

Externally Valid Treatment Choice

Christopher Adjaho* Timothy Christensen†

April 22, 2022

Abstract

We consider the problem of learning treatment (or policy) rules that are externally valid in the sense that they have welfare guarantees in target populations that are similar to, but possibly different from, the experimental population. We allow for shifts in both the distribution of potential outcomes and covariates between the experimental and target populations. This paper makes two main contributions. First, we provide a formal sense in which policies that maximize social welfare in the experimental population remain optimal for the “worst-case” social welfare when the distribution of potential outcomes (but not covariates) shifts. Hence, policy learning methods that have good regret guarantees in the experimental population, such as empirical welfare maximization, are externally valid with respect to a class of shifts in potential outcomes. Second, we develop methods for policy learning that are robust to shifts in the joint distribution of potential outcomes and covariates. Our methods may be used with experimental or observational data.

Keywords: Policy learning, Individualized treatment rules, Empirical welfare maximization, External validity, Distributionally robust optimization, Optimal transport.

*Department of Economics, New York University. ca2384@nyu.edu

†Department of Economics, New York University. timothy.christensen@nyu.edu

1 Introduction

Over the past two decades, a large literature has developed in economics, machine learning, and biostatistics on learning policy rules from experimental or quasi-experimental data.¹ The typical objective is to choose a policy rule mapping individual covariates into a treatment choice to maximize average welfare, mean response, or expected reward. The maintained assumption underlying this literature is that the *target* population in which the policy is to be implemented is the same as the *experimental* (or training) population from which the data are sampled.² While policy learning algorithms may have good welfare guarantees when the target and experimental populations are the same, little is known about the performance of the learned policies in target populations that are similar to, but meaningfully different from, the experimental population.

There are many reasons why the target and experimental populations may differ. For instance, data may be collected in a randomized controlled trial (RCT) under artificial experimental conditions that differ from real-world settings in which the policy is to be implemented. Alternatively, there may be a delay between collecting the data and implementing the policy, allowing for the possibility that the distribution of potential outcomes and/or covariates may drift. The RCT may also define outcomes in terms of certain measurable dimensions (e.g. test scores) while the policy maker may be concerned with more difficult to quantify outcomes that are similar to, but slightly different from, those measured in the RCT (e.g. overall academic achievement). Moreover, due to budget constraints, experiments may be run in a few sub-populations, while the policy is to be implemented in additional similar sub-populations beyond those studied in the experiment (e.g. neighboring school districts).

In this paper, we consider the problem of learning treatment (or policy) rules that are *externally valid* in the sense that they have welfare guarantees in target populations that are similar to, but possibly different from, the experimental population. We allow for shifts in both the distribution of potential outcomes and covariates between the experimental and target populations. We propose methods for learning externally valid rules using experimental or observational data (where treatment is possibly endogenous).

¹In economics, the literature goes back to [Manski \(2004\)](#) and includes [Bhattacharya and Dupas \(2012\)](#), [Stoye \(2012\)](#), [Kitagawa and Tetenov \(2018\)](#), [Mbakop and Tabord-Meehan \(2021\)](#), and [Athey and Wager \(2021\)](#). References in machine learning and biostatistics include [Qian and Murphy \(2011\)](#), [Zhao, Zeng, Rush, and Kosorok \(2012\)](#), [Swaminathan and Joachims \(2015\)](#), and [Kallus \(2017, 2021\)](#), among many others.

²We refer to the distribution from which the data are sampled as “experimental”, with the understanding that the analyst could in fact be using observational (i.e., non-experimental) data.

More formally, consider a class of policies τ mapping an individual’s covariates X into a binary outcome $\tau(X)$, where $\tau(X) = 1$ indicates that treatment is to be assigned to the individual and $\tau(X) = 0$ indicates otherwise. Following Manski (2004), treatment rules are typically evaluated using a *social welfare* criterion

$$W(\tau) = \mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))], \quad (1)$$

where Y_0 and Y_1 denote the individual’s untreated and treated potential outcomes, and $\mathbb{E}_P[\cdot]$ denotes expectation under the distribution P of (X, Y_0, Y_1) in the experimental population.³ The typical objective is to learn a policy that maximizes W over a class \mathcal{T} of policies that may incorporate functional-form, budget, fairness, or other constraints.

Our objective is to derive externally valid treatment rules that maximize worst-case social welfare over a family \mathcal{Q} of distributions of (X, Y_0, Y_1) representing target populations that are similar to, but possibly different from, the experimental population. To this end, we replace criterion (1) with the *robust welfare* criterion

$$RW(\tau) = \inf_{Q \in \mathcal{Q}} \mathbb{E}_Q[Y_1\tau(X) + Y_0(1 - \tau(X))]. \quad (2)$$

We seek to learn an optimal policy τ^* that maximizes RW over \mathcal{T} , given data sampled under P . This max-min approach ensures that τ^* delivers welfare guarantees uniformly over all potential target populations $Q \in \mathcal{Q}$.

A novel aspect of our approach is that we use a class of Wasserstein metrics to define \mathcal{Q} . This has three main advantages. First, it serves as a unifying framework in which to study robustness to shifts in the distribution of potential outcomes only, and shifts in the distribution of both potential outcomes and covariates. Second, it leads to a tractable characterization of the robust welfare objective (2). Third, the size parameter ε used to define \mathcal{Q} is precisely the maximum that the average treatment effect (ATE) can differ between the experimental population P and target populations $Q \in \mathcal{Q}$, as we show formally in Section 2. As such, the size of \mathcal{Q} is very interpretable and may easily be calibrated by the analyst.

Section 3 considers robustness to shifts in the distribution of potential outcomes. Proposition 3.1 characterizes the robust welfare criterion (2) in this case. A number of implications then follow from this characterization. In particular, rules that are optimal

³This criterion is often referred to as the mean response in biostatistics and expected reward in machine learning.

under the social welfare criterion (1) are also optimal under the robust criterion (2). Moreover, policy learning methods that have good regret guarantees under criterion (1), such as those proposed by Manski (2004), Qian and Murphy (2011), Kitagawa and Tetenov (2018), and Mbakop and Tabord-Meehan (2021), to name a few, also enjoy good regret guarantees under the robust welfare criterion (2).⁴

To handle known covariate shifts between the experimental and target populations, Kitagawa and Tetenov (2018, Remark 2.2), Uehara, Kato, and Yasui (2020), and Kallus (2021), amongst others, study a re-weighted version of (1). Under their re-weighting, the argument of the expectation is multiplied by the ratio $q(X)/p(X)$ of the densities of X in the target (q) and experimental (p) populations. If the conditional average treatment effect $\mathbb{E}[Y_1 - Y_0|X = x]$ is the same in both populations, then the re-weighted criterion corresponds to social welfare in the target population. However, if there is a shift in potential outcomes as well, then this equivalence can break down. We show that rules that are preferred under the re-weighted criterion are also preferred under a robust criterion that considers the worst-case welfare over a class of distributions with unknown shift in potential outcomes and known shift (from p to q) in covariates. Hence, re-weighted criteria are more robust to covariate shifts than they may otherwise appear.

In Section 4 we then consider the problem of robustness to shifts in both potential outcomes and covariates. The robust welfare criterion (2) is derived in Proposition 4.1. In this case, the robust criterion may induce a different ordering over treatment rules than criterion (1). The robust criterion depends on two functions representing “distance to non-treatment” and “distance to treatment” under τ . We show how to compute these functions for linear eligibility score rules, threshold rules, and decision trees.

Section 4.2 discusses empirical implementation. Unlike the case of shifts in potential outcomes only, here the robust welfare criterion depends on the joint distribution of (X, Y_0, Y_1) in the experimental population. As only the marginals for (X, Y_0) and (X, Y_1) are identified without further assumptions even in the context of a RCT,⁵ we provide a non-exhaustive set of methods for constructing empirical robust welfare criteria under different identifying assumptions from the literature on distributional treatment effects (see, e.g., Abbring and Heckman (2007) for a review). Our methods allow for both experimental and observational data. For example, we show that the empirical ro-

⁴We show in Appendix A that these findings extend to when welfare is measured using the criterion $\mathbb{E}_P[Y_1\tau(X) - Y_0]$ considered by Athey and Wager (2021).

⁵These marginals may be identified under additional assumptions with observational data.

bust welfare criterion may be easily computed by linear programming under a standard monotonicity (or sorting) assumption. Further, we show that the estimated rule $\hat{\tau}$ that maximizes the empirical criterion is asymptotically optimal, in the sense that the regret $\sup_{\tau \in \mathcal{T}} \text{RW}(\tau) - \text{RW}(\hat{\tau})$ converges in probability to zero for all policy classes \mathcal{T} under general conditions.

The different identifying assumptions in Section 4.2 correspond to different couplings of the conditional distributions of Y_0 and Y_1 given X in the experimental population. In Section 4.3 we take a conservative approach in the spirit of Imbens and Menzel (2021) and study a criterion under a *least-favorable* coupling. This criterion only requires knowledge of the marginals for (X, Y_0) and (X, Y_1) in the experimental population. Rules based on this criterion have welfare guarantees with respect to both adversarial shifts in the joint distribution of potential outcomes and covariates and adversarial couplings of the conditional distribution of Y_0 and Y_1 given X in the experimental population. We provide an empirical counterpart which can be constructed from experimental or observational data and can be easily computed by linear programming.

We then conclude in Section 5 with an empirical application that revisits the Job Training Partnership act study analyzed by Bloom et al. (1997). For the class of threshold rules considered by Kitagawa and Tetenov (2018), our findings show that the empirical welfare maximizing rule is fairly robust to shifts in both potential outcomes and covariates, at least relative to other rules in this class.

As a sensitivity analysis, researchers could report plots similar to those we present in Section 5, which compare the worst-case welfare of different rules across different sized families of target populations.

Our focus on robustness to shifts in potential outcomes (and possibly covariates) complements recent work by Mo, Qi, and Liu (2021) and Spini (2021) on treatment rules that are robust to shifts in covariates only. These papers propose using a criterion based on the worst-case re-weighting of the covariate distribution over a f -divergence neighborhood. A key assumption underlying their approach is that the conditional average treatment effect does not change between the experimental and target populations. Si, Zhang, Zhou, and Blanchet (2020) study robust policy choice in contextual bandit problems over Kullback–Leibler neighborhoods. The different neighborhood constructions used in these papers leads to a very different robust criterion function than we derive. Unlike Wasserstein metrics, Kullback–Leibler and other f -divergences prohibit extrapolation beyond the support of the experimental population, which may be restrictive in applications. The notion of neighborhood size is also not as easily interpretable in

these works as it is in ours. Qi, Pang, and Liu (2022) propose choosing treatment rules to maximize a type of value-at-risk criterion, which they show has robustness properties to a class of distribution shifts.⁶ None of these papers studies robustness with respect to shifts in the distribution of potential outcomes only, or implementation with observational data (i.e., where treatment is possibly endogenous).

2 Wasserstein Neighborhoods

We consider families of target populations \mathcal{Q} that are neighborhoods of the experimental population P under appropriate choices of Wasserstein metric. This section introduces the Wasserstein neighborhoods we work with and shows that the radius of \mathcal{Q} has a clear interpretation as the maximal change in the ATE between the experimental and target populations.

Let \mathcal{Z} denote the support of $Z := (X, Y_0, Y_1)$ and let $d : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}_+ \cup \{+\infty\}$ be a (pseudo-)metric on \mathcal{Z} . The *Wasserstein (pseudo-)metric* of order p between P and Q is

$$W_p(P, Q) = \inf_{\pi \in \Pi(P, Q)} \mathbb{E}_\pi [d(Z, \tilde{Z})^p]^{\frac{1}{p}}, \quad (1 \leq p < \infty),$$

where $\Pi(P, Q)$ denotes all probability distributions for the coupling (Z, \tilde{Z}) with marginals P for Z and Q for \tilde{Z} . We will focus on the metric with $p = 1$, so we drop the p subscript and write $W(P, Q)$ in what follows. Some generalizations to $p > 1$ are presented in Appendix A. In the following sections, we define Wasserstein neighborhoods using

$$d((x, y_0, y_1), (\tilde{x}, \tilde{y}_0, \tilde{y}_1)) = |y_0 - \tilde{y}_0| + |y_1 - \tilde{y}_1| + b(x, \tilde{x}), \quad (3)$$

for a metric b . We shall use different choices of b to separately handle robustness with respect to shifts in potential outcomes (as in Section 3) or shifts in potential outcomes and covariates (as in Section 4).

We seek to maximize worst-case welfare (2) over

$$\mathcal{Q} = \{Q : W(P, Q) \leq \varepsilon\}, \quad (4)$$

where $\varepsilon > 0$ is a measure of neighborhood size. The parameter ε has a clear interpretation

⁶The class of shifts does not allow extrapolation outside the support of P and has a less interpretable notion of “size”.

as the maximum change in the ATE as Q varies over \mathcal{Q} . This interpretation holds irrespective of whether we are seeking robustness with respect to shifts in potential outcomes only or both potential outcomes and covariates. First, we introduce some notation. Let \mathcal{Y} denote the support⁷ of the potential outcomes, and let $\underline{Y} = \inf \mathcal{Y}$ and $\bar{Y} = \sup \mathcal{Y}$. For instance, $\mathcal{Y} = \{0, 1\}$, $\underline{Y} = 0$, and $\bar{Y} = 1$ for binary outcomes. We allow $\underline{Y} = -\infty$ and ($\bar{Y} = +\infty$) for outcomes that are unbounded from below (from above). We say that potential outcomes are *unbounded* if $\underline{Y} = -\infty$ or $\bar{Y} = +\infty$ (or both).

