

Collegio Carlo Alberto



Equality-Minded Treatment Choice

Toru Kitagawa
Aleksy Tetenov

No. 496
February 2017

Carlo Alberto Notebooks

www.carloalberto.org/research/working-papers

Equality-Minded Treatment Choice*

Toru Kitagawa[†] and Aleksey Tetenov[‡]

February 18, 2017

Abstract

The goal of many randomized experiments and quasi-experimental studies in economics is to inform policies that aim to raise incomes and reduce economic inequality. A policy maximizing the sum of individual incomes may not be desirable if it magnifies economic inequality and post-treatment redistribution of income is infeasible. This paper develops a method to estimate the optimal treatment assignment policy based on observable individual covariates when the policy objective is to maximize an *equality-minded rank-dependent social welfare function*, which puts higher weight on individuals with lower-ranked outcomes. We estimate the optimal policy by maximizing a sample analog of the rank-dependent welfare over a properly constrained set of policies. Although an analytical characterization of the optimal policy under a rank-dependent social welfare is not available even with the knowledge of potential outcome distributions, we show that the average social welfare attained by our estimated policy converges to the maximal attainable welfare at $n^{-1/2}$ rate uniformly over a large class of data distributions. We also show that this rate is minimax optimal. We provide an application of our method using the data from the National JTPA Study.

Keywords: Program evaluation, Treatment choice, Social welfare, Inequality index, Gini coefficient

*We would like to thank Ignacio Monzón and John Weymark for useful comments. We have also benefited from presenting this work at Academia Sinica, 2015 Bristol Econometric Study Group, RES, Cowles Conference on Econometrics of Randomized Experiments, IAAE, SCW, AMES, Cemmap Advances in Econometrics, and CEU. Financial support from the ESRC through the ESRC Centre for Microdata Methods and Practice (CeMMAP) (grant number RES-589-28-0001) and the European Research Council (Starting grant No. 715940) is gratefully acknowledged. Toru Kitagawa thanks the Visiting Researcher Fellowship from Collegio Carlo Alberto.

[†]Cemmap/University College London, Department of Economics. Email: t.kitagawa@ucl.ac.uk

[‡]University of Bristol and Collegio Carlo Alberto, Email: a.tetenov@bristol.ac.uk

1 Introduction

In causal inference studies analyzing experimental or quasi-experimental data, treatment response generally varies with individual observable characteristics. Learning about such heterogeneity from the data is essential for designing *individualized treatment rules* that assign treatments on the basis of individual observable characteristics. The optimal individualized treatment rule maximizes a social welfare criterion representing the policy maker’s preferences over population distributions of post-treatment outcomes. The literature on statistical treatment choice initiated by Manski (2004) emphasizes this perspective of welfare-based empirical policy design and pursues statistically sound ways to estimate the optimal treatment assignment rule.

Research on statistical treatment rules typically focuses on the *utilitarian social welfare* criterion defined by the mean of the outcomes in the population,¹ even though welfare economics offers a variety of alternative criteria. The utilitarian social welfare criterion offers analytical and computational convenience because it is additive across subgroups of the population and depends only on the mean of the outcome distribution. The optimal treatment rule then depends only on the conditional average treatment effect. Empirical researchers studying causal impacts of social programs have stressed the importance of evaluating distributional impacts, which are overlooked when only mean outcomes are considered (e.g., Bitler et al. (2006)). The distributional impact of a policy is especially important when the policy maker is concerned about *economic inequality* in the population.

We study the problem of treatment assignment that aims to maximize a rank-dependent *social welfare function (SWF)*, which has the form

$$W \equiv \int Y_i \cdot \omega(\text{Rank}(Y_i)) di, \tag{1.1}$$

where Y_i is individual i ’s outcome, $\text{Rank}(Y_i)$ is the outcome rank of i from the bottom of the outcome distribution, and $\omega(\cdot)$ is a non-negative weight assigned to each rank. The class of *generalized Gini* SWFs proposed by Mehran (1976) and Weymark (1981) consists of SWFs of the form (1.1) with *non-increasing* $\omega(\cdot)$. It closely relates to *income inequality*

¹More generally, the utilitarian welfare is defined as the mean of a function of applied to each individual outcome.

indices, including the widely-used Gini index. Blackorby and Donaldson (1978) show that, given a specification of $\omega(\cdot)$, the rank-dependent SWF can be written as a product of the average outcome and one minus the generalized relative index of inequality, e.g., Gini. This implies that these SWFs generate a ranking of outcome distributions that is increasing in the average outcome and decreasing in the chosen index of inequality. While inequality measures are predominantly applied to net income, our analysis allows Y_i to denote other outcome variables of interest, including functions of income, consumption, wealth, or human capital. We will therefore refer to Y_i simply as “the outcome” in this paper.

The goal is to choose a treatment rule δ that assigns individuals to one of two treatments $d \in \{0, 1\}$ depending on their observable pre-treatment covariates $X \in \mathcal{X}$. This choice is made after experimental data has been collected and analyzed. We do not consider the problem of optimal experimental design in this paper, taking the design as given. We assume that an individual’s treatment outcome does not depend on treatments received by others. The policy-maker in our setup can only impact the distribution of outcomes through the choice of a treatment assignment rule and cannot combine it with other redistributive policies.

Finding a policy that maximizes a rank-dependent SWF is a non-trivial problem without a closed-form solution even if the conditional distributions of potential outcomes ($P(Y_0|X)$ and $P(Y_1|X)$) are known. A utilitarian SWF ($\int u(Y_i)di$) is additive across subgroups of the population, so it is optimal to assign for each subgroup the treatment with the highest conditional mean $E(u(Y_d)|X)$. In contrast, a rank-dependent SWF is not separable across subgroups, as the ranking of treatment assignment rules for a given covariate value may change depending on the treatment assignment of other subgroups.

