \left\{ (1/\mu)D_2(R||P) + \sup_{f \in \mathcal{F}} \{ \mathbb{E}_Q[f] - \mathbb{E}_R[f] \} \right\}$$

B General Statement and Proof for Remark 3.1

Theorem (Birrell et al). *Let ϕ be a convex function such that $\phi(1) = 0$ with convex conjugate ϕ^* such that $\{\phi^* < \infty\} = \mathbb{R}$. Let ϕ_μ denote the weighted function $\phi_\mu(x) = (1/\mu)\phi(x)$ and ϕ_μ^* its convex conjugate. Then under Assumption 2, for $\delta > 0$, $\exists \mu \geq 0$ such that the optimization problem,*

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} D_\phi(R||P)$$

such that $IPM_{\mathcal{F}}(Q, R) \leq \delta$,

has a solution,

$$R^* = P \text{ when } \mu = 0,$$

$$\frac{dR^*}{dP} = (\phi_\mu^*)'(f^* - \lambda^*) \text{ otherwise,}$$

where f^ and λ^* achieve the supremum,*

$$\sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_Q[f] - \inf_{\lambda \in \mathbb{R}} \{ \lambda + \mathbb{E}_P[\phi_\mu^*(f - \lambda)] \} \right\}.$$

The proof begins with an argument identical to 1.5 above which gives the equivalent optimization problem:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \{ D_{\phi_\mu}(R||P) + IPM_{\mathcal{F}}(Q, R) \}.$$

From here, conceptually, the proof is similar to above, except we cannot apply our proof directly because general ϕ loses some of the nice properties of the quadratic. Instead, using the theory of infimal convolutions, Birrell et al. (2020a) prove that for Polish \mathcal{X} and any ϕ -divergence such that $\{\phi^* < \infty\} = \mathbb{R}$:

$$\inf_{\substack{R \in \mathcal{M}(P) \\ \mathbb{E}_R[1]=1}} \{ D_\phi(R||P) + IPM_{\mathcal{F}}(Q, R) \}$$

$$= \sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_Q[f] - \inf_{\lambda \in \mathbb{R}} \{ \lambda + \mathbb{E}_P[\phi^*(f - \lambda)] \} \right\}$$

and so in particular, it holds for ϕ_μ above.

This result follows from Theorems 2.15 and 3.3 in Birrell et al. (2020a) with one minor modification: we do not require that $\lim_{y \rightarrow -\infty} \phi^*(y) < \infty$ which results in R^* no longer being a probability measure because we lose statement (174) in their proof of Theorem C.6. This is in the spirit of the arguments in (Broniatowski and Keziou, 2006). In fact, since the original problem is a projection of P in ϕ -divergence onto an IPM, we can interpret this as a version of (Broniatowski and Keziou, 2006) Theorem 5.1 which applies for the case of finitely-many linear inequality constraints, generalized to the case of a linear inequality constraint for each f in a convex and closed set \mathcal{F} .

C Extension with non-negative weights

We can use the general Theorem to immediately get results for the case where we require non-negative weights. This is identical to taking $\phi(x) = x^2 - 1$, but restricting the domain to $[0, \infty)$. Then we have

$$\phi_\mu^*(y) = \frac{\mu}{4}y^2\mathbf{1}(y \geq 0) + \frac{1}{\mu}$$

and applying the theorem, we get the dual formulation:

$$\frac{dR^*}{dP} = \frac{\mu}{2}(f^* - \lambda^*)\mathbf{1}(f^* \geq \lambda^*)$$

such that:

$$\mathbb{E}_P \left[\frac{\mu}{2}(f^* - \lambda^*)\mathbf{1}(f^* \geq \lambda^*) \right] = 1$$

and f^* and λ^* solve:

$$\sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_Q[f] - \inf_{\lambda \in \mathbb{R}} \left\{ \lambda + \frac{\mu}{4} \mathbb{E}_P [(f - \lambda)^2 \mathbf{1}(f \geq \lambda)] \right\} + \frac{1}{\mu} \right\}.$$

The minimax weights are extremely similar to the minimax weights from Theorem 3.1. In Theorem 3.1 we found a function f^* , such that we de-meaned it, rescaled it, and then shifted the result to get weights with expectation equal to 1. In the case where the weights are non-negative, we can no longer just de-mean and add 1. Instead, we have to optimize over all shifts λ which give expectation 1 *after* truncating at 0.