Proposition 2.1 *Suppose that \mathcal{Q} is defined using the Wasserstein metric $W(P, Q)$ induced by (3). Then*

$$\begin{aligned}\inf_{Q \in \mathcal{Q}} \mathbb{E}_Q [Y_1 - Y_0] &= \max \{ \mathbb{E}_P [Y_1 - Y_0] - \varepsilon, \underline{Y} - \bar{Y} \}, \\ \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q [Y_1 - Y_0] &= \min \{ \mathbb{E}_P [Y_1 - Y_0] + \varepsilon, \bar{Y} - \underline{Y} \}.\end{aligned}$$

If potential outcomes are unbounded, then

$$\begin{aligned}\inf_{Q \in \mathcal{Q}} \mathbb{E}_Q [Y_1 - Y_0] &= \mathbb{E}_P [Y_1 - Y_0] - \varepsilon, \\ \sup_{Q \in \mathcal{Q}} \mathbb{E}_Q [Y_1 - Y_0] &= \mathbb{E}_P [Y_1 - Y_0] + \varepsilon.\end{aligned}\tag{5}$$

Remark 2.1 The bounds (5) apply for unbounded outcomes and also for bounded outcomes provided ε is small enough that $\mathbb{E}_P [Y_1 - Y_0]$ is more than ε from $\underline{Y} - \bar{Y}$ and $\bar{Y} - \underline{Y}$. Moreover, inspection of the proof of Proposition 2.1 shows these bounds are sharp: there exist distributions $Q \in \mathcal{Q}$ for which $\mathbb{E}_Q [Y_1 - Y_0] = \mathbb{E}_P [Y_1 - Y_0] \pm \varepsilon$. As such, Proposition 2.1 provides a formal sense in which the neighborhood size ε is precisely the maximum that the ATE can vary between the experimental population P and target populations $Q \in \mathcal{Q}$.

3 Shifts in Potential Outcomes

In this section, we consider robustness with respect to shifts in the distribution of potential outcomes (but not covariates) between the experimental and target populations.

⁷By “support” we mean the set of all values the potential outcomes could conceivably take, as distinct from the measure-theoretic notion of support.

This is relevant in a number of scenarios. First, it is relevant for designing treatment rules to maximize welfare for the experimental population when welfare is measured using outcomes (e.g. proxies) that are similar to, but possibly different from, the outcomes measured in the RCT. For instance, the policy maker may care about overall educational attainment, while the RCT may measure outcomes in terms of a specific test score. It is also relevant when the experimental conditions differ from the real-world setting in which the policy is to be implemented.⁸ Thus, subjects’ responses may be similar to, but slightly different from, their responses under the experimental conditions. This approach also accommodates settings without covariates as a special case.

We start in Section 3.1 by introducing an appropriate neighborhood construction to handle this problem. Proposition 3.1 derives the robust welfare criterion in this case. We then discuss several implications of this result. In particular, rules that are preferred under the social welfare criterion (1) are also preferred under the robust criterion. Moreover, policy learning methods that have good regret guarantees under criterion (1) also enjoy regret guarantees under the robust criterion.

As we discuss in Section 3.2, this approach and findings extend to handle settings with shifts in covariates of a known form but unspecified shift in potential outcomes. For instance, the RCT may sample one school district while the analyst wishes to design a policy for a larger set of neighboring school districts. The distribution of covariates in both districts may be known e.g. from census or administrative data, and we might reasonably assume that the distribution of potential outcomes is similar, but not the same, across districts.

3.1 Robust Welfare Criterion

To allow for shifts in the distribution of potential outcomes while holding the distribution of covariates fixed, we define Wasserstein neighborhoods (4) using the metric

$$d((x, y_0, y_1), (\tilde{x}, \tilde{y}_0, \tilde{y}_1)) = |y_0 - \tilde{y}_0| + |y_1 - \tilde{y}_1| + \infty \times \mathbb{I}[x \neq \tilde{x}]. \quad (6)$$

This metric combines the ℓ^1 norm to penalize the shifts in the distribution of potential outcomes and a discrete (extended) metric on the support \mathcal{X} of X to prohibit any shift

⁸In this case and the previous one, we could think about individuals being characterized by $(X, Y_0, Y_1, Y_0^*, Y_1^*)$, where (Y_0, Y_1) are the experimental-condition outcomes (or measured outcomes) and (Y_0^*, Y_1^*) are the real-world setting outcomes (or policy-relevant outcomes). The distribution P is the marginal for (X, Y_0, Y_1) while Q is the marginal for (X, Y_0^*, Y_1^*) .

in the distribution of covariates. In models without covariates we simply have

$$d((x, y_0, y_1), (\tilde{x}, \tilde{y}_0, \tilde{y}_1)) = |y_0 - \tilde{y}_0| + |y_1 - \tilde{y}_1|.$$

In this case, the metric $W(P, Q)$ reduces to the classical Wasserstein metric W_1 between P and Q .

The following result characterizes the robust welfare criterion (2) when the family of target populations \mathcal{Q} is defined using (6). Recall from Section 2 that \mathcal{Y} is the support of the potential outcomes and $\underline{Y} = \inf \mathcal{Y}$. We allow $\underline{Y} = -\infty$ for outcomes that are unbounded from below. To simplify the proofs when $\underline{Y} = -\infty$, we assume that \mathcal{Y} is *equispaced*, i.e., there is finite $C > 0$ such that for any $y \in \mathcal{Y}$ the set $[y - C, y) \cap \mathcal{Y}$ is nonempty. Note that commonly used supports such as \mathbb{R} and \mathbb{Z} satisfy this condition.

Proposition 3.1 *Suppose that \mathcal{Q} is defined using the Wasserstein metric $W(P, Q)$ induced by (6). Then for any treatment rule τ*

$$\text{RW}(\tau) = \max \left\{ \mathbb{E}_P [Y_1 \tau(X) + Y_0 (1 - \tau(X))] - \varepsilon, \underline{Y} \right\}.$$

If potential outcomes are unbounded from below, then the criterion simplifies to

$$\text{RW}(\tau) = \mathbb{E}_P [Y_1 \tau(X) + Y_0 (1 - \tau(X))] - \varepsilon.$$

Remark 3.1 Proposition 3.1 implies that any treatment rule maximizing the social welfare criterion (1) for the experimental population must also maximize the robust welfare criterion (2), irrespective of the neighborhood size ε . Indeed, if $W(\tau) \geq W(\tau')$ then $\text{RW}(\tau) \geq \text{RW}(\tau')$ for any treatment rules $\tau, \tau' \in \mathcal{T}$. Moreover, it follows from Proposition 3.1 that the regret of any (possibly data-dependent) rule $\hat{\tau}$ under the robust criterion is bounded by its regret under criterion (1):

$$\sup_{\tau \in \mathcal{T}} \text{RW}(\tau) - \text{RW}(\hat{\tau}) \leq \sup_{\tau \in \mathcal{T}} W(\tau) - W(\hat{\tau}),$$

irrespective of ε . Hence, policy learning methods that have good regret guarantees under criterion (1), such as those of Manski (2004), Qian and Murphy (2011), Kitagawa and Tetenov (2018), and Mbakop and Tabord-Meehan (2021), to name a few, also enjoy good regret guarantees under the robust welfare criterion (2).

Remark 3.2 In some settings \mathcal{Y} will be a strict subset of \mathbb{R} . For instance, $\mathcal{Y} = \{0, 1\}$ for binary outcomes. In this case, there may be many distributions $Q \in \mathcal{Q}$ under which the support of potential outcomes is different from \mathcal{Y} . This raises the concern that the neighborhoods \mathcal{Q} are “too large”, in the sense that they contain distributions with supports that the analyst would never confront in any realistic target population. Inspection of the proof of Proposition 3.1 shows that if either $\min \mathcal{Y} = \underline{Y} > -\infty$ or $\underline{Y} = -\infty$ hold, then the *worst-case* distributions that solve the minimization problem (2) have support \mathcal{Y} . The neighborhoods \mathcal{Q} are therefore not too large, because the worst-case distributions that are being guarded against are precisely those that respect this support condition.

3.2 Extension: Known Covariate Shifts

In some cases the analyst may be able to estimate (or have prior knowledge of) the marginal distribution of covariates in the target population. For instance, a RCT may sample one school district while the analyst wishes to design a policy for a larger set of neighboring school districts. The distribution of covariates in both districts may be known e.g. from census or administrative data, and we might reasonably assume that the distribution of potential outcomes is similar, but possibly not the same, across districts.

In this case, one could consider the re-weighted social welfare criterion

$$\mathbb{E}_P [(Y_1\tau(X) + Y_0(1 - \tau(X))) \rho(X)] \tag{7}$$

where $\rho(x) = q(x)/p(x)$ is the ratio of the marginal densities for X in the target and experimental populations, respectively. This re-weighted criterion (or sample versions thereof) was recently considered by [Kitagawa and Tetenov \(2018, Remark 2.2\)](#), [Uehara et al. \(2020\)](#), and [Kallus \(2021\)](#), amongst others, for policy learning in the presence of a known covariate shift between the experimental and target populations.

Criterion (7) corresponds to social welfare in the target population provided the conditional average treatment effect (CATE) is the same under P and Q . Otherwise, criterion (7) may not equal social welfare in the target population. Note that a shift in the CATE between the experimental and target populations implies the distribution of potential outcomes shifts between the experimental and target populations.

To study the implications of shifts in potential outcomes in this setting, we consider the weighted robust criterion

$$\inf_{Q \in \mathcal{Q}'} \mathbb{E}_Q [(Y_1\tau(X) + Y_0(1 - \tau(X))) \rho(X)] \tag{8}$$

where $\mathcal{Q}' = \{Q : W'(P, Q) \leq \varepsilon\}$ with

$$W'(P, Q) = \inf_{\pi \in \Pi(P, Q)} \mathbb{E}_\pi[d(Z, \tilde{Z})\rho(X)]$$

for $Z = (X, Y_0, Y_1)$, and where d is the metric (6). The robust criterion (8) gives the worst-case social welfare over distributions Q where X has marginal density q and where the marginals for the potential outcomes are similar to, but different from, the experimental population. An identical argument to Proposition 3.1 yields

$$\begin{aligned} & \inf_{Q \in \mathcal{Q}'} \mathbb{E}_Q [(Y_1\tau(X) + Y_0(1 - \tau(X))) \rho(X)] \\ & = \max \left\{ \mathbb{E}_P [(Y_1\tau(X) + Y_0(1 - \tau(X))) \rho(X)] - \varepsilon, \underline{Y} \right\}. \end{aligned}$$

The implications discussed in Remark 3.1 also carry over to this re-weighted criterion. In particular, any treatment rule that maximizes the re-weighted social welfare criterion (7) must also maximize its robust counterpart (8).

4 Shifts in Potential Outcomes and Covariates

We now turn to the problem of external validity when we allow for shifts in the joint distribution of covariates X and potential outcomes (Y_0, Y_1) .

We begin in Section 4.1 by characterizing the robust welfare criterion in this setting. The criterion takes the form of a scalar optimization problem involving an expectation over potential outcomes and certain functions representing the distance of covariates to treatment and non-treatment regions. We show how to compute these distance functions for linear eligibility score rules, threshold rules, and decision trees.

Section 4.2 discusses the empirical implementation of externally valid treatment rules in this setting, allowing for experimental and observational data. Empirical implementation is more complicated here because the robust welfare criterion depends on the full joint distribution P of (X, Y_0, Y_1) , which is not identified—even within the context of a RCT—without further assumptions. We establish asymptotic optimality of a number of empirical approaches to estimating externally valid treatment rules under different sets of identifying assumptions, allowing for experimental and observational data.

Finally, in Section 4.3 we consider a criterion based on a least-favorable coupling of the conditional distribution of Y_0 and Y_1 given X . This criterion can be estimated

without any further identifying assumptions. Rules maximizing this criterion are robust with respect to both shifts in potential outcomes and covariate, and adversarial couplings of the conditional distribution of Y_0 and Y_1 given X . This approach and some of those from Section 4.2 can be easily implemented via linear programming.

4.1 Robust Welfare Criterion

To allow for shifts in the distribution of potential outcomes and covariates, we define Wasserstein neighborhoods (4) using

$$d((x, y_0, y_1), (\tilde{x}, \tilde{y}_0, \tilde{y}_1)) = |y_0 - \tilde{y}_0| + |y_1 - \tilde{y}_1| + \|x - \tilde{x}\|, \quad (9)$$

where $\|\cdot\|$ is a norm on \mathcal{X} .

The following result derives the robust welfare criterion (2) when the family of target distributions \mathcal{Q} is defined using (6). We again let \underline{Y} denote the (sharp) lower bound on the supports of the potential outcomes Y_0 and Y_1 .