We show in Theorem 2.1 that an equality-minded rank-dependent SWF is always maximized by a non-randomized treatment rule (assigning the same treatment to all individuals with identical covariates). This result greatly simplifies the space of treatment rules that need to be considered. It also allows us to index treatment rules by their *decision sets* $G \subset \mathcal{X}$, denoting all values of the covariates $\{X \in G\}$ for which treatment 1 is assigned.

Our aim is to estimate from the sample data a treatment assignment rule $\hat{G} \in \mathcal{G}$ belonging to a constrained (but generally large) *set of feasible policies* $\mathcal{G} \equiv \{G \subset \mathcal{X}\}$, which is a collection of non-randomized treatment rules indexed by their decision sets. Policy makers

often face legal, ethical, or political constraints that restrict how individual characteristics can be used to determine treatment assignment. One of the advantages of our framework is that it easily incorporates such exogenous restrictions. Our analytical results also require \mathcal{G} to satisfy a certain complexity restriction (a finite *VC-dimension*) to prevent overfitting. As argued in Kitagawa and Tetenov (2015), this is not restrictive for public policy applications and provide rich examples of treatment rule classes that satisfy this complexity restriction.

We propose estimating the treatment rule \hat{G} by maximizing a sample analog of $W(G)$, the SWF evaluated at the population distribution of Y that would realize if treatment assignment rule G is implemented. This idea of estimating a policy by maximizing an empirical welfare criterion is in line with the method proposed in Kitagawa and Tetenov (2015) for the utilitarian welfare case. Following the terminology suggested there, we also refer to the method proposed in the current paper for the rank-dependent social welfare as *Empirical Welfare Maximization (EWM)*. The maximum welfare achievable by treatment rules from the constrained class \mathcal{G} is $W_{\mathcal{G}}^* \equiv \sup_{G \in \mathcal{G}} W(G)$. We assess the statistical performance of \hat{G} in terms of the average welfare it achieves (with respect to the sampling distribution P^n of a size n sample). Specifically, we derive a non-asymptotic and distribution-free upper bound on regret $W_{\mathcal{G}}^* - E_{P^n}[W(\hat{G})]$ in terms of the sample size n and a measure of complexity of \mathcal{G} , and show that it converges to zero at $n^{-1/2}$ rate. We also show that this rate is minimax optimal over a minimally constrained class of population distributions, ensuring that no other data-driven treatment rule can lead to a faster welfare loss convergence rate uniformly over the class of distributions. Even though a rank-dependent SWF depends on the entire conditional distributions of potential outcomes given covariates, the welfare convergence rate is the same as the minimax optimal rate that the EWM rule attains in the utilitarian case (Theorems 2.1 and 2.2 in Kitagawa and Tetenov (2015)).

The remainder of this paper is organized as follows. Section 1.1 provides an overview of related literature. Section 2 discusses the properties of equality-minded rank-dependent social welfare functions and their application to the analysis of treatment choice. In Section 3, we introduce the general analytical framework and show the convergence rate properties of the EWM rule for rank-dependent welfare. Section 4 provides several extensions of the model including the cases where propensity scores need to be estimated or where a cost or

capacity constraint is present. In Section 5, we apply our method to the experimental data from the National JTPA study. We collect all the proofs in the appendix.

1.1 Related Literature

The analysis of statistical treatment rules was pioneered by Manski (2004), and is a growing area of research in econometrics. Important recent developments can be found in Dehejia (2005), Hirano and Porter (2009), Stoye (2009, 2012), Chamberlain (2011), Bhattacharya and Dupas (2012), Tetenov (2012), Kasy (2014, 2016) and Kitagawa and Tetenov (2015), among others. All the existing works on treatment choice except for Kasy (2016) exclusively posit a utilitarian welfare criterion (the sum of the outcomes in the population) as the objective function of the policy maker. Motivated by the policy concerns about economic inequality, the current paper analyzes the treatment choice problem for a class of rank-dependent social welfare functions that embody inequality aversion. The main contributions of the current paper are (i) introducing a rank-dependent social welfare function into the treatment choice problem, (ii) developing an estimation method for the optimal policy, and (iii) characterizing its optimality in terms of the convergence rate of the average welfare loss. Several features distinguish the current analysis from the EWM approach for the utilitarian welfare case considered in Kitagawa and Tetenov (2015): (i) the analytical characterization of the optimal assignment policy is not available even with the knowledge of the distribution of (Y_0, Y_1, X) , (ii) the empirical welfare criterion involves a nonlinear transformation of the empirical distribution, and (iii) as illustrated in Section 2.1 below, the rank-dependent welfare criterion is non-decomposable, so optimal assignment for one subpopulation depends on what treatment is assigned to other subpopulations. These features also distinguish this paper from machine learning and statistics literature on empirical risk minimization problems (Vapnik (1998)), where the empirical risk criterion always takes the form of a sample average.

Within the econometrics literature on policy choice, Kasy (2016) is the only existing paper that analyzes a class of social welfare functions including the rank-dependent social welfare. Our approach differs from his in several aspects. First, Kasy (2016) considers a linear approximation of the rank-dependent welfare function around a status-quo policy in order to discuss (partial) identification of a welfare-improving local policy change. We instead

focus on a globally optimal policy without invoking approximations. Second, we assume that the welfare criterion is point-identified by the sampling process, while Kasy (2016) focuses on partial identification of the welfare criterion and construction of social planner’s incomplete preference ordering over policies. Third, we study estimation of an optimal policy and examine optimality of the estimator in terms of welfare regret convergence rate, while Kasy (2016) studies asymptotically valid inference on the welfare rankings.

The causal impact on non-utilitarian social welfare has been rarely considered in program evaluation studies. Aaberge et al. (2016) estimate a rank-dependent social welfare function of two policy alternatives: with and without the uniform implementation of the treatment. Firpo and Pinto (2015) estimate the impact of uniform implementation of the treatment on measures of inequality, including the Gini coefficient. Instead, the focus of the current paper is on estimating the optimal treatment rule from a large class of individualized assignment rules.