D Connection to surrogate loss for density ratio estimation

A closely related literature implements regularized estimators of the density ratio via a surrogate loss as first proposed in Nguyen et al. (2010). Similar connections have been made for balancing weights; see the relation to propensity score estimation in Zhao (2019); Ben-Michael et al. (2021). Here, we compare our dual formulation based on variational representations to density ratio estimation with a surrogate loss.

Consider the variational representation (10). Specializing to $\phi(x) = \mu(x^2 - 1)$ for $\mu > 0$, we can write the weighted χ^2 divergence between Q and P as:

$$\frac{1}{\mu} D_2(Q||P) = \sup_f \left\{ \mathbb{E}_Q[f] - \mathbb{E}_P[f] - \frac{\mu}{4} \text{Var}_P[f] \right\},$$

where the supremum is over all real-valued measurable functions. If overlap holds then the supremum is achieved by $2\mu(dQ/dP)$, otherwise $D_2(Q||P) = \infty$. The variational representation is identical to (11), except that the problem in (11) is restricted to functions in \mathcal{F} . We immediately have the following corollary:

Corollary D.1. *Under the conditions in Theorem 4.1,*

$$\frac{2}{\mu} \frac{dQ}{dP} \in \mathcal{F} \implies f^* = \frac{2}{\mu} \frac{dQ}{dP} \text{ and } \frac{dR^*}{dP} = \frac{dQ}{dP}.$$

If the density ratio exists and a scaled version belongs to the outcome function class, then it is minimax optimal to reweight so that there is zero bias. If \mathcal{F} is sufficiently flexible, then this will always hold as $\mu \rightarrow \infty$, as in the example from Section 5.1.

There is a clear similarity to density ratio estimation using ϕ -divergences as in Nguyen et al. (2010). The form of their final estimator looks the same as ours: maximizing a variational representation over a function class. However, our motivations and assumptions are entirely different. Nguyen et al. (2010) assumes that the density ratio exists and that it belongs to \mathcal{F} . We do not make any assumption about the density ratio. It could have some arbitrary functional form. We show that the shape of the optimal weights is determined by the shape of f_0 , *not* by the shape of dQ/dP .

If the density ratio exists, and our dual problem is solved over all measurable functions, the only unique solution will be dQ/dP . In that sense, once we restrict to optimization over \mathcal{F} , it might be helpful to think of the balancing weights as a projection of the density ratio onto our outcome function class. However, in general, we do not need a functional form for the density ratio to be specified. In fact, we do not need even require that the density ratio exists.

E RKHS optimization problem

The optimization problem that implements our dual formulation using an RKHS for the IHDP dataset is:

$$\begin{aligned} & \sup_{f \in \tilde{\mathcal{F}}_{\mathcal{H}}^B} \left\{ \mathbb{E}_Q[f] - \mathbb{E}_P[f] - \frac{\mu}{4} \text{Var}_P[f] \right\} \\ &= \sup_{\substack{f=K\alpha \\ \alpha \in \mathbb{R}^n: \alpha^T K \alpha \leq B}} \left\{ f^T e_q - f^T e_p - (f^T I_p f - (f^T e_p)^2) \right\} \end{aligned}$$

where:

- $e_p \in \mathbb{R}^n$ is a vector equal to $1/n_0$ for those indices corresponding to control group data points and 0 otherwise
- $e_q \in \mathbb{R}^n$ is a vector equal to $1/n_0$ for those indices corresponding to treatment group data points and 0 otherwise
- $I_p \in \mathbb{R}^{n \times n}$ is the matrix with e_p on the diagonal.

We solve this problem using the scipy Python package.