Proposition 4.1 *Suppose that \mathcal{Q} is defined using the Wasserstein metric $W(P, Q)$ induced by (9) and that $\mathbb{E}_P[\|X\|]$ is finite. Then for any treatment rule τ ,*

$$\text{RW}(\tau) = \max \left\{ \sup_{\eta \geq 1} \mathbb{E}_P [\min \{Y_0 + \eta h_0(X), Y_1 + \eta h_1(X)\}] - \eta \varepsilon, \underline{Y} \right\}, \quad (10)$$

where

$$h_0(x) = \inf_{\tilde{x} \in \mathcal{X}: \tau(\tilde{x})=0} \|x - \tilde{x}\|, \quad h_1(x) = \inf_{\tilde{x} \in \mathcal{X}: \tau(\tilde{x})=1} \|x - \tilde{x}\|,$$

with the understanding that $h_0(x) = +\infty$ or $h_1(x) = +\infty$ if the infimum runs over an empty set. If potential outcomes are unbounded from below, then criterion (10) simplifies:

$$\text{RW}(\tau) = \sup_{\eta \geq 1} \mathbb{E}_P [\min \{Y_0 + \eta h_0(X), Y_1 + \eta h_1(X)\}] - \eta \varepsilon.$$

Remark 4.1 The functions $h_0(x)$ and $h_1(x)$ represent the “distance to non-treatment” and “distance to treatment” under τ for an individual with covariates $X = x$. Different treatment rules τ correspond to different h_0 and h_1 . We suppress dependence of h_0 and h_1 on τ to simplify notation. Note that at most one of $h_0(x)$ and $h_1(x)$ is nonzero for each $x \in \mathcal{X}$.

As we now show, the functions h_0 and h_1 have closed-form expressions or are otherwise efficient to compute for some popular classes of treatment rule. To simplify exposition, we take $\mathcal{X} = \mathbb{R}^d$ and let $\|\cdot\|$ be the Euclidean norm.

Example: Linear Eligibility Score Rules. Consider a treatment rule of the form

$$\tau(x) = \mathbb{1}\{\beta_0 + x'\beta_1 \geq 0\},$$

parameterized by β_0 (a scalar) and β_1 (a vector), where at least one element of β_1 is non-zero. The interpretation of the rule is that treatment is assigned when the eligibility score $x'\beta_1$ exceeds the threshold $-\beta_0$. Here

$$h_0(x) = \inf_{\tilde{x}:\beta_0+\tilde{x}'\beta_1<0} \|x - \tilde{x}\|, \quad h_1(x) = \inf_{\tilde{x}:\beta_0+\tilde{x}'\beta_1\geq 0} \|x - \tilde{x}\|.$$

The values $h_0(x)$ and $h_1(x)$ are the minimum distances from x to the half-spaces $\{\tilde{x} : \beta_0 + \tilde{x}'\beta_1 \leq 0\}$ and $\{\tilde{x} : \beta_0 + \tilde{x}'\beta_1 \geq 0\}$, namely

$$h_0(x) = \frac{(\beta_0 + x'\beta_1)_+}{\|\beta_1\|}, \quad h_1(x) = \frac{(\beta_0 + x'\beta_1)_-}{\|\beta_1\|},$$

where $(a)_- = -\min\{a, 0\}$ and $(a)_+ = \max\{a, 0\}$. □

Example: Threshold (or Quadrant) Rules. Threshold rules assign treatment when the values of certain covariates are above or below given thresholds. For illustrative purposes we consider a rule that depends on the first two components of X :

$$\tau(x) = \mathbb{1}\{s_1(x_1 - \beta_1) \geq 0 \text{ and } s_2(x_2 - \beta_2) \geq 0\}$$

for $s_1 = -1$ and $s_2 = 1$ and $\beta_1, \beta_2 \in \mathbb{R}$. For this rule, $\tau(x) = 1$ if and only if $x_1 - \beta_1 \leq 0$ and $x_2 - \beta_2 \geq 0$. Similarly, $\tau(x) = 0$ if and only if $x_1 - \beta_1 < 0$ or $x_2 - \beta_2 > 0$. It follows that

$$h_0(x) = \sqrt{\min\{(x_1 - \beta_1)_-^2, (x_2 - \beta_2)_+^2\}}, \quad h_1(x) = \sqrt{(x_1 - \beta_1)_+^2 + (x_2 - \beta_2)_-^2}.$$

Derivations of h_0 and h_1 for rules with fewer or more thresholds follow similarly. □

Example: Decision Trees. Decision trees are defined by recursively partitioning \mathcal{X} into a set of regions for which the values of certain covariates lie above or below given

thresholds. Treatment is then assigned depending on whether covariates lie in particular subsets of the resulting partition. Rules based on decision trees may be expressed as

$$\tau(x) = \mathbb{1}\{x \in \cup_{k=1}^K C_k\}$$

where each C_k is a hypercube. Different decision trees, and hence different rules τ , correspond to different sets $\cup_{k=1}^K C_k$. Note the non-treatment region $\mathcal{X} \setminus (\cup_{k=1}^K C_k)$ is itself the union of finitely many nonempty hypercubes, say $\cup_{l=1}^L \tilde{C}_l$. The squared distance from x to a hypercube C can be efficiently solved via quadratic programming:

$$\text{dist}(x, C)^2 = \min_y \|x - y\|^2 \quad \text{subject to} \quad l_i \leq y_i \leq u_i, \quad i = 1, \dots, d,$$

where l_i and u_i , $i = 1, \dots, d$, define the boundaries of C (note we may have $l_i = -\infty$ or $u_i = +\infty$ for some i). The values of $h_0(x)$ and $h_1(x)$ are simply the minimum distance from x to the treatment and non-treatment hypercubes, respectively:

$$h_0(x) = \min_{1 \leq k \leq K} \text{dist}(x, C_k), \quad h_1(x) = \min_{1 \leq l \leq L} \text{dist}(x, \tilde{C}_l).$$

The functions h_0 and h_1 are unaffected by whether the decision trees are defined using weak or strict inequalities (equivalently, open or closed hypercubes). \square

4.2 Empirical Strategies

The main challenge in the empirical implementation of externally valid treatment rules is to estimate the expectation

$$\mathbb{E}_P [\min \{Y_0 + \eta h_0(X), Y_1 + \eta h_1(X)\}] \tag{11}$$

appearing in the robust criterion. The presence of the minimum as an argument of the expectation means that the expectation depends joint distribution P of (Y_0, Y_1, X) , which is not identified—even within the context of a RCT—without further assumptions.

In this section, we give some strategies for implementing externally valid treatment rules under different identifying assumptions. We emphasize that we only require these dependence assumptions to hold for the experimental population P , and not for any target population $Q \in \mathcal{Q}$. The examples we give are also non-exhaustive: any other set of credible assumptions that permits identification and estimation of (11) may be used.

In particular, any assumptions identifying the distribution of individual causal effects may be used; we refer the reader to the handbook chapters [Abbring and Heckman \(2007\)](#) and [Heckman and Vytlacil \(2007\)](#) for a comprehensive treatment.

For each example, we derive an empirical robust welfare criterion RW_n and discuss computation. Externally valid treatment rules can be estimated by maximizing RW_n over \mathcal{T} . In what follows, we say that $\hat{\tau}$ *approximately maximizes* RW_n over \mathcal{T} if

$$RW_n(\hat{\tau}) \geq \sup_{\tau \in \mathcal{T}} RW_n(\tau) - \eta_n, \quad \eta_n \geq 0, \quad \eta_n \rightarrow_p 0,$$

where η_n should be interpreted as an asymptotically negligible optimization error. We say an estimated rule $\hat{\tau}$ is *asymptotically optimal* if its regret under the robust criterion is asymptotically zero in probability:

$$\sup_{\tau \in \mathcal{T}} RW(\tau) - RW(\hat{\tau}) \rightarrow_p 0.$$

4.2.1 Constant Treatment Effects

A simple approach to identifying P is to assume that treatment effects are constant, in the sense that conditional on $X = x$ we have

$$Y_0 - Y_1 = \Delta(x) \tag{12}$$

for some function $\Delta : \mathcal{X} \rightarrow \mathbb{R}$. This assumption allows treatment effect heterogeneity across individuals with different covariates, but requires treatment effect homogeneity across individuals with the same covariates. For the empirical implementation in this case, note that the expectation (11) may be written

$$\mathbb{E}_P [Y_0 + \min \{ \eta h_0(X), \Delta(X) + \eta h_1(X) \}],$$

so we only need to identify $\Delta(x)$ on the support of P . This is possible under a variety of conditions.

For example, suppose the analyst observes data $(X_i, Y_i, D_i)_{i=1}^n$ where D_i is a binary treatment indicator and $Y_i = D_i Y_{1i} + (1 - D_i) Y_{0i}$. Then Δ is identified under the unconfoundedness (or selection on observables) condition

$$(Y_0, Y_1) \perp\!\!\!\perp D \mid X \tag{13}$$

and overlap condition

$$0 < \mathbb{E}[D|X = x] < 1 \quad \text{for all } x \in \text{supp}(P) \quad (14)$$

(Rosenbaum and Rubin, 1983; Imbens and Rubin, 2015). We may estimate Δ using standard nonparametric techniques or modern methods; see, e.g, Imai and Ratkovic (2013), Wager and Athey (2018), Nie and Wager (2020), and references therein.⁹

As an alternative to assuming unconfoundedness, suppose the analyst observes data $(X_i, Y_i, D_i, Z_i)_{i=1}^n$ where Z_i is an instrumental variable satisfying appropriate regularity conditions (see, e.g., Abadie (2003)). Then in view of (12), Δ is identified as the ratio

$$\Delta(x) = \frac{\text{Cov}(Y_i, Z_i|X_i = x)}{\text{Cov}(D_i, Z_i|X_i = x)}$$

and may be estimated using a variety of nonparametric instrumental variables methods.

In either case, given estimators $\hat{\Delta}$ of Δ and \bar{Y}_0 of $\mathbb{E}_P[Y_0]$ and a random sample X_1, \dots, X_n from P , we choose $\hat{\tau}$ by maximizing the empirical criterion

$$\text{RW}_n(\tau) = \max \left\{ \bar{Y}_0 + \sup_{\eta \geq 1} \frac{1}{n} \sum_{i=1}^n \min \left\{ \eta h_0(X_i), \hat{\Delta}(X_i) + \eta h_1(X_i) \right\} - \eta \varepsilon, \underline{Y} \right\}$$

with respect to $\tau \in \mathcal{T}$. For each τ , the optimization over η may be efficiently performed using linear programming:

$$\begin{aligned} & \sup_{\eta \geq 1} \frac{1}{n} \sum_{i=1}^n \min \left\{ \eta h_0(X_i), \hat{\Delta}(X_i) + \eta h_1(X_i) \right\} - \eta \varepsilon \\ &= \sup_{\eta, (t_i)_{i=1}^n} \frac{1}{n} \sum_{i=1}^n t_i - \eta \varepsilon \\ & \text{subject to } \eta \geq 1, \quad t_i \leq \eta h_0(X_i), \quad t_i \leq \hat{\Delta}(X_i) + \eta h_1(X_i), \quad i = 1, \dots, n. \end{aligned} \quad (15)$$

Sample code for solving this problem is presented in Appendix B.

Asymptotic optimality of $\hat{\tau}$ may be established under a variety of different regularity conditions. The following result provides one example, allowing for both experimental and observational data.

⁹Note by restricting \mathcal{T} we can allow for X in (13) and (14) to contain additional variables beyond those that the treatment rule τ depends upon, which may be important to ensure unconfoundedness is credibly satisfied in observational settings.

Proposition 4.2 *Suppose that the conditions of Proposition 4.1 hold, condition (12) holds, \mathcal{X} is bounded, $\frac{1}{n} \sum_{i=1}^n |\hat{\Delta}(X_i) - \Delta(X_i)| \rightarrow_p 0$, and $\bar{Y}_0 \rightarrow_p \mathbb{E}_P[Y_0]$. Then any rule $\hat{\tau}$ that approximately maximizes $\text{RW}_n(\tau)$ over \mathcal{T} is asymptotically optimal.*

Remark 4.2 Proposition 4.2 shows the regret $\sup_{\tau \in \mathcal{T}} \text{RW}(\tau) - \text{RW}(\hat{\tau})$ converges to zero in probability for any (and hence all) classes of treatment rules. The price to pay for this generality is the assumption of a bounded covariate space \mathcal{X} , which is used to restrict the bracketing number of the class $\{\min\{\eta h_0(x), \Delta(x) + \eta h_1(x)\} : \tau \in \mathcal{T}\}$. This assumption can be relaxed for particular classes of treatment rules \mathcal{T} for which the bracketing number (or related complexity measures) can be controlled by other methods. It is also possible to derive explicit convergence rates for the regret of $\hat{\tau}$ under additional conditions on \mathcal{T} and the estimators $\hat{\Delta}$ and \bar{Y}_0 .

4.2.2 Perfect Positive Dependence

A weaker identifying assumption that permits heterogeneity in treatment effects among individuals with the same covariates is perfect positive dependence (or rank invariance); see Heckman, Smith, and Clements (1997) and Chernozhukov and Hansen (2005), among others.

Let $F_0(\cdot|x)$ and $F_1(\cdot|x)$ denote the conditional distributions of Y_0 and Y_1 given $X = x$ under P . Perfect positive dependence assumes that the ranking of individuals' potential outcomes in both distributions is the same. If $F_0(\cdot|x)$ and $F_1(\cdot|x)$ are continuous and increasing for each x ,¹⁰ then $\phi_{0,x}(\cdot) = F_0^{-1}(F_1(\cdot|x)|x)$ maps the treated outcome of an individual with covariates $X = x$ into their untreated outcome. Similarly, $\phi_{1,x}(\cdot) = F_1^{-1}(F_0(\cdot|x)|x)$ maps untreated potential outcomes into treated outcomes.