We consider social welfare functions that satisfy the axiom of *anonymity*, i.e., they are functionals of the distribution of the outcomes after treatment assignment and indifferent to reshuffling of the outcomes between individuals. Thus our objective does not depend on the distribution of individual treatment effects $P(Y_1 - Y_0)$, which has also received attention recently in the program evaluation literature (Heckman et al. (1997b), Firpo and Ridder (2008), Fan and Park (2010)).

Building social welfare function satisfying the Pigou-Dalton principle of transfers is one of the central themes in the literature of inequality measurement and welfare economics (see Cowell (1995, 2000), Lambert (2001)). Currently, there are two widely-used classes of social welfare functions that meet the Pigou-Dalton principle. The first one is the class of Atkinson-type SWFs (Atkinson (1970)), $W(F) = \int_0^{-\infty} U(y)dF(y)$, where $F(y)$ is the cumulative distribution function (cdf) of the outcome and $U(\cdot)$ is a concave non-decreasing function. Since the Atkinson-type social welfare function is linear in F , the EWM approach of Kitagawa and Tetenov (2015) can be readily applied by defining the outcome observations as $U(Y)$. The second class, which is this paper’s main focus, is the class of rank-dependent social welfare functions introduced by Mehran (1976), Blackorby and Donaldson (1978) and Weymark (1981) and axiomatized by Yaari (1988). The key axiom of Yaari (1988) that

distinguishes rank-dependent social welfare from Atkinson-type social welfare is *Invariance under a Rank-Preserving Lump-Sum Change of Incomes at the Upper Tail*, which means that the preference ordering between two income distributions F and F' is invariant to any fixed lump-sum increase (decrease) in income of all those above (below) the τ -th quantile of F and F' for any $\tau \in (0, 1)$. On the other hand, the key axiom that characterizes the Atkinson-type social welfare is the independence axiom: the preference ordering between F and F' is invariant to any mixing with another common income distribution.² These rich and insightful works in welfare economics have not yet been well linked to econometrics and empirical analysis for policy design. One of the main aims of the current paper is to establish a link between these two important literatures.

2 Treatment Choice with Equality-Minded Social Welfare Functions

We call a SWF *equality minded* if it satisfies the *Pigou-Dalton Principle of Transfers*: A transfer of income from a higher ranked individual to a lower ranked individual is always desirable when it does not change their ranks in the income distribution. The equality-minded SWFs analyzed in this paper are the rank-dependent SWFs with decreasing welfare weights (also called generalized Gini SWFs), introduced by Mehran (1976) and Weymark (1981) and axiomatized by Yaari (1988). An equality-minded rank-dependent SWF admits the following representation:

$$W(F) = \int_{\mathbb{R}} \Lambda(F(y)) dy, \tag{2.1}$$

where $\Lambda(\cdot) : [0, 1] \rightarrow [0, 1]$ is a non-increasing, non-negative function with $\Lambda(0) = 1$ and $\Lambda(1) = 0$. The rank-dependent SWF meets the Pigou-Dalton Principle of Transfers if and only if $\Lambda(\cdot)$ is convex.

²As noted in Machina (1982), the rank-dependent social welfare function generalizes the Atkinson-type social welfare exactly as the rank-dependent expected utility theory generalizes the classical von Neumann-Morgenstern expected utility theory (Machina (1982) and Quiggin (1982)) by relaxing the controversial independence axiom.

The term *rank-dependent* is due to an equivalent representation of (2.1) as a weighted sum of incomes. If $\Lambda(\cdot)$ is almost everywhere differentiable, then by integration by parts $W(F)$ can be equivalently expressed as

$$W(F) = \int_0^1 F^{-1}(t)\omega(t)dt, \quad (2.2)$$

where $F^{-1}(t) \equiv \inf\{y : F(y) \geq t\}$ is the outcome of individuals at t -th quantile and $\omega(t) = \frac{d[1-\Lambda(t)]}{dt}$ is the welfare weight assigned to that quantile. Thus $W(F)$ is a weighted average of outcomes Y , where the $\omega(\cdot)$ specifies the *rank-specific* weight. If the SWF is equality-minded then $\Lambda(\cdot)$ is convex, hence $\omega(\cdot)$ is non-increasing and assigns larger welfare weights to individuals with lower-ranked outcomes.

Throughout the paper, we consider equality-minded SWFs satisfying the following assumption:

Assumption 2.1 (SWF).

The policy-maker's SWF has representation (2.1), where $\Lambda(\cdot) : [0, 1] \rightarrow [0, 1]$ is a non-increasing, convex function with $\Lambda(0) = 1$ and $\Lambda(1) = 0$. We also assume that $\Lambda(\cdot)$ is almost everywhere differentiable and its right derivative at the origin denoted by $\Lambda'(0)$ is bounded, $|\Lambda'(0)| < \infty$.

Assumption 2.1 states that the rank specific weight function $\omega(\cdot)$ defined in (2.2) does not asymptote at the origin, implying that the rank specific weight assigned to the lowest rank is bounded. For instance, this assumption is satisfied for the extended Gini family of social welfare functions $\{W_k(F) = \int_0^\infty (1 - F(y))^{k-1} : 2 < k < \infty\}$ considered in Donaldson and Weymark (1980, 1983) and Aaberge et al. (2016). On the other hand, the Rawlsian social welfare that can be approximated by $\lim_{k \rightarrow \infty} W_k(F)$ is ruled out from our analysis.