For the empirical implementation in this case, suppose the analyst observes a random sample $(X_i, Y_i, D_i)_{i=1}^n$ with D_i a treatment dummy and $Y_i = D_i Y_{1i} + (1 - D_i) Y_{0i}$, and the conventional unconfoundedness (13) and overlap (14) conditions hold. The conditional CDFs $F_0(\cdot|x)$ and $F_1(\cdot|x)$ are then nonparametrically identified as the conditional CDFs of Y given $D = d, X = x$ for $d = 0, 1$, respectively. These can be estimated using standard nonparametric methods, from which we may construct $\hat{\phi}_{0,x}$ and $\hat{\phi}_{1,x}$. Alternatively, in the case of endogenous treatments, we refer the reader to Vuong and Xu (2017) and

¹⁰Following the literature on nonseparable models (see, e.g., Matzkin (2003), Chesher (2003), and Chernozhukov and Hansen (2005)), suppose that $Y_d = m(d, X, U)$ where $m(d, x, \cdot)$ is strictly increasing and the conditional distribution of $U|X = x$ is absolutely continuous for all $x \in \mathcal{X}$ and $d \in \{0, 1\}$. Then perfect positive dependence holds and $F_0(\cdot|x)$ and $F_1(\cdot|x)$ are strictly monotone for all $x \in \mathcal{X}$.

Feng, Vuong, and Xu (2020) for identification and estimation of $\phi_{0,x}$ and $\phi_{1,x}$ using instrumental variables.

Given an estimator \bar{Y}_0 of $\mathbb{E}_P[Y_0]$ and estimators $\hat{\phi}_{1,x}$ and $\hat{\phi}_{0,x}$ of $\phi_{1,x}$ and $\phi_{0,x}$, we choose $\hat{\tau}$ by maximizing the empirical criterion

$$\text{RW}_n(\tau) = \max \left\{ \bar{Y}_0 + \sup_{\eta \geq 1} \frac{1}{n} \sum_{i=1}^n \min \left\{ \eta h_0(X_i), \hat{\Delta}_i + \eta h_1(X_i) \right\} - \eta \varepsilon, \underline{Y} \right\}$$

with respect to $\tau \in \mathcal{T}$, where

$$\hat{\Delta}_i = D_i \left(Y_i - \hat{\phi}_{0,X_i}(Y_i) \right) + (1 - D_i) \left(\hat{\phi}_{1,X_i}(Y_i) - Y_i \right)$$

is an estimate of the individual treatment effect $\Delta_i := Y_{1i} - Y_{0i}$. The optimization over η can be solved by linear programming as in display (15), replacing $\hat{\Delta}(X_i)$ with $\hat{\Delta}_i$.

Proposition 4.3 *Suppose that the conditions of Proposition 4.1 hold, perfect positive dependence holds, \mathcal{X} is bounded, $\frac{1}{n} \sum_{i=1}^n |\hat{\Delta}_i - \Delta_i| \rightarrow_p 0$, and $\bar{Y}_0 \rightarrow_p \mathbb{E}_P[Y_0]$. Then any rule $\hat{\tau}$ that approximately maximizes RW_n over \mathcal{T} is asymptotically optimal.*

As before, it is possible to relax the assumption of bounded \mathcal{X} for particular classes of rule \mathcal{T} . It is also possible to derive explicit convergence rates for the regret of $\hat{\tau}$ under additional conditions on \mathcal{T} and the estimators $\hat{\Delta}_i$ and \bar{Y}_0 .

4.2.3 Conditional Independence

An alternative identifying assumption is to posit the existence of a random variable C such that all dependence between Y_0 and Y_1 given X comes through C :

$$Y_0 \perp\!\!\!\perp Y_1 \mid X, C \tag{16}$$

(Abbring and Heckman, 2007, Section 2.5.1), and that a suitably modified version of unconfoundedness holds:

$$(Y_0, Y_1) \perp\!\!\!\perp D \mid (X, C). \tag{17}$$

The variables C may correspond to those in X , but can be distinct. For example, Y_0 and Y_1 may be conditionally independent given X and a latent common factor F , and C could be a perfect proxy for F constructed from X and other observables. A special case are group fixed effects where potential outcomes (and possibly the assignment mechanism)

has a group-specific component α_C . Using a linear model for simplicity, we might have

$$\begin{aligned} Y_0 &= X'\beta_0 + \lambda_0\alpha_C + u_0, \\ Y_1 &= X'\beta_1 + \lambda_1\alpha_C + u_1, \end{aligned}$$

with C denoting the individual's group membership (assumed known) and where u_0 and u_1 are conditionally independent given (X, C) .

Suppose the analyst observes a random sample $(X_i, Y_i, D_i, C_i)_{i=1}^n$ for which conditions (16) and (17) hold and a suitably modified version of the overlap condition (14) holds. The conditional CDFs $F_0(\cdot|x, c)$ and $F_1(\cdot|x, c)$ are then nonparametrically identified as the conditional CDFs of Y given $D = d, X = x, C = c$ for $d = 0, 1$, respectively. Assumption (16) then permits identification of the expectation (11) for all η, h_0 , and h_1 from $F_0(\cdot|x, c)$ and $F_1(\cdot|x, c)$. To see this, note the conditional CDF $G(\cdot|x, c)$ of $Z := \min\{Y_0 + \eta h_0(X), Y_1 + \eta h_1(X)\}$ given $X = x$ and $C = c$ is

$$G_{\eta, h_0, h_1}(z|x, c) = 1 - (1 - F_0(z - \eta h_0(x)|x, c))(1 - F_1(z - \eta h_1(x)|x, c)),$$

where F_0 and F_1 are the conditional CDFs of Y_0 and Y_1 given $X = x$ and $C = c$. It follows by iterated expectations that

$$\mathbb{E}_P[\min\{Y_0 + \eta h_0(X), Y_1 + \eta h_1(X)\}] = \int \int z dG_{\eta, h_0, h_1}(z|x, c) dP_{X, C}(x, c),$$

where $P_{X, C}$ is the distribution of (X, C) in the experimental population.

Given estimators $\hat{F}_0(y|x, c)$ and $\hat{F}_1(y|x, c)$, we estimate G_{η, h_0, h_1} using

$$\hat{G}_{\eta, h_0, h_1}(z|x, c) = 1 - (1 - \hat{F}_0(z - \eta h_0(x)|x, c))(1 - \hat{F}_1(z - \eta h_1(x)|x, c)).$$

A nice feature of \hat{G} is that it is uniformly consistent in (z, η, h_0, h_1) whenever $\hat{F}_0(y|x, c)$ and $\hat{F}_1(y|x, c)$ are uniformly consistent in y . We choose $\hat{\tau}$ by maximizing the empirical criterion

$$\text{RW}_n(\tau) = \max \left\{ \sup_{\eta \geq 1} \frac{1}{n} \sum_{i=1}^n \int z d\hat{G}_{\eta, h_0, h_1}(z|X_i, C_i) - \eta \varepsilon, \underline{Y} \right\}$$

with respect to $\tau \in \mathcal{T}$.

Asymptotic optimality may be derived under a variety of regularity conditions. The following result provides one example. Let $\bar{F}_d(y) = \frac{1}{n} \sum_{i=1}^n \hat{F}_d(y|X_i, C_i)$ for $d = 0, 1$, let a, A be finite positive constants, and let wpa1 denote with probability approaching one.

Proposition 4.4 *Suppose that the conditions of Proposition 4.1 hold, the conditional independence condition (16) holds, \mathcal{X} is bounded, and $\int |y|^{1+a} d\bar{F}_d(y) \leq A$ wpa1 and $\frac{1}{n} \sum_{i=1}^n \sup_y |\hat{F}_d(y|X_i, C_i) - F_d(y|X_i, C_i)| \rightarrow_p 0$ for $d = 0, 1$. Then any rule $\hat{\tau}$ that approximately maximizes RW_n over \mathcal{T} is asymptotically optimal.*

4.3 Criterion Based on a Least-Favorable Coupling

In this section, we take a conservative approach in the spirit of Imbens and Menzel (2021) and consider a criterion based on coupling of the conditional distribution of Y_0 and Y_1 given X that is *least-favorable* from the perspective of the robust welfare criterion. Under this coupling, the resulting distribution P^* of (X, Y_0, Y_1) induces the same marginal distributions for (X, Y_0) and (X, Y_1) as under P , but the robust welfare criterion is (weakly) smaller under P^* than under any other distribution of (X, Y_0, Y_1) with these marginals. Rules maximizing this criterion have welfare guarantees with respect to both adversarial shifts in the joint distribution of potential outcomes and covariates and adversarial couplings of the conditional distribution of Y_0 and Y_1 given X in the experimental population. We propose an empirical version of this criterion that can be estimated with experimental or observational data and can be computed by linear programming.

4.3.1 Least-Favorable Coupling and Criterion

Our starting point is to use additive separability to rewrite the expectation (11) as

$$\mathbb{E}_P [Y_0] + \mathbb{E}_P [\min\{\eta h_0(X), \Delta + \eta h_1(X)\}] .$$

We will use stochastic dominance relations on the distribution of $\Delta = Y_1 - Y_0$ given X to deduce a (sharp) lower bound on (11). Let P^* be a joint distribution for (X, Y_0, Y_1) such that the marginal distributions of (X, Y_0) and (X, Y_1) are the same as under P , and the conditional distribution of (Y_0, Y_1) given X under P^* satisfies perfect negative dependence:

$$Y_1 = F_1^{-1}(1 - F_0(Y_0|X)|X), \quad Y_0 = F_0^{-1}(1 - F_1(Y_1|X)|X),$$

where we assume as in Section 4.2.2 that the conditional distributions $F_d(\cdot|x)$ of Y_d given $X = x$ are continuous for each x and $d = 0, 1$. By Lemma 2.2 of Fan and Park (2010), the conditional distribution of Δ given X under P (weakly) second-order stochastically

dominates the corresponding conditional distribution under P^* . It follows by concavity of $u \mapsto \min\{a, u + b\}$ that

$$\mathbb{E}_P [\min\{\eta h_0(X), \Delta + \eta h_1(X)\}] \geq \mathbb{E}_{P^*} [\min\{\eta h_0(X), \Delta + \eta h_1(X)\}].$$

Hence, P^* minimizes $\text{RW}(\tau)$ among all joint distributions of (X, Y_0, Y_1) with the same marginals for (X, Y_0) and (X, Y_1) as P .¹¹

Define the robust welfare criterion

$$\text{RW}^*(\tau) = \max \left\{ \sup_{\eta \geq 1} \mathbb{E}_{P^*} [\min\{Y_0 + \eta h_0(X), Y_1 + \eta h_1(X)\}] - \eta \varepsilon, \underline{Y} \right\}. \quad (18)$$

Evidently $\text{RW}^*(\tau)$ provides a lower bound for the robust welfare criterion $\text{RW}(\tau)$ from Proposition 4.1. The lower bound is sharp because $\text{RW}(\tau) = \text{RW}^*(\tau)$ holds when the conditional distributions of Y_0 and Y_1 given $X = x$ are perfectly negatively dependent for P -almost every x .

4.3.2 Empirical Implementation

Under perfect negative dependence, the functions $\phi_{1,x}^*(x) = F_1^{-1}(1 - F_0(\cdot|x)|x)$ and $\phi_{0,x}^*(x) = F_0^{-1}(1 - F_1(\cdot|x)|x)$ map the treated outcome of an individual with covariates $X = x$ into their untreated outcome and vice versa. These functions may be estimated by standard nonparametric conditional density and quantile estimation methods if the analyst observes $(X_i, Y_i, D_i)_{i=1}^n$ and the unconfoundedness (13) and overlap (14) conditions hold. In the case of endogenous treatments, one can estimate these functions by suitably modifying the approach of Feng et al. (2020) to feature perfect negative dependence (as opposed to perfect positive dependence).

Given estimates $\hat{\phi}_{1,x}^*$ and $\hat{\phi}_{0,x}^*$ of $\phi_{1,x}^*$ and $\phi_{0,x}^*$, one may construct the estimate

$$\hat{\Delta}_i^* = D_i \left(Y_i - \hat{\phi}_{0,X_i}^*(Y_i) \right) + (1 - D_i) \left(\hat{\phi}_{1,X_i}^*(Y_i) - Y_i \right)$$

of $\Delta_i^* := D_i(Y_i - \phi_{0,X_i}^*(Y_i)) + (1 - D_i)(\phi_{1,X_i}^*(Y_i) - Y_i)$, which represents the treatment effect for individual i under perfect negative dependence. Given an estimator \bar{Y}_0 of $\mathbb{E}_P[Y_0]$, one

¹¹Similarly, the case of perfect positive dependence produces the largest $\text{RW}(\tau)$.

may then choose $\hat{\tau}$ by maximizing

$$\text{RW}_n^*(\tau) = \max \left\{ \bar{Y}_0 + \sup_{\eta \geq 1} \frac{1}{n} \sum_{i=1}^n \min \left\{ \eta h_0(X_i), \hat{\Delta}_i^* + \eta h_1(X_i) \right\} - \eta \varepsilon, \underline{Y} \right\}$$

with respect to $\tau \in \mathcal{T}$. The optimization over η can be solved by linear programming as in display (15), replacing $\hat{\Delta}(X_i)$ with $\hat{\Delta}_i^*$.

Analogously to the previous section, we refer to an empirical rule $\hat{\tau}$ as *asymptotically optimal* if $\sup_{\tau \in \mathcal{T}} \text{RW}^*(\tau) - \text{RW}^*(\hat{\tau}) \rightarrow_p 0$. Asymptotic optimality of empirical rules $\hat{\tau}$ that maximize RW_n^* may be established under a variety of regularity conditions. To this end, the following proposition again assumes \mathcal{X} is bounded.