One prominent example from this class is the *Gini social welfare function* (Blackorby and Donaldson (1978), Weymark (1981)),

$$W_{Gini}(F) \equiv E(Y)(1 - I_{Gini}(F)), \quad (2.3)$$

where $I_{Gini}(F) = 1 - \frac{\int_0^1 F^{-1}(t)2(1-t)dt}{E(Y)}$ is the Gini inequality index. This social welfare can be represented in the form of (2.2) with $\omega(t) = 2(1 - t)$, implying $\Lambda(t) = (1 - t)^2$.

We consider the problem of choosing a policy that assigns individuals to one of two treatments $d \in \{0, 1\}$ in order to maximize the chosen SWF. A treatment assignment rule $\delta : \mathcal{X} \rightarrow [0, 1]$ specifies the proportion of individuals with observable pre-treatment covariates $X \in \mathcal{X} \subset \mathbb{R}^{d_x}$ who will be assigned to treatment 1 by the policy-maker. The policy randomly assigns individuals with covariates X to the two treatments with probabilities $1 - \delta(X)$ and $\delta(X)$. The population distribution of outcomes induced by treatment rule δ has cdf

$$F_\delta(y) \equiv \int_{\mathcal{X}} [(1 - \delta(x))F_{Y_0|X=x}(y) + \delta(x)F_{Y_1|X=x}(y)] dP_X(x), \quad (2.4)$$

where Y_0 and Y_1 denote the potential outcomes of the two treatments with conditional distributions $F_{Y_0|X}$ and $F_{Y_1|X}$ given X and P_X is the marginal distribution of X .

If the population distribution of (Y_0, Y_1, X) were known, the optimal policy maximizing the social welfare function (2.1) would be

$$\delta^* \in \arg \max_{\delta} W(F_\delta). \quad (2.5)$$

For the utilitarian welfare function (the mean of Y), the welfare maximization problem simplifies to

$$\delta_{util}^* \in \arg \max_{\delta} \int_{\mathcal{X}} [(1 - \delta(x))E(Y_0|X = x) + \delta(x)E(Y_1|X = x)] dP_X(x). \quad (2.6)$$

The utilitarian social welfare is additive across covariates and depends on the outcome distributions only through their conditional means $E(Y_d|X)$. The optimal utilitarian policy is

$$\delta_{util}^* = 1 \{x \in \mathcal{X} : E(Y_1|X = x) > E(Y_0|X = x)\}.$$

In contrast, the optimal rule for a rank-dependent welfare function (2.1) depends on the whole conditional distributions of potential outcomes $F_{Y_0|X}$ and $F_{Y_1|X}$, not only on their means. The optimal rule can differ from the utilitarian one if there is no first-order stochastic dominance relationship between $F_{Y_0|X}$ and $F_{Y_1|X}$ for some covariate values.

Even with the knowledge of the distribution of (Y_0, Y_1, X) , a simple characterization of the optimal rule does not seem available for rank-dependent SWFs. The following theorem mitigates this complication by substantially reducing the set of candidate treatment rules that need to be considered.

Theorem 2.1. *If $W(\cdot)$ satisfies Assumption 2.1, then for every measurable treatment rule $\delta : \mathcal{X} \rightarrow [0, 1]$, there exists a non-randomized treatment rule $\delta_G(x) \equiv 1\{x \in G\}$ for some Borel set $G \subset \mathcal{X}$, such that $W(F_{\delta_G}) \geq W(F_\delta)$.*

If all upper level sets of δ belong to a collection \mathcal{G} of Borel subsets of \mathcal{X} :

$$\{x : \delta(x) \geq t\} \in \mathcal{G}, \quad \forall t \in \mathbb{R},$$

then there exists $\delta_G(x)$, $G \in \mathcal{G}$, such that $W(F_{\delta_G}) \geq W(F_\delta)$.

Proof. See Appendix A. □

This theorem shows that a treatment assignment rule maximizing an equality-minded rank-dependent welfare is non-randomized (assigns all individuals with the same covariates to the same treatment). We can therefore restrict our search for an optimal policy to the set of non-randomized rules that can be succinctly characterized by their *decision sets* $G \subset \mathcal{X}$. Decision set G determines the group of individuals $\{X \in G\}$ to whom treatment 1 is assigned. With abuse of notation, we denote the welfare of a non-randomized treatment rule with decision set G by $W(G)$, suppressing the cumulative distribution function in its argument,

$$\begin{aligned} W(G) &\equiv W(F_G), \\ F_G(y) &\equiv \int_{\mathcal{X}} [F_{Y_0|X=x}(y)1\{x \notin G\} + F_{Y_1|X=x}(y)1\{x \in G\}] dP_X(x). \end{aligned} \quad (2.7)$$

Our goal is to estimate from the sample data a treatment assignment rule that attains the maximum level of social welfare $W_G^* \equiv \sup_{G \in \mathcal{G}} W(G)$ over the *set of feasible policies* $\mathcal{G} \equiv \{G \subset \mathcal{X}\}$, which is a collection of non-randomized treatment rules (subsets of the covariate space \mathcal{X}). An important feature of our empirical welfare maximization approach is that the complexity of \mathcal{G} is constrained by a finite *Vapnik-Cervonenkis (VC) dimension*:

Assumption 2.2 (VC).

The class of decision sets \mathcal{G} has a finite VC-dimension $v < \infty$.

The VC-dimension is a restriction on the complexity of the set of feasible policies. Without it, maximizing a sample analog of $W(G)$ over G can lead to arbitrarily complicated policies (overfitting) and prevent us from learning the optimal policy on the basis of a finite

number of observations. It does not require \mathcal{G} to be finite and allows for very large classes of treatment rules. For example, a class of treatment rules defined by a linear equation in functions of x , $\mathcal{G} \equiv \{G = \{x : \sum_{j=1}^m \beta_j f_j(x) \geq 0\}, \beta \in \mathbb{R}^m\}$ has a finite VC-dimension. See Kitagawa and Tetenov (2015) for other examples of classes \mathcal{G} that satisfy Assumption 2.2.