Proposition 4.5 *Suppose that the conditions of Proposition 4.1 hold, \mathcal{X} is bounded, $\frac{1}{n} \sum_{i=1}^n |\hat{\Delta}_i^* - \Delta_i^*| \rightarrow_p 0$, and $\bar{Y}_0 \rightarrow_p \mathbb{E}_P[Y_0]$. Then any rule $\hat{\tau}$ that approximately maximizes RW_n^* over \mathcal{T} is asymptotically optimal.*

Remark 4.3 As with Propositions 4.2 and 4.3, asymptotic optimality is established for any (and hence all) classes of treatment rules. Boundedness of \mathcal{X} can be relaxed for particular classes of treatment rules \mathcal{T} . It is also possible to derive explicit convergence rates for particular classes \mathcal{T} .

5 Empirical Illustration

We conclude the paper with an empirical illustration using experimental data from the Job Training Partnership Act (JTPA) study (Bloom et al., 1997), as in Kitagawa and Tetenov (2018), Mbakop and Tabord-Meehan (2021), and several other recent works. We refer to these papers for a detailed description of the study and data.

The sample consists of 9,223 individuals. Each individual was randomly assigned to a job training program with probability 2/3. The outcome variable Y measure individual earnings in the 30-month period following assignment. As covariates we take individuals' years of education and pre-program annual earnings.¹²

We focus on the class of threshold rules that assign treatment depending on whether education and pre-program earnings are above or below particular cut-offs. The empirical welfare-maximizing (EWM) rule is to treat if education does not exceed 15 years and pre-program earnings do not exceed \$19,670 (Kitagawa and Tetenov, 2018).

¹²When computing the robust criterion, we rescale the education variable by the regression coefficient of pre-program earnings on education so that units are comparable across variables.

To examine the robustness of this rule to shifts in potential outcomes and covariates, we first proceed as in Section 4.2.1 and assume treatment effects are only heterogeneous in terms of education and pre-program earnings. For simplicity, we assume $\Delta(x)$ is linear in x and estimate it by regression. Figure 1 plots the robust criterion for the EWM rule and rules for 10 other randomly selected parameterizations of threshold rules. We plot over a range of neighborhood sizes from 0 to 2000. Note, for reference, that the estimated ATE is \$1,057, so the neighborhood size of 2000 includes populations with an ATE between roughly -\$943 and \$3,057. Similarly, a neighborhood size of 1000 includes populations with an ATE between roughly \$57 and \$2,057.

In this case, the EWM rule performs the best relative to the 10 other random parameterizations. Two of these are highlighted in Figure 1. The rule highlighted in orange corresponds to treating if pre-program earnings \leq \$21,160 and education \geq 8 years. Under this rule, only the small fraction of individuals with highest pre-program earnings or lowest education (7 years) are not treated. This rule is less robust than the EWM rule, with the gap between the worst-case welfare of the rules widening to \$155 once the neighborhood size exceeds 150 (or roughly 15% of the experimental-population ATE). The rule highlighted in blue shows even less robustness. This rule assigns treatment if pre-program earnings are \leq \$5,000 and education \leq 12 years. There is a large mass of individuals with exactly 12 years education and for whom pre-program earnings are in the ballpark of \$5,000. For this rule, adversarial covariate shifts push those at the boundary who would benefit from treatment into non-treatment regions which, when coupled with adversarial shifts in potential outcomes, leads to a large decline in welfare.

We next consider the case of perfect ranking, as in Section 4.2.2. We estimate the conditional CDFs $F_0(\cdot|x)$ and $F_1(\cdot|x)$ nonparametrically using kernels, then use these to estimate $\hat{\phi}_{0,X_i}$ and $\hat{\phi}_{1,X_i}$. The median individual treatment effect estimate is \$1,089, which is close to the ATE of \$1,057. The robust criterion is plotted in Figure 2. In this case the EWM rule performs near-optimally relative to the other rules, but is slightly dominated by the orange rule for $\varepsilon \geq 400$. The rule highlighted in blue is again very fragile, with performance deteriorating rapidly over small neighborhoods.

Overall, our findings show that the EWM rule is fairly robust to shifts in both potential outcomes and covariates, at least relative to other rules in this class.

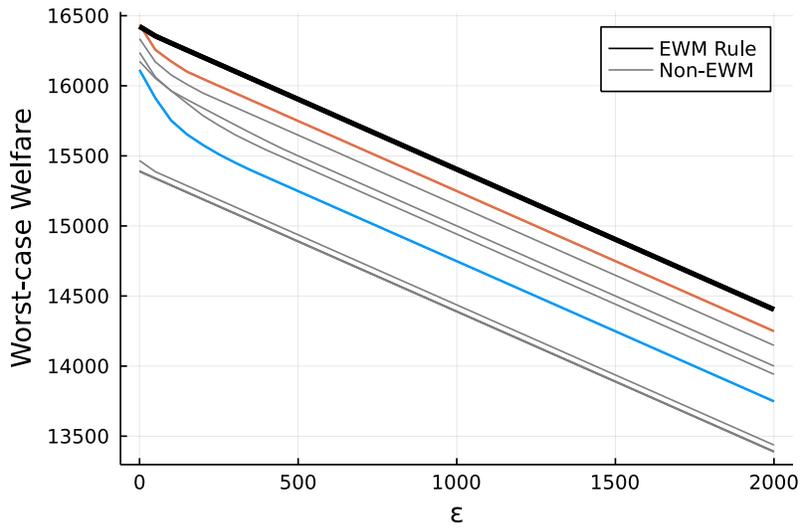


Figure 1: Worst-case welfare assuming treatment effects are linear in education and pre-program earnings. *Note:* colored lines highlight two non-EWM rules.

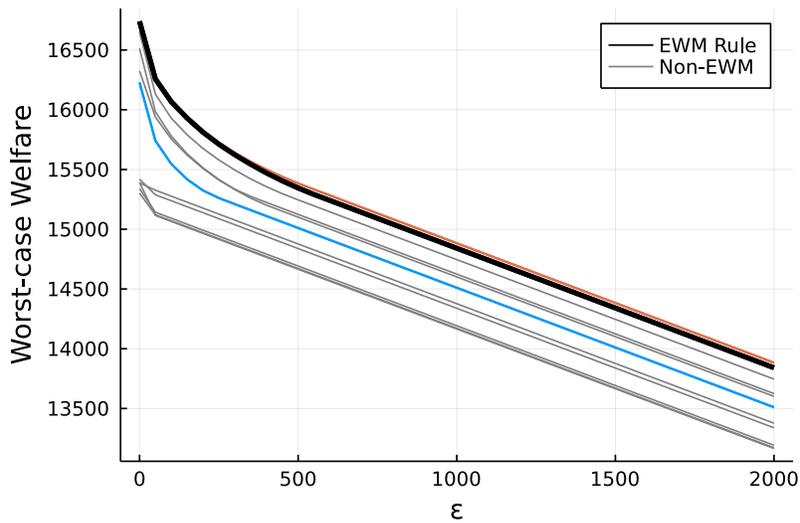


Figure 2: Worst-case welfare assuming perfect ranking conditional on education and pre-program earnings. *Note:* colored lines highlight two non-EWM rules.

A Extensions and Complements

A.1 Other Wasserstein Metrics

In this section, we show that the conclusion of Proposition 3.1 is not specific to our choice of Wasserstein metrics of order 1. We focus on metrics of order 2 as they are the most tractable; extensions to metrics of order p is straightforward, but with more complicated notation. Let $\mathcal{Q} = \{Q : W_2(P, Q) \leq \varepsilon\}$ where $W_2(P, Q)$ is the Wasserstein metric of order 2 induced by

$$d((x, y_0, y_1), (\tilde{x}, \tilde{y}_0, \tilde{y}_1)) = (|y_0 - \tilde{y}_0|^2 + |y_1 - \tilde{y}_1|^2 + \infty \times \mathbb{I}[x \neq \tilde{x}])^{1/2}. \quad (19)$$

The following result gives a tractable reformulation of the robust welfare criterion (2) when $\mathcal{Y} = \mathbb{R}$.

Proposition A.1 *Suppose that \mathcal{Q} is defined using the Wasserstein metric $W_2(P, Q)$ induced by (19) and that Y_0 and Y_1 have finite second moments under P . Then for any treatment rule τ , the worst-case social welfare over \mathcal{Q} is*

$$\inf_{Q \in \mathcal{Q}} \mathbb{E}_Q [Y_1 \tau(X) + Y_0(1 - \tau(X))] = \mathbb{E}_P [Y_1 \tau(X) + Y_0(1 - \tau(X))] - \varepsilon.$$

The robust welfare criterion in Proposition A.1 is the same as the criterion derived in Proposition 3.1 when $\underline{Y} = -\infty$. As such, Propositions 3.1 and A.1 provide a stronger sense in which rules with good guarantees under criterion (1) also have good external validity guarantees with respect to shifts in potential outcomes.

A.2 Alternative Welfare Criteria

We now extend Proposition 3.1 to the welfare criterion

$$W'(\tau) = \mathbb{E}_P [Y_1 \tau(X) - Y_0]$$

considered by Athey and Wager (2021). This criterion measures the additional utility incurred by treatment rule τ relative to treating no one. We define the robust version of this criterion as

$$\text{RW}'(\tau) = \inf_{Q \in \mathcal{Q}} \mathbb{E}_Q [Y_1 \tau(X) - Y_0].$$

Proposition A.2 *Suppose that \mathcal{Q} is defined using the Wasserstein metric $W(P, Q)$ induced by (6). Then for any treatment rule τ*

$$RW'(\tau) = \max \left\{ \mathbb{E}_P [Y_1 \tau(X) - Y_0] - \varepsilon, \underline{Y} \mathbb{E}_P [\tau(X)] - \bar{Y} \right\}.$$

If potential outcomes are unbounded or if $\underline{Y} = 0$, as is the case for binary and, more generally, non-negative outcomes, then Proposition A.2 shows that rules that are optimal under $W'(\tau)$ are also optimal under $RW'(\tau)$. The other implications discussed in Remark 3.1 also carry over in this case. It follows that empirical methods for learning rules that have good regret guarantees under $W'(\tau)$, such as those of [Athey and Wager \(2021\)](#), also have good regret guarantees under the robust criterion $RW'(\tau)$.

B Code for Implementation

This appendix presents example code to compute the empirical robust welfare criterion RW_n from Sections 4.2.1 and 4.2.2. Computing the robust welfare requires solving the linear program (15), where $\hat{\Delta}(X_i)$ is replaced with $\hat{\Delta}_i$ in the case of perfect ranking. The code is written in Julia and uses the JuMP and HiGHS packages.

The inputs are the neighborhood size `epsilon` and vectors `H0`, `H1`, and `Delta`, each of length n . The vectors `H0` and `H1` collect $h_0(X_i)$ and $h_1(X_i)$ across each of the n observations. The vector `Delta` collects $\hat{\Delta}(X_i)$ (for Section 4.2.1) or $\hat{\Delta}_i$ (for Section 4.2.2) across each of the n observations.

```
model = Model(HiGHS.Optimizer)
@variable(model, eta >= 1.0)
@variable(model, t[1:length(H0)])
for i in 1:length(H0)
    @constraint(model, t[i] <= eta * H0[i])
    @constraint(model, t[i] <= Delta[i] + eta * H1[i])
end
@objective(model, Max, sum(t) / length(H0) - epsilon * eta)
optimize!(model)
```

Solving the linear program in the empirical application (sample size $n = 9,223$) took between 0.2 and 0.6 seconds on a 2.7GHz MacBook Pro with 16GB memory, depending on the values of `epsilon`, `H0`, `H1`, and `Delta`.