2.1 Illustrating Example

In this section, we illustrate the properties of rank-dependent SWFs in comparison with the utilitarian one in a simple setting with the Gini SWF, $W_{Gini}(F) = \int_0^\infty (1 - F(y))^2 dy$. We first compare the welfare ordering on the parametric family of log-normal outcome distributions. Second, we consider a simple treatment choice problem with binary X in order to illustrate how the optimal rules fundamentally differ between the two SWFs.

First, consider the welfare ordering over the family of log-normal distributions of outcomes, $Y \sim \log N(\mu, \sigma^2)$, ignoring the treatment choice problem. The mean of Y is given by $E(Y) = \exp(\mu + \sigma^2/2)$. The Gini inequality coefficient for $\log N(\mu, \sigma^2)$ is given by $2\Phi(\sigma/\sqrt{2}) - 1$ (see, e.g., Cowell (1995)), where $\Phi(\cdot)$ is the cdf of the standard normal distribution. By (2.3), we have

$$W_{Gini}(F) = 2 \exp\left(\mu + \frac{\sigma^2}{2}\right) \left[1 - \Phi\left(\frac{\sigma}{\sqrt{2}}\right)\right]. \quad (2.8)$$

This welfare function is increasing in μ , whereas it is not monotonic in σ . For instance, when $\mu = 0$, $W_{Gini}(F_Y)$ is decreasing in σ for $\sigma < 0.87$ and increasing for $\sigma > 0.87$. See Figure 1 for a plot of $W_{Gini}(F)$ over $\sigma \in [0, 2]$ holding $\mu = 0$ fixed. The U-shape of the Gini social welfare indicates that for $\sigma < 0.87$, the negative contribution to the social welfare from an increase in the Gini coefficient dominates the positive contribution from an increase in the mean, while for $\sigma > 0.87$, this relationship reverses. In Figure 2, we plot the densities of the log-normal distributions for $\sigma = 0.25, 0.5$, and 1 . Since $E(Y)$ is monotonically increasing both in μ and σ , higher σ is always preferable in terms of the utilitarian social welfare. In contrast, as shown in the welfare values plotted in Figure 1, the Gini social welfare yields the complete opposite welfare ordering over the three log-normal distributions in Figure 2.

Consider now the treatment choice problem. Suppose there is only one binary covariate $X \in \{a, b\}$ with $\Pr(X = a) = \Pr(X = b) = 1/2$. Consider the following parameterization of

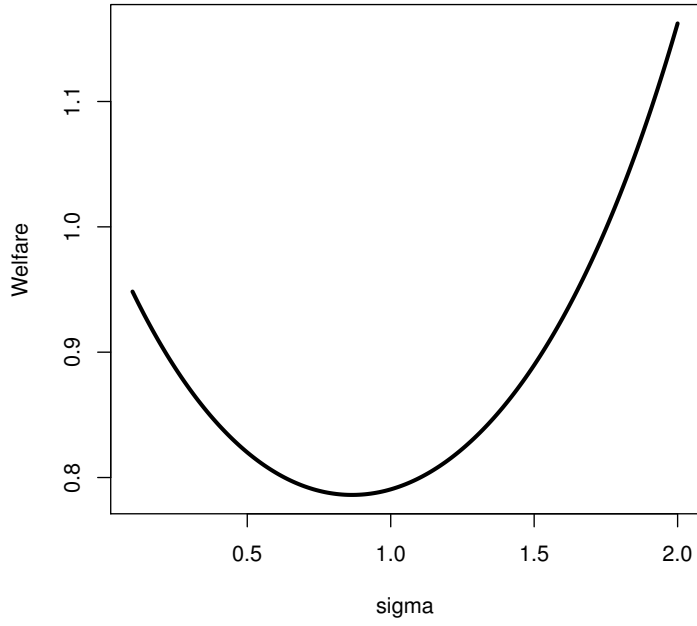


Figure 1: Equality-minded welfare for $\log N(0, \sigma^2)$.

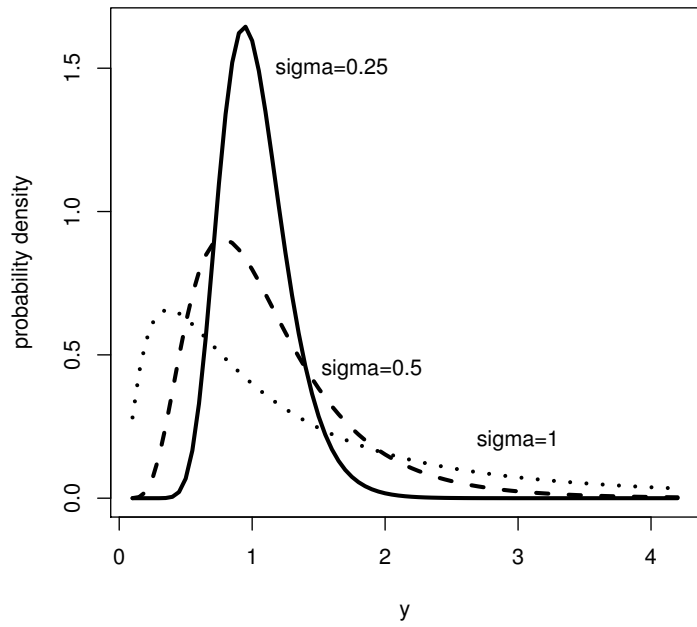


Figure 2: Density of $\log N(0, \sigma^2)$.

the potential outcome distributions:

$$\begin{aligned} Y_1|(X = a) &\sim \log N(\mu_a, \sigma_a^2), & Y_0|(X = a) &\sim \log N(0, 0.8^2), \\ Y_1|(X = b) &\sim \log N(\mu_b, \sigma_b^2), & Y_0|(X = b) &\sim \log N(0, 0.8^2). \end{aligned} \quad (2.9)$$

According to Theorem 2.1, it suffices to consider non-randomized rules to search for an optimal one. We therefore consider ranking the following four policies: $\mathcal{G} = \{\emptyset, \{a\}, \{b\}, \{a, b\}\} \equiv \{G_\emptyset, G_a, G_b, G_{ab}\}$.