C Proofs

Proof of Proposition 2.1. As X does not appear in the objective and because the metric (3) is additively separable, it is without loss of generality to drop covariates by setting $b = 0$ and letting P and Q be distributions over (Y_0, Y_1) . Consider first the minimization problem. The Lagrangian is

$$L = \inf_Q \sup_{\eta \geq 0} \mathbb{E}_Q[Y_1 - Y_0] + \eta(W(P, Q) - \varepsilon).$$

The Lagrangian dual is

$$\begin{aligned} L^* &= \sup_{\eta \geq 0} \inf_Q (\mathbb{E}_Q[Y_1 - Y_0] + \eta(W(P, Q) - \varepsilon)) \\ &= \sup_{\eta \geq 0} \inf_Q \inf_{\pi \in \Pi(P, Q)} \mathbb{E}_\pi \left[\tilde{Y}_1 - \tilde{Y}_0 + \eta \left(|Y_0 - \tilde{Y}_0| + |Y_1 - \tilde{Y}_1| - \varepsilon \right) \right]. \end{aligned}$$

Note the iterated infimum over Q and couplings $\pi \in \Pi(P, Q)$ is equivalent to infimizing over all joint distributions for $(Y_0, Y_1, \tilde{Y}_0, \tilde{Y}_1)$ with marginal P for (Y_0, Y_1) . Hence,

$$L^* = \sup_{\eta \geq 0} \inf_{\{F_{(\tilde{Y}_0, \tilde{Y}_1)|(Y_0, Y_1)}\}} \mathbb{E}_P \left[\mathbb{E}_{F_{(\tilde{Y}_0, \tilde{Y}_1)|(Y_0, Y_1)}} \left[\tilde{Y}_1 - \tilde{Y}_0 + \eta \left(|Y_0 - \tilde{Y}_0| + |Y_1 - \tilde{Y}_1| - \varepsilon \right) \middle| Y_0, Y_1 \right] \right],$$

where the infimum is over all conditional distribution for $(\tilde{Y}_0, \tilde{Y}_1)$ given (Y_0, Y_1) , for each (Y_0, Y_1) in the support of P . As it is without loss of generality to optimize over point masses, we obtain

$$L^* = \sup_{\eta \geq 0} \mathbb{E}_P \left[\inf_{\tilde{y}_0, \tilde{y}_1} (\tilde{y}_1 - \tilde{y}_0 + \eta (|Y_0 - \tilde{y}_0| + |Y_1 - \tilde{y}_1| - \varepsilon)) \right].$$

First suppose $\eta \in [0, 1)$. As $y \mapsto y + \eta|Y - y|$ is minimized by taking $y = \underline{Y}$ if $\underline{Y} > -\infty$ or $y \rightarrow -\infty$ otherwise, while $y \mapsto y - \eta|Y - y|$ is maximized by taking $y = \bar{Y}$, or $y \rightarrow +\infty$ otherwise. We therefore have

$$\begin{aligned} &\mathbb{E}_P \left[\inf_{\tilde{y}_0, \tilde{y}_1} (\tilde{y}_1 - \tilde{y}_0 + \eta (|Y_0 - \tilde{y}_0| + |Y_1 - \tilde{y}_1| - \varepsilon)) \right] \\ &= \begin{cases} -\infty & \text{if } \underline{Y} = -\infty \text{ or } \bar{Y} = +\infty, \\ (1 - \eta) (\underline{Y} - \bar{Y}) + \eta \mathbb{E}_P[Y_1 - Y_0] - \eta \varepsilon & \text{otherwise.} \end{cases} \end{aligned}$$

If $\eta \geq 1$, the function $y \mapsto y + \eta|Y - y|$ is minimized by taking $y = Y$ while $y \mapsto y - \eta|Y - y|$

is maximized also by taking $y = Y$, so

$$\mathbb{E}_P \left[\inf_{\tilde{y}_0, \tilde{y}_1} (\tilde{y}_1 - \tilde{y}_0 + \eta (|Y_0 - \tilde{y}_0| + |Y_1 - \tilde{y}_1| - \varepsilon)) \right] = \mathbb{E}_P[Y_1 - Y_0] - \eta\varepsilon.$$

Combining the preceding two displays and maximizing with respect to η , we obtain

$$L^* = \max \{ \mathbb{E}_P[Y_1 - Y_0] - \varepsilon, \underline{Y} - \bar{Y} \}.$$

We now show $L = L^*$. Note (Y_0, Y_1) has finite first moment under P by virtue of our maintained assumptions on the potential outcomes. It follows by Theorem 1 of [Yue, Kuhn, and Wiesemann \(2021\)](#) that \mathcal{Q} is weakly compact. Moreover, $Q \mapsto W(P, Q)$ and $Q \mapsto \mathbb{E}_Q[Y_1 - Y_0]$ are weakly lower-semicontinuous (see the proofs of Theorems 1 and 3 of [Yue et al. \(2021\)](#), respectively). It follows by Sion's minimax theorem that we may interchange the inf and sup in L , hence $L = L^*$.

The proof of the upper bound follows by symmetric arguments. ■

Proof of Proposition 3.1. The Lagrangian is

$$L = \inf_Q \sup_{\eta \geq 0} (\mathbb{E}_Q[Y_1\tau(X) + Y_0(1 - \tau(X))] + \eta(W(P, Q) - \varepsilon)).$$

The Lagrangian dual is

$$\begin{aligned} L^* &= \sup_{\eta \geq 0} \inf_Q (\mathbb{E}_Q[Y_1\tau(X) + Y_0(1 - \tau(X))] + \eta(W(P, Q) - \varepsilon)) \\ &= \sup_{\eta \geq 0} \inf_Q \inf_{\pi \in \Pi(P, Q)} \mathbb{E}_\pi \left[\tilde{Y}_1\tau(\tilde{X}) + \tilde{Y}_0(1 - \tau(\tilde{X})) \right. \\ &\quad \left. + \eta (|Y_0 - \tilde{Y}_0| + |Y_1 - \tilde{Y}_1| + \infty \times \mathbb{I}[X \neq \tilde{X}] - \varepsilon) \right]. \end{aligned}$$

As before, the iterated infimum over Q and couplings $\pi \in \Pi(P, Q)$ is equivalent to infimizing over all joint distributions for (Z, \tilde{Z}) with marginal P for Z . Note also that it suffices to consider distribution for which $X = \tilde{X}$ almost surely. We therefore have

$$\begin{aligned} L^* &= \sup_{\eta \geq 0} \inf_{\{F_{(\tilde{Y}_0, \tilde{Y}_1)|Z}\}} \mathbb{E}_P \left[\mathbb{E}_{F_{(\tilde{Y}_0, \tilde{Y}_1)|Z}} \left[\tilde{Y}_1\tau(X) + \tilde{Y}_0(1 - \tau(X)) \right. \right. \\ &\quad \left. \left. + \eta (|Y_0 - \tilde{Y}_0| + |Y_1 - \tilde{Y}_1| - \varepsilon) \middle| Z \right] \right], \end{aligned}$$

where the infimum is over all conditional distribution for $(\tilde{Y}_0, \tilde{Y}_1)$ given Z , for each Z in the support of P . As it is without loss of generality to optimize over point masses, we obtain

$$L^* = \sup_{\eta \geq 0} \mathbb{E}_P \left[\inf_{\tilde{y}_0, \tilde{y}_1} (\tilde{y}_1 \tau(X) + \tilde{y}_0 (1 - \tau(X)) + \eta (|Y_0 - \tilde{y}_0| + |Y_1 - \tilde{y}_1| - \varepsilon)) \right].$$

The remainder of the proof will differ depending on whether $\underline{Y} > -\infty$ or $\underline{Y} = -\infty$.

Case 1: $\underline{Y} > -\infty$. The inner infimization may be solved in closed form for each fixed $Z = (X, Y_0, Y_1)$. Suppose $\eta \in [0, 1)$. Then $y \mapsto y + \eta|Y - y|$ is minimized by setting $y = \underline{Y}$, with the minimizing value being $\underline{Y} + \eta(Y - \underline{Y})$. Therefore, if $\tau(X) = 1$, then the minimum is attained with $(\tilde{y}_0, \tilde{y}_1) = (Y_0, \underline{Y})$. Conversely, if $\tau(X) = 0$, then the minimum is attained with $(\tilde{y}_0, \tilde{y}_1) = (\underline{Y}, Y_1)$. For $\eta \in [0, 1)$, the dual objective therefore becomes

$$\mathbb{E}_P [\underline{Y} + \eta((Y_0 - \underline{Y}) + (Y_1 - Y_0)\tau(X) - \varepsilon)].$$

Maximizing this term with respect to $\eta \in [0, 1)$, we see that the optimal value is

$$\begin{cases} \mathbb{E}_P [Y_0 + (Y_1 - Y_0)\tau(X) - \varepsilon] & \text{if } \mathbb{E}_P [(Y_0 - \underline{Y}) + (Y_1 - Y_0)\tau(X)] \geq \varepsilon, \\ \underline{Y} & \text{otherwise.} \end{cases}$$

Now suppose that $\eta \geq 1$. In this case, $y + \eta|Y - y|$ is minimized by setting $y = Y$, so the dual objective is

$$\mathbb{E}_P [Y_0 + (Y_1 - Y_0)\tau(X) - \eta\varepsilon]$$

Combining these results, it follows that

$$L^* = \max \left\{ \mathbb{E}_P [Y_1 \tau(X) + Y_0 (1 - \tau(X))] - \varepsilon, \underline{Y} \right\}.$$

It remains to show that $L = L^*$. We cannot invoke the argument used in the proof of Proposition 2.1 as the metric d in (6) is not proper, which Yue et al. (2021) require. We therefore provide a constructive proof. By weak duality ($L^* \leq L$), it suffices to show $L \leq L^*$. Partition $\mathcal{X} = \mathcal{X}_0 \cup \mathcal{X}_1$ with $\mathcal{X}_0 = \tau^{-1}(\{0\})$ and $\mathcal{X}_1 = \tau^{-1}(\{1\})$. Let Q_0 denote the distribution of $T(Z)$ with $Z \sim P$, where

$$T(x, y_0, y_1) = \begin{cases} (x, \underline{Y}, y_1) & \text{if } x \in \mathcal{X}_0, \\ (x, y_0, \underline{Y}) & \text{if } x \in \mathcal{X}_1. \end{cases}$$

Then $\mathbb{E}_{Q_0}[Y_1\tau(X) + Y_0(1 - \tau(X))] = \underline{Y}$. Consider the coupling $(Z, T(Z)) \sim \pi$ for $Z \sim P$. Then under the metric (6), we have

$$\begin{aligned} W(P, Q_0) &\leq \mathbb{E}_\pi[d((X, Y_0, Y_1), (\tilde{X}, \tilde{Y}_0, \tilde{Y}_1))] \\ &= \mathbb{E}_P[(Y_0 - \underline{Y})\mathbb{I}[X \in \mathcal{X}_0] + (Y_1 - \underline{Y})\mathbb{I}[X \in \mathcal{X}_1]] \\ &= \mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \underline{Y}. \end{aligned} \tag{20}$$

It follows that whenever $\mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \varepsilon < \underline{Y}$, we have

$$\begin{aligned} L &\leq \sup_{\eta \geq 0} (\mathbb{E}_{Q_0}[Y_1\tau(X) + Y_0(1 - \tau(X))] + \eta(W(P, Q_0) - \varepsilon)) \\ &\leq \underline{Y} + \sup_{\eta \geq 0} \eta (\mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \underline{Y} - \varepsilon) = L^*, \end{aligned}$$

as required.

Now suppose that $\mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \varepsilon \geq \underline{Y}$. Let $Q_1 = wQ_0 + (1 - w)P$ be a mixture distribution with weight

$$w = \frac{\varepsilon}{\mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \underline{Y}} \tag{21}$$

on Q_0 . Then $\mathbb{E}_{Q_1}[Y_1\tau(X) + Y_0(1 - \tau(X))] = \mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \varepsilon$. Moreover, by convexity of W (see, e.g., Villani (2009, Theorem 4.8)),

$$W(P, Q_1) \leq wW(P, Q_0) + (1 - w)W(P, P) = wW(P, Q_0) \leq \varepsilon,$$

where the final equality follows from (20) and (21). Therefore,

$$\begin{aligned} L &\leq \sup_{\eta \geq 0} (\mathbb{E}_{Q_1}[Y_1\tau(X) + Y_0(1 - \tau(X))] + \eta(W(P, Q_1) - \varepsilon)) \\ &\leq \mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \varepsilon + \sup_{\eta \geq 0} \eta(W(P, Q_1) - \varepsilon) = L^*, \end{aligned}$$

as required.

Case 2: $\underline{Y} = -\infty$. Suppose $\eta \in [0, 1)$. Then $y + \eta|Y - y|$ is minimized by taking $y \rightarrow -\infty$, and the minimizing value is $-\infty$. Conversely, if $\eta \geq 1$ then $y + \eta|Y - y|$ is minimized by setting $y = Y$, so the dual objective is

$$\mathbb{E}_P[Y_0 + (Y_1 - Y_0)\tau(X) - \eta\varepsilon],$$

which is maximized over $\eta \geq 1$ at $\eta = 1$. The dual is therefore

$$L^* = \mathbb{E}_P [Y_1 \tau(X) + Y_0(1 - \tau(X))] - \varepsilon.$$

It remains to show that $L \leq L^*$. As $\underline{Y} = -\infty$, for each y we let $(y)_\varepsilon$ denote an element of \mathcal{Y} for which $y - C \leq (y)_\varepsilon < y - \varepsilon$ for some constant $C > 0$ (we can always choose such a C and $(y)_\varepsilon$ because \mathcal{Y} is assumed to be equispaced). Let Q_0 denote the distribution of $T(Z)$ with $Z \sim P$, where

$$T(x, y_0, y_1) = \begin{cases} (x, (y_0)_\varepsilon, y_1) & \text{if } x \in \mathcal{X}_0, \\ (x, y_0, (y_1)_\varepsilon) & \text{if } x \in \mathcal{X}_1. \end{cases}$$

Then

$$\begin{aligned} \mathbb{E}_{Q_0}[Y_1 \tau(X) + Y_0(1 - \tau(X))] &= \mathbb{E}_P[(Y_1)_\varepsilon \tau(X) + (Y_0)_\varepsilon(1 - \tau(X))] \\ &< \mathbb{E}_P[Y_1 \tau(X) + Y_0(1 - \tau(X))] - \varepsilon, \end{aligned}$$

where $\mathbb{E}_{Q_0}[Y_1 \tau(X) + Y_0(1 - \tau(X))] > \mathbb{E}_P[Y_1 \tau(X) + Y_0(1 - \tau(X))] - C$, ensuring the left-hand side expectation is finite. Moreover, with $(Z, T(Z)) \sim \pi$ for $Z \sim P$, we have

$$\begin{aligned} W(P, Q_0) &\leq \mathbb{E}_\pi[d((X, Y_0, Y_1), (\tilde{X}, \tilde{Y}_0, \tilde{Y}_1))] \\ &= \mathbb{E}_P[(Y_0 - (Y_0)_\varepsilon)\mathbb{I}[X \in \mathcal{X}_0] + (Y_1 - (Y_1)_\varepsilon)\mathbb{I}[X \in \mathcal{X}_1]] \\ &= \mathbb{E}_P[Y_1 \tau(X) + Y_0(1 - \tau(X))] - \mathbb{E}_{Q_0}[Y_1 \tau(X) + Y_0(1 - \tau(X))]. \end{aligned} \quad (22)$$

Let $Q_1 = wQ_0 + (1 - w)P$ be a mixture distribution with weight

$$w = \frac{\varepsilon}{\mathbb{E}_P[Y_1 \tau(X) + Y_0(1 - \tau(X))] - \mathbb{E}_{Q_0}[Y_1 \tau(X) + Y_0(1 - \tau(X))]} \quad (23)$$

on Q_0 . Then $\mathbb{E}_{Q_1}[Y_1 \tau(X) + Y_0(1 - \tau(X))] = \mathbb{E}_P[Y_1 \tau(X) + Y_0(1 - \tau(X))] - \varepsilon$. Moreover,

$$W(P, Q_1) \leq wW(P, Q_0) + (1 - w)W(P, P) = wW(P, Q_0) \leq \varepsilon,$$

by convexity of W and (22) and (23). Therefore,

$$\begin{aligned} L &\leq \sup_{\eta \geq 0} (\mathbb{E}_{Q_1}[Y_1 \tau(X) + Y_0(1 - \tau(X))] + \eta(W(P, Q_1) - \varepsilon)) \\ &\leq \mathbb{E}_P[Y_1 \tau(X) + Y_0(1 - \tau(X))] - \varepsilon + \sup_{\eta \geq 0} \eta(W(P, Q_1) - \varepsilon) = L^*, \end{aligned}$$

as required. ■

Proof of Proposition 4.1. We argue as in the proof of Proposition 3.1. The Lagrangian is

$$L = \inf_Q \sup_{\eta \geq 0} (\mathbb{E}_Q[Y_1 \tau(X) + Y_0(1 - \tau(X))] + \eta(W(P, Q) - \varepsilon))$$

and its dual is

$$L^* = \sup_{\eta \geq 0} \mathbb{E}_P \left[\inf_{(\tilde{x}, \tilde{y}_0, \tilde{y}_1)} \left(\tilde{y}_1 \tau(\tilde{x}) + \tilde{y}_0(1 - \tau(\tilde{x})) + \eta (|Y_0 - \tilde{y}_0| + |Y_1 - \tilde{y}_1| + \|X - \tilde{x}\| - \varepsilon) \right) \right].$$

Consider the inner infimization at any fixed $Z = (X, Y_0, Y_1)$. We first fix \tilde{x} and optimize with respect to $(\tilde{y}_0, \tilde{y}_1)$, then optimize with respect to \tilde{x} . There are two cases to consider.

Case 1: $\underline{Y} > -\infty$. Suppose $\eta \in [0, 1)$. Then $y + \eta|Y - y|$ is minimized by setting $y = \underline{Y}$ and the minimizing value is $\underline{Y} + \eta(Y - \underline{Y})$. If $\tau(\tilde{x}) = 1$, then the infimum (at fixed \tilde{x}) is attained with $(\tilde{y}_0, \tilde{y}_1) = (Y_0, \underline{Y})$. Conversely, if $\tau(\tilde{x}) = 0$ then the infimum is attained with $(\tilde{y}_0, \tilde{y}_1) = (\underline{Y}, Y_1)$. For $\eta \in [0, 1)$, the objective therefore becomes

$$(1 - \eta)\underline{Y} + \eta \mathbb{E}_P \left[Y_0 + \inf_{\tilde{x}} ((Y_1 - Y_0)\tau(\tilde{x}) + \|X - \tilde{x}\|) \right] - \eta\varepsilon.$$

Maximizing with respect to $\eta \in [0, 1)$ yields

$$\max \left\{ \mathbb{E}_P \left[Y_0 + \inf_{\tilde{x}} ((Y_1 - Y_0)\tau(\tilde{x}) + \|X - \tilde{x}\|) \right] - \varepsilon, \underline{Y} \right\}.$$

Now suppose that $\eta \geq 1$. In this case, $y + \eta|Y - y|$ is minimized by setting $y = Y$ so the dual objective reduces to

$$\mathbb{E}_P \left[Y_0 + \inf_{\tilde{x}} ((Y_1 - Y_0)\tau(\tilde{x}) + \eta (\|X - \tilde{x}\|)) \right] - \eta\varepsilon. \quad (24)$$

Case 2: $\underline{Y} = -\infty$. If $\eta \in [0, 1)$, then $y + \eta|Y - y|$ is minimized by taking $y \rightarrow -\infty$ and the minimizing value is $-\infty$. Conversely if $\eta \geq 1$ then $y + \eta|Y - y|$ is minimized by setting $y = Y$ so the dual objective again reduces to (24).

Combining these results, we obtain

$$L^* = \max \left\{ \sup_{\eta \geq 1} \mathbb{E}_P \left[Y_0 + \inf_{\tilde{x}} ((Y_1 - Y_0)\tau(\tilde{x}) + \eta \|X - \tilde{x}\|) \right] - \eta\varepsilon, \underline{Y} \right\}.$$

Note that we may split the infimization up into separate infimizations over $\{\tilde{x} : \tau(\tilde{x}) = 0\}$ and $\{\tilde{x} : \tau(\tilde{x}) = 1\}$, then take the minimum:

$$\begin{aligned} L^* &= \max \left\{ \sup_{\eta \geq 1} \mathbb{E}_P \left[\min \left\{ Y_0 + \inf_{\tilde{x}:\tau(\tilde{x})=0} \eta \|X - \tilde{x}\|, Y_1 + \inf_{\tilde{x}:\tau(\tilde{x})=1} \eta \|X - \tilde{x}\| \right\} \right] - \eta\varepsilon, \underline{Y} \right\} \\ &= \max \left\{ \sup_{\eta \geq 1} \mathbb{E}_P [\min \{Y_0 + \eta h_0(X), Y_1 + \eta h_1(X)\}] - \eta\varepsilon, \underline{Y} \right\}. \end{aligned}$$

To complete the proof, it remains to show $L = L^*$. Note Z has finite first moment under P by virtue of our maintained assumptions on the potential outcomes and the condition $\mathbb{E}_P[\|X\|] < \infty$ in the statement of the result. It follows by Theorem 1 of [Yue et al. \(2021\)](#) that \mathcal{Q} is weakly compact. Moreover, $Q \mapsto W(P, Q)$ and $Q \mapsto \mathbb{E}_Q[Y_1\tau(X) + Y_0(1 - \tau(X))]$ are weakly lower-semicontinuous (see the proofs of Theorems 1 and 3 of [Yue et al. \(2021\)](#), respectively). It follows by Sion's minimax theorem that we may interchange the inf and sup in L , hence $L = L^*$. ■

Proof of Proposition 4.2. It is sufficient to show $\sup_{\tau \in \mathcal{T}} |\text{RW}_n(\tau) - \text{RW}(\tau)| = o_p(1)$. Note that in view of (12), we have

$$\text{RW}(\tau) = \max \left\{ \mathbb{E}_P[Y_0] + \sup_{\eta \geq 1} \mathbb{E}_P [\min \{\eta h_0(X), \Delta(X) + \eta h_1(X)\}] - \eta\varepsilon, \underline{Y} \right\}.$$

By the functional form of RW_n and the condition $\bar{Y}_0 \rightarrow_p \mathbb{E}_P[Y_0]$, it suffices to show

$$\begin{aligned} \sup_{\tau \in \mathcal{T}} \left| \sup_{\eta \geq 1} \frac{1}{n} \sum_{i=1}^n \min \left\{ \eta h_0(X_i), \hat{\Delta}(X_i) + \eta h_1(X_i) \right\} - \eta\varepsilon \right. \\ \left. - \left(\sup_{\eta \geq 1} \mathbb{E}_P [\min \{\eta h_0(X), \Delta(X) + \eta h_1(X)\}] - \eta\varepsilon \right) \right| = o_p(1). \end{aligned}$$

To this end, first note if $\tau \equiv 1$ then $h_0(x) = +\infty$ and $h_1(x) = 0$ for all x . The inner term in the preceding display therefore reduces to

$$\left| \frac{1}{n} \sum_{i=1}^n \hat{\Delta}(X_i) - \mathbb{E}_P[\Delta(X)] \right| + o_p(1),$$

which is $o_p(1)$ by the condition $\frac{1}{n} \sum_{i=1}^n |\hat{\Delta}(X_i) - \Delta(X_i)| = o_p(1)$ and the LLN. Similarly, if $\tau \equiv 0$ then $h_0(x) = 0$ and $h_1(x) = +\infty$ for all x , so the inner term is zero. In what follows, it therefore suffices to treat \mathcal{T} as a class of rules that are not identically 0 or 1.

It is straightforward to deduce by way of contradiction that there is a finite positive constant C such that wpa1 the argsup of both problems is an element of $[1, C]$, uniformly in τ . Moreover, by the assumed consistency condition for $\hat{\Delta}$ and the fact that the max and min functions are Lipschitz, we have

$$\begin{aligned} & \sup_{\tau \in \mathcal{T}} \left| \sup_{\eta \geq 1} \frac{1}{n} \sum_{i=1}^n \min \left\{ \eta h_0(X_i), \hat{\Delta}(X_i) + \eta h_1(X_i) \right\} - \eta \varepsilon \right. \\ & \quad \left. - \left(\sup_{\eta \geq 1} \mathbb{E}_P [\min \{ \eta h_0(X), \Delta(X) + \eta h_1(X) \}] - \eta \varepsilon \right) \right| \\ & \leq \sup_{\tau \in \mathcal{T}, \eta \in [1, C]} \left| \frac{1}{n} \sum_{i=1}^n \min \{ \eta h_0(X_i), \Delta(X_i) + \eta h_1(X_i) \} \right. \\ & \quad \left. - \mathbb{E}_P [\min \{ \eta h_0(X), \Delta(X) + \eta h_1(X) \}] \right| + o_p(1). \end{aligned}$$

To complete the proof, we invoke a ULLN to control the right-hand side of the preceding display. By the triangle inequality, we may deduce $|h_0(x) - h_0(x')| \leq \|x - x'\|$ for all $x, x' \in \mathcal{X}$, and similarly for h_1 . Therefore, $\{\eta h_d : \tau \in \mathcal{T}, \eta \in [1, C]\}$, $d = 0, 1$, are subsets of the class of Lipschitz functions on \mathcal{X} with Lipschitz constant at most C . As \mathcal{X} is bounded, these classes have finite $L^1(P)$ bracketing numbers (see, e.g., Corollary 2.7.2 of [van der Vaart and Wellner \(1996\)](#)). Note that $|\min\{\eta h_0(x), \Delta(x) + \eta h_1(x)\}| \leq |\Delta(x)|$ holds uniformly in τ and η , where $\mathbb{E}_P[|\Delta(X)|] \leq \mathbb{E}_P[|Y_0| + |Y_1|] < \infty$. It follows that the class $\{\min\{\eta h_0, \Delta + \eta h_1\} : \tau \in \mathcal{T}, \eta \in [1, C]\}$ has finite $L^1(P)$ bracketing numbers. The desired ULLN therefore holds by the Glivenko–Cantelli theorem (see, e.g., Theorem 2.4.1 of [van der Vaart and Wellner \(1996\)](#)). ■

Proof of Proposition 4.3. Follows by identical arguments to Proposition 4.2, replacing $\Delta(X_i)$ with Δ_i and $\Delta(X)$ with $Y_1 - Y_0$. ■

Proof of Proposition 4.4. We proceed as in the proof of Propositions 4.2 and 4.3. To simplify notation, let $\hat{G}_{\eta, h_0, h_1}^{(i)}(z) = \hat{G}_{\eta, h_0, h_1}(z | X_i, C_i)$, and similarly for G_{η, h_0, h_1} , \hat{F}_0 , \hat{F}_1 ,

F_0 , and F_1 . By the functional form of RW_n , it suffices to show

$$\sup_{\tau \in \mathcal{T}} \left| \sup_{\eta \geq 1} \frac{1}{n} \sum_{i=1}^n \int z \, d\hat{G}_{\eta, h_0, h_1}^{(i)}(z) - \eta \varepsilon - \left(\sup_{\eta \geq 1} \int \int z \, dG_{\eta, h_0, h_1}(z|x, c) \, dP_{X, C}(x, c) - \eta \varepsilon \right) \right| = o_p(1).$$

We first show that there is a sufficiently large constant B such that for all $\tau \in \mathcal{T}$, the argsup of both problems is an element of $[1, B]$ wpa1. For the first problem suppose the assertion is false. Then for some $\tau \in \mathcal{T}$ and $\eta \geq B$, we have

$$\frac{1}{n} \sum_{i=1}^n \int z \, d(\hat{G}_{\eta, h_0, h_1}^{(i)}(z) - \hat{G}_{1, h_0, h_1}^{(i)}(z)) \geq (\eta - 1)\varepsilon \geq (B - 1)\varepsilon. \quad (25)$$

Fix $1 \leq i \leq n$ and suppose that $h_0(X_i) = 0$. Then

$$\begin{aligned} & \left| \int z \, d(\hat{G}_{\eta, h_0, h_1}^{(i)}(z) - \hat{G}_{1, h_0, h_1}^{(i)}(z)) \right| \\ &= \left| \int z \, d \left(\left(\hat{F}_1^{(i)}(z - \eta h_1(X_i)) - \hat{F}_1^{(i)}(z - h_1(X_i)) \right) \left(1 - \hat{F}_0^{(i)}(z) \right) \right) \right| \leq 2 \int |z| \, d\hat{F}_0^{(i)}(z). \end{aligned}$$

A symmetric argument applies when $h_1(X_i) = 0$. Hence, the left-hand side of (25) is bounded above by

$$2 \int |z| \, d\bar{F}_0(z) + 2 \int |z| \, d\bar{F}_1(z).$$

It follows from the integrability condition in the statement of the proposition that B can be chosen sufficiently large so that inequality (25) is violated wpa1. That the argsup of the second problem is an element of $[1, B]$ may be deduced similarly by contradiction.