Suppose $\sigma_a = \sigma_b = 0.8$ and $\mu_a, \mu_b > 0$. Then, in each subpopulation of $X = a$ and $X = b$, the distribution of Y_1 stochastically dominates the distribution of Y_0 . Since the rank-dependent social welfare is clearly monotonic in the first-order stochastic dominance relationship, treating both $\{X = a\}$ and $\{X = b\}$ maximizes the Gini social welfare. This optimal rule indeed coincides with that of the utilitarian welfare case. In general, when the stochastic dominance relationships between $Y_1|X$ - and $Y_0|X$ -distributions are present for all X , the optimal rule for the rank-dependent social welfare agrees with the utilitarian one and can be obtained by solving the treatment choice problem separately in each subpopulation.

These results change drastically once we let $\sigma_a \neq \sigma_b$. Suppose we fix $\mu_a = \mu_b = 0$, while we vary both σ_a and σ_b over $[0.1, 1.6]$. As the mean of a log normal random variable is increasing in σ , the optimal treatment rule for the utilitarian welfare is obtained by

$$G_{util}^* = \begin{cases} G_\emptyset & \text{if } \sigma_a < 0.8 \text{ and } \sigma_b < 0.8, \\ G_a & \text{if } \sigma_a \geq 0.8 \text{ and } \sigma_b < 0.8, \\ G_b & \text{if } \sigma_a < 0.8 \text{ and } \sigma_b \geq 0.8, \\ G_{ab} & \text{if } \sigma_a \geq 0.8 \text{ and } \sigma_b \geq 0.8. \end{cases}$$

In Figure 3, we plot the utilitarian optimal treatment rule at each grid point of $(\sigma_a, \sigma_b) \in [0.1, 1.6]^2$. Since the utilitarian social welfare is additive over the subpopulations, a treatment preferable for one subpopulation does not depend on the treatment assigned to the other subpopulation. The regions in which different rules from \mathcal{G} are optimal form a quadrant partition, as shown in Figure 3.

In Figure 4, we plot the optimal policies in terms of the Gini social welfare. The regions in which different rules from \mathcal{G} are optimal are strikingly different compared with the utilitarian

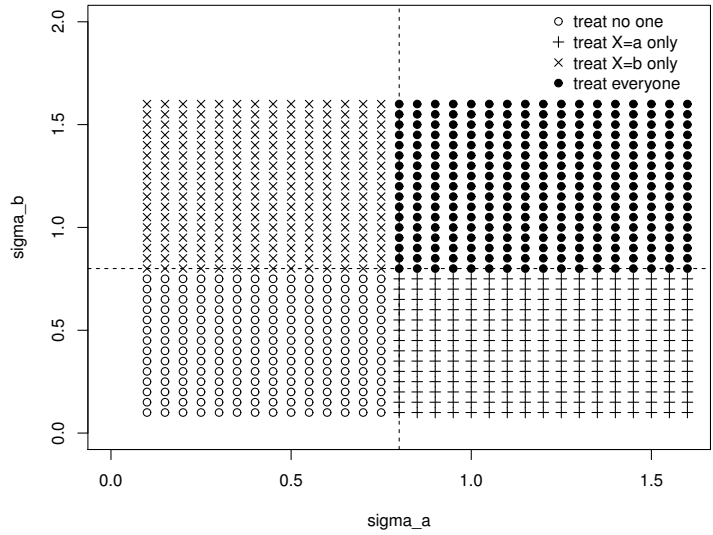


Figure 3: Optimal policies under the utilitarian welfare. Log-normal potential outcome distributions with $\mu_a = \mu_b = 0$.

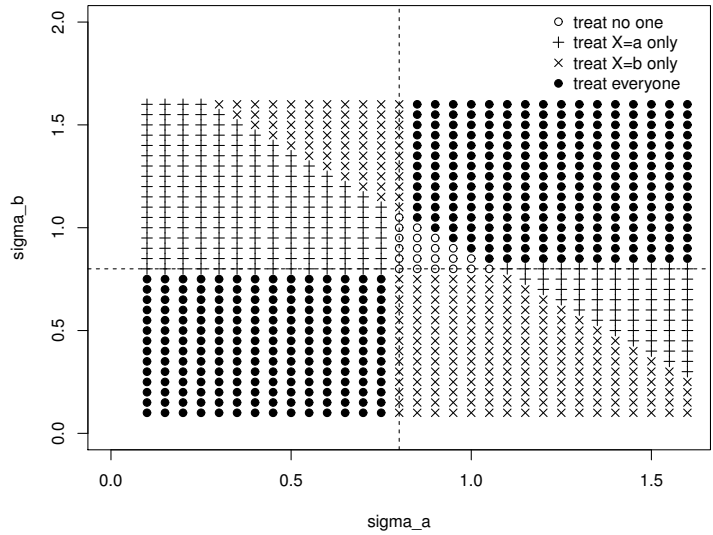


Figure 4: Optimal policies under the Gini welfare. Log-normal potential outcome distributions with $\mu_a = \mu_b = 0$.

welfare case (G_{util}^*) shown in Figure 3. In the neighborhood of $(\sigma_a, \sigma_b) = (0.8, 0.8)$, the subpopulations to be treated under the Gini social welfare are the converse of those to be treated under the utilitarian welfare. This is because the Gini social welfare is decreasing in σ in the neighborhood of $\sigma = 0.8$ (Figure 1), while the utilitarian welfare is monotonically increasing in σ . Another notable difference is that in contrast to the quadrant partition observed in the utilitarian welfare case, the partition in the equality-minded welfare case is more complex. Some treatment rules are optimal in disconnected regions, e.g., G_{ab} is optimal in the south-west and the north-east regions of the plot. Furthermore, the region in which G_a is optimal can border the region in which G_b is optimal. On the border between these regions, the policy maker chooses whether to treat $X = a$ only or $X = b$ only, rather than whether to additionally treat the other subpopulation.

3 EWM for Equality-Minded Welfare

We proceed to propose our method of estimating the treatment rule in finite samples and analyze its properties. Let the data be a size n random sample of $Z_i = (Y_i, D_i, X_i)$, where X_i refers to the observable pre-treatment covariates of individual i , $D_i \in \{0, 1\}$ is a binary indicator of the individual's experimental treatment, and $Y_i \in \mathbb{R}_+$ is her/his observed post-treatment outcome. The population from which the sample is drawn is characterized by P , a joint distribution of $(Y_{0i}, Y_{1i}, D_i, X_i)$, where Y_{0i} and Y_{1i} are the potential outcomes that would be observed if i 's treatment status were $D_i = 0$ and $D_i = 1$, respectively. We assume that the propensity score $e(X) = \Pr(D = 1|X = x)$ is known. Section 4.2 extends the analysis to include propensity score estimation. Based on this data, the policy-maker has to choose a conditional treatment rule $G \in \mathcal{G}$ that determines whether individuals with covariates X in a target population will be assigned to treatment 0 or to treatment 1. The following are our maintained assumptions about the class of population distributions of (Y_0, Y_1, D, X) :

Assumption 3.1.

(UCF) *Unconfoundedness*: $(Y_0, Y_1) \perp D|X$.

(BO) *Bounded Outcomes*: There exists $M < \infty$ such that the support of Y is contained in $[0, M]$.

(SO) *Strict Overlap*: There exist $\kappa \in (0, 1/2]$ such that the propensity score satisfies $e(x) \in [\kappa, 1 - \kappa]$ for all $x \in \mathcal{X}$.

These assumptions generally hold if the data come from an experiment with randomized treatment assignment. In observational studies, on the other hand, *Unconfoundedness* rules out situations in which the observed treatment assignments depend on subjects' unobserved characteristics that can be associated with their potential outcomes. *Strict Overlap* can be also violated in an observational study if only one of the treatments is assigned in the sampling process for some covariate values. We do not constrain any feature of the joint distribution of (Y_0, Y_1, X) except that Y_0 and Y_1 have a bounded support. The outcome variable and the covariates can be discrete, continuous, or their combination, and the support of X does not have to be bounded.

We estimate the treatment rule by maximizing a sample analog of the population SWF. The *equality-minded EWM treatment rule* \widehat{G} maximizes a sample analog $\widehat{W}(G)$ of the welfare criterion over the set of feasible rules $G \in \mathcal{G}$. The unknown outcome distribution F_G induced by treatment rule G in (2.7) could be estimated by

$$\widehat{F}_G(y) \equiv \frac{1}{n} \sum_{i=1}^n \left[\frac{D_i \cdot 1\{Y_i \leq y\}}{e(X_i)} \cdot 1\{X_i \in G\} + \frac{(1 - D_i) \cdot 1\{Y_i \leq y\}}{1 - e(X_i)} \cdot 1\{X_i \notin G\} \right]. \quad (3.1)$$

The sample analog of welfare (equation (2.1)) is defined as³

$$\widehat{W}(G) \equiv \int_0^M \Lambda(\widehat{F}_G(y) \wedge 1) dy. \quad (3.2)$$

The equality-minded EWM treatment rule is then

$$\widehat{G} \in \arg \max_{G \in \mathcal{G}} \widehat{W}(G). \quad (3.3)$$

The cdf sample analog $\widehat{F}_G(y)$ is increasing in y , but will not be a proper cdf when $\widehat{F}_G(M) \neq 1$. It could be rescaled to a proper cdf as follows:

$$\widehat{F}_G^R(y) \equiv \begin{cases} \widehat{F}_G(y) / \widehat{F}_G(\widehat{M}) & \text{if } \widehat{F}_G(\widehat{M}) > 0, \\ 1\{y \geq \widehat{M}\} & \text{if } \widehat{F}_G(\widehat{M}) = 0, \end{cases} \quad (3.4)$$

³The minimum (\wedge) of $\widehat{F}_G(y)$ and 1 is taken because $\widehat{F}_G(y)$ may take values greater than 1, for which $\Lambda(\cdot)$ is not defined.

where $\widehat{M} \equiv \max_{1 \leq i \leq n} Y_i$ is the largest outcome observed in the sample. We also consider properties of the equality-minded EWM rule using rescaled cdfs $\widehat{F}_G^R(y)$

$$\widehat{G}^R \equiv \arg \max_{G \in \mathcal{G}} \widehat{W}^R(G), \text{ where } \widehat{W}^R(G) \equiv \int_0^{\widehat{M}} \Lambda(\widehat{F}_G^R(y)) dy. \quad (3.5)$$

Rescaling also eliminates the need to specify the bound M ex ante.

3.1 Rate Optimality of EWM

The next theorem derives a uniform upper bound of the average welfare loss of the EWM rule.

Theorem 3.1. *Under Assumptions 2.1 and 2.2, the average welfare loss of treatment rules \widehat{G} and \widehat{G}^R satisfies*

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[W_G^* - W(\widehat{G}) \right] \leq C |\Lambda'(0)| \frac{M}{\kappa} \sqrt{\frac{v}{n}}, \quad (3.6)$$

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[W_G^* - W(\widehat{G}^R) \right] \leq C^R |\Lambda'(0)| \frac{M}{\kappa} \sqrt{\frac{v}{n}}, \quad (3.7)$$

where \mathcal{P} is the class of all distributions satisfying Assumption 3.1 and $C > 0$, $C^R > 0$ are universal constants.

Proof. See Appendix A.

This theorem shows that for a large class of data generating processes characterized by Assumption 3.1, the welfare loss of the EWM rule is guaranteed to converge to the maximal attainable welfare no slower than at $n^{-1/2}$ rate. The uniform convergence rate of $n^{-1/2}$ coincides with that of the EWM rule for the utilitarian welfare shown in Theorem 2.1 of Kitagawa and Tetenov (2015). This is a nontrivial result, given that the rank-dependent welfare function depends on the whole conditional distributions of the potential outcomes given covariates, rather than only on their conditional means, as is the case for the utilitarian welfare criterion.

The next proposition provides a universal lower bound for the worst-case average welfare loss of any treatment rule.

Theorem 3.2. *Suppose that Assumptions 2.1 and 2.2 hold with $v \geq 2$, then for any non-randomized treatment choice rule \widehat{G} that is a function of the sample, and for any $t^* \in (0, 1]$ at which $\Lambda(\cdot)$ is differentiable, it holds*

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[W_{\mathcal{G}}^* - W(\widehat{G}) \right] \geq \frac{e^{-4}}{2} M |\Lambda'(t^*)| \sqrt{t^*} \sqrt{\frac{v-1}{n}} \text{ for all } n \geq 4(v-1)/t^*, \quad (3.8)$$

where \mathcal{P} is the class of all distributions satisfying Assumption 3.1.

Proof. See Appendix A.

Since $\Lambda(\cdot)$ is convex and $|\Lambda'(0)| > 0$, there also exists some $t^* > 0$ for which $|\Lambda'(t^*)| > 0$. Hence the bound (3.8) is always positive for some $t^* > 0$. A comparison of the lower bound of this theorem with the welfare loss upper bound of the EWM rule obtained in Theorem 3.1 shows that the EWM rule is *minimax rate optimal* over the class of data generating processes satisfying Assumption 3.1. We can therefore claim that in the absence of any additional restrictions other than Assumption 3.1, no other data-driven procedure for obtaining a non-randomized rule can outperform the EWM rule in terms of the uniform convergence rate over \mathcal{P} . This optimality claim is analogous to that of the EWM rule for the utilitarian welfare case (Theorems 2.1 and 2.2 in Kitagawa and Tetenov (2015)), and the minimax optimal rate attained by the equality-EWM rule is the same as the optimal rate in the utilitarian welfare case. It is remarkable to see that even in the absence of any analytical characterization of the true optimal assignment rule in terms of the population distribution of (Y_0, Y_1, X) , maximizing the empirical welfare leads to a policy, by implementing which the social welfare is guaranteed to reach the attainable maximum welfare at the minimax optimal rate.

It is also worth noting that the VC-dimension of \mathcal{G} appears in the same order both in the upper and lower bound expressions of Theorems 3.1 and 3.2. Since these bounds are non-asymptotic, we can let v increase with the sample size, and we can still conclude the minimax rate optimality of the equality-EWM rule. This insight is similar to the EWM rule for the utilitarian welfare case (Remark 2.2 in Kitagawa and Tetenov (2015)).

4 Extensions

4.1 Social Welfare is Defined on a Population Larger than the Sampled Population

As illustrated in Section 2.1, one of the distinguishing features of the rank-dependent social welfare is that it is *not* additive over the multiple populations. This implies that if the population for which the policy intervention takes place (e.g., unemployed workers) is only a subset of the population on which the rank-dependent SWF is defined (e.g., population of a country), it is important to explicitly take into account the outcome distribution of the non-targeted subset of the population in estimating the optimal assignment rule.

Suppose that the social welfare function is defined on the population whose outcome distribution is given by a mixture of two subpopulations

$$J = \eta F + (1 - \eta)H, \quad \eta \in (0, 1), \quad (4.1)$$

where F is the outcome distribution from which the experimental data (Y, D, X) with size n are sampled and for which a treatment assignment policy is estimated and to be implemented. H is the outcome distribution of those that were excluded from the sampling process and are out of scope of the treatment assignment policy intervention. The mixture weight η represents the size of subpopulation F . For simplicity, we assume that η and H are known to the social planner. We also assume that the outcome distribution H is invariant to a treatment assignment policy applied to the subpopulation F , e.g., no spill-over or general equilibrium effect across F and H .

Suppose that the social welfare is defined on J . Hence, implementing treatment assignment rule $\{X \in G\}$ on subpopulation F attains

$$W(J_G) \equiv \int_0^\infty \Lambda(\eta F_G(y) + (1 - \eta)H(y))dy,$$

where $F_G(\cdot)$ is the cdf defined in (2.7). The empirical welfare maximization method in the current case is set up as maximizing a sample analog of $W(J_G)$,

$$\hat{G} \in \arg \max_{G \in \mathcal{G}} W(\eta \hat{F}_G(y) + (1 - \eta)H(y)),$$

where $\widehat{F}_G(y)$ is as defined in (3.1).

The uniform convergence property of Theorem 3.1 carries over to the current case except for a minor change in the constant term, as shown in the next corollary. Its proof can be obtained similarly to Theorem 3.1.

Corollary 4.1. *Under Assumptions 2.1, 2.2 and 3.1,*

$$\sup_{P \in \mathcal{P}} E_{P^n} \left[\sup_{G \in \mathcal{G}} W(J_G) - W(J_{\widehat{G}}) \right] \leq C |\Lambda'(0)| \eta \frac{M}{\kappa} \sqrt{\frac{v}{n}}, \quad (4.2)$$

where $C > 0$ is a universal constant defined in Theorem 3.1.

4.2 EWM with Estimated Propensity Score

Unknown propensity score is common in observational studies. This section considers that the equality-minded EWM approach with estimated propensity scores and investigates the influence of the lack of knowledge on propensity scores to the uniform convergence rate of the welfare loss criterion.