It therefore suffices to show

$$\sup_{\tau \in \mathcal{T}, \eta \in [1, B]} \left| \frac{1}{n} \sum_{i=1}^n \int z \, d\hat{G}_{\eta, h_0, h_1}^{(i)}(z) - \int \int z \, dG_{\eta, h_0, h_1}(z|x, c) \, dP_{X, C}(x, c) \right| = o_p(1).$$

By similar arguments to the above, we may use the integrability condition in the statement of the proposition to deduce that for any $\delta > 0$ there exists a finite constant M such that wpa1 the inequalities

$$\frac{1}{n} \sum_{i=1}^n \int_{|z| > M} |z| \, d\hat{G}_{\eta, h_0, h_1}(z|X_i, C_i) \leq \delta, \quad \int \int_{|z| > M} |z| \, dG_{\eta, h_0, h_1}(z|x, c) \, dP_{X, C}(x, c) \leq \delta$$

hold uniformly in (η, h_0, h_1) . Now consider the remaining terms:

$$T_1 + T_2 := \left(\frac{1}{n} \sum_{i=1}^n \int_{-M}^M z \, d\hat{G}_{\eta, h_0, h_1}^{(i)}(z) - \frac{1}{n} \sum_{i=1}^n \int_{-M}^M z \, dG_{\eta, h_0, h_1}^{(i)}(z) \right) \\ + \left(\frac{1}{n} \sum_{i=1}^n \int_{-M}^M z \, dG_{\eta, h_0, h_1}^{(i)}(z) - \int \int_{-M}^M z \, dG_{\eta, h_0, h_1}(z|x, c) \, dP_{X, C}(x, c) \right).$$

For T_1 we may deduce

$$\left| \int_{-M}^M z \, d \left(\hat{G}_{\eta, h_0, h_1}^{(i)}(z) - G_{\eta, h_0, h_1}^{(i)}(z) \right) \right| \leq 8M^2 \max_{d=0,1} \sup_y \left| \hat{F}_d^{(i)}(y|X_i, C_i) - F_d^{(i)}(y) \right|.$$

It follows by the uniform consistency condition for \hat{F}_0 and \hat{F}_1 in the statement of the proposition that $\sup_{\tau \in \mathcal{T}, \eta \in [1, B]} |T_1| \rightarrow_p 0$. To establish the corresponding result for T_2 , it suffices to show that a ULLN holds for the class of functions

$$\{\mathbb{E}[a(X, Y_0, Y_1) \mathbb{I}[|a(X, Y_0, Y_1)| \leq M] | X = x, C = c] : a \in \mathcal{A}\},$$

where $\mathcal{A} = \{\min\{Y_0 + \eta h_0(X), Y_1 + \eta h_1(X)\} : \tau \in \mathcal{T}, \eta \in [1, B]\}$, which follows by similar arguments to the proof of Proposition 4.2. ■

Proof of Proposition 4.5. Follows by identical arguments to Proposition 4.3, replacing Δ_i with Δ_i^* . ■

Proof of Proposition A.1. The Lagrangian is

$$L = \inf_Q \sup_{\eta \geq 0} (\mathbb{E}_Q[Y_1 \tau(X) + Y_0(1 - \tau(X))] + \eta(W_2(P, Q)^2 - \varepsilon^2)).$$

Following the same steps as the proof of Proposition 3.1, we deduce that the Lagrangian dual is

$$L^* = \sup_{\eta \geq 0} \mathbb{E}_P \left[\inf_{\tilde{y}_0, \tilde{y}_1} (\tilde{y}_1 \tau(X) + \tilde{y}_0(1 - \tau(X)) + \eta((Y_0 - \tilde{y}_0)^2 + (Y_1 - \tilde{y}_1)^2 - \varepsilon^2)) \right].$$

Recall that we take $\mathcal{Y} = \mathbb{R}$. When $\eta = 0$ the inner infimum is $-\infty$ and $L^* = -\infty$. Now suppose $\eta > 0$. Fix $Z = (X, Y_0, Y_1)$. If $\tau(X) = 1$, then the inner infimum over \tilde{y}_0 is attained at $\tilde{y}_0 = Y_0$. The inner infimum over \tilde{y}_1 reduces to minimizing $y \mapsto y + \eta(Y_1 - y)^2$, which is achieved at $y = Y_1 - \frac{1}{2\eta}$. A symmetric argument applies when $\tau(X) = 0$. The

dual value is therefore

$$L^* = \sup_{\eta > 0} \mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \frac{1}{4\eta} - \eta\varepsilon^2.$$

The global maximum occurs at $\eta = 1/(2\varepsilon) > 0$, which implies a dual value of

$$L^* = \mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \varepsilon.$$

We now show that $L \leq L^*$. Define $\tilde{Z} = T(Z)$ where

$$T(x, y_0, y_1) = \begin{cases} (x, y_0 - \varepsilon, y_1) & \text{if } x \in \mathcal{X}_0, \\ (x, y_0, y_1 - \varepsilon) & \text{if } x \in \mathcal{X}_1. \end{cases}$$

Let Q_0 and π denote the distribution of $T(Z)$ and $(Z, T(Z))$, respectively, with $Z \sim P$. Then $\mathbb{E}_{Q_0}[Y_1\tau(X) + Y_0(1 - \tau(X))] = \mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \varepsilon$ and under the metric (6), we have

$$W_2(P, Q_0)^2 \leq \mathbb{E}_\pi[d((X, Y_0, Y_1), (\tilde{X}, \tilde{Y}_0, \tilde{Y}_1))^2] = \mathbb{E}_P[\varepsilon^2\mathbb{I}[X \in \mathcal{X}_0] + \varepsilon^2\mathbb{I}[X \in \mathcal{X}_1]] = \varepsilon^2.$$

As Q_0 is feasible for the outer minimization over $Q \in \mathcal{Q}$, we obtain

$$\begin{aligned} L &\leq \sup_{\eta \geq 0} (\mathbb{E}_{Q_0}[Y_1\tau(X) + Y_0(1 - \tau(X))] + \eta(W_2(P, Q_0)^2 - \varepsilon^2)) \\ &= \mathbb{E}_P[Y_1\tau(X) + Y_0(1 - \tau(X))] - \varepsilon = L^*, \end{aligned}$$

as required. ■

Proof of Proposition A.2. We follow similar arguments to the proof of Proposition 3.1, and state only the necessary modifications. The Lagrangian dual is

$$L^* = \sup_{\eta \geq 0} \mathbb{E}_P \left[\inf_{\tilde{y}_0, \tilde{y}_1} (\tilde{y}_1\tau(X) - \tilde{y}_0 + \eta(|Y_0 - \tilde{y}_0| + |Y_1 - \tilde{y}_1| - \varepsilon)) \right].$$

If $\eta \in [0, 1)$, the inner infimum is achieved by taking $\tilde{y}_0 = \bar{Y}$ (or $\rightarrow +\infty$ if $\bar{Y} = +\infty$) and $\tilde{y}_1 = Y_1$ when $\tau(X) = 0$ and $\tilde{y}_1 = \underline{Y}$ (or $\rightarrow -\infty$ if $\underline{Y} = -\infty$) when $\tau(X) = 1$. Similarly, if $\eta \geq 1$ then it is optimal to take $\tilde{y}_0 = Y_0$ and $\tilde{y}_1 = Y_1$. Optimizing with respect to η then yields the desired result. ■

References

- Abadie, A. (2003). Semiparametric instrumental variable estimation of treatment response models. *Journal of Econometrics* 113(2), 231–263.
- Abbring, J. H. and J. J. Heckman (2007). Econometric evaluation of social programs, Part III: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation. Volume 6 of *Handbook of Econometrics*, Chapter 72, pp. 5145–5303. Elsevier.
- Athey, S. and S. Wager (2021). Policy learning with observational data. *Econometrica* 89(1), 133–161.
- Bhattacharya, D. and P. Dupas (2012). Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics* 167(1), 168–196.
- Bloom, H. S., L. L. Orr, S. H. Bell, G. Cave, F. Doolittle, W. Lin, and J. M. Bos (1997). The benefits and costs of JTPA Title II-A programs: Key findings from the national job training partnership act study. *The Journal of Human Resources* 32(3), 549–576.
- Chernozhukov, V. and C. Hansen (2005). An IV model of quantile treatment effects. *Econometrica* 73(1), 245–261.
- Chesher, A. (2003). Identification in nonseparable models. *Econometrica* 71(5), 1405–1441.
- Fan, Y. and S. S. Park (2010). Sharp bounds on the distribution of treatment effects and their statistical inference. *Econometric Theory* 26(3), 931–951.
- Feng, Q., Q. Vuong, and H. Xu (2020). Estimation of heterogeneous individual treatment effects with endogenous treatments. *Journal of the American Statistical Association* 115(529), 231–240.
- Heckman, J. J., J. Smith, and N. Clements (1997). Making The Most Out Of Programme Evaluations and Social Experiments: Accounting For Heterogeneity in Programme Impacts. *The Review of Economic Studies* 64(4), 487–535.
- Heckman, J. J. and E. J. Vytlacil (2007). Econometric evaluation of social programs, part ii: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments. Volume 6 of *Handbook of Econometrics*, Chapter 71, pp. 4875–5143. Elsevier.
- Imai, K. and M. Ratkovic (2013). Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics* 7(1), 443 – 470.
- Imbens, G. and K. Menzel (2021). A causal bootstrap. *The Annals of Statistics* 49(3), 1460–1488.

- Imbens, G. W. and D. B. Rubin (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.
- Kallus, N. (2017, 06–11 Aug). Recursive partitioning for personalization using observational data. In D. Precup and Y. W. Teh (Eds.), *Proceedings of the 34th International Conference on Machine Learning*, Volume 70 of *Proceedings of Machine Learning Research*, pp. 1789–1798. PMLR.
- Kallus, N. (2021). More efficient policy learning via optimal retargeting. *Journal of the American Statistical Association* 116(534), 646–658.
- Kitagawa, T. and A. Tetenov (2018). Who should be treated? empirical welfare maximization methods for treatment choice. *Econometrica* 86(2), 591–616.
- Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica* 72(4), 1221–1246.
- Matzkin, R. L. (2003). Nonparametric estimation of nonadditive random functions. *Econometrica* 71(5), 1339–1375.
- Mbakop, E. and M. Tabord-Meehan (2021). Model selection for treatment choice: Penalized welfare maximization. *Econometrica* 89(2), 825–848.
- Mo, W., Z. Qi, and Y. Liu (2021). Learning optimal distributionally robust individualized treatment rules. *Journal of the American Statistical Association* 116(534), 659–674.
- Nie, X. and S. Wager (2020). Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika* 108(2), 299–319.
- Qi, Z., J.-S. Pang, and Y. Liu (2022). On robustness of individualized decision rules. *Journal of the American Statistical Association* 0(0), 1–15.
- Qian, M. and S. A. Murphy (2011). Performance guarantees for individualized treatment rules. *The Annals of Statistics* 39(2), 1180–1210.
- Rosenbaum, P. R. and D. B. Rubin (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* 70(1), 41–55.
- Si, N., F. Zhang, Z. Zhou, and J. Blanchet (2020, 13–18 Jul). Distributionally robust policy evaluation and learning in offline contextual bandits. In H. D. III and A. Singh (Eds.), *Proceedings of the 37th International Conference on Machine Learning*, Volume 119 of *Proceedings of Machine Learning Research*, pp. 8884–8894. PMLR.
- Spini, P. E. (2021). Robustness, heterogeneous treatment effects and covariate shifts. *arXiv preprint arXiv:2112.09259*.
- Stoye, J. (2012). Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics* 166, 138–156.

- Swaminathan, A. and T. Joachims (2015). Batch learning from logged bandit feedback through counterfactual risk minimization. *Journal of Machine Learning Research* 16(52), 1731–1755.
- Uehara, M., M. Kato, and S. Yasui (2020). Off-policy evaluation and learning for external validity under a covariate shift. In H. Larochelle, M. Ranzato, R. Hadsell, M.-F. Balcan, and H.-T. Lin (Eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- van der Vaart, A. W. and J. A. Wellner (1996). *Weak convergence and empirical processes*. Springer.
- Villani, C. (2009). *Optimal Transport: Old and New*. Springer.
- Vuong, Q. and H. Xu (2017). Counterfactual mapping and individual treatment effects in nonseparable models with binary endogeneity. *Quantitative Economics* 8(2), 589–610.
- Wager, S. and S. Athey (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523), 1228–1242.
- Yue, M.-C., D. Kuhn, and W. Wiesemann (2021). On linear optimization over Wasserstein balls. *Mathematical Programming*.
- Zhao, Y., D. Zeng, A. J. Rush, and M. R. Kosorok (2012). Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association* 107(499), 1106–1118. PMID: 23630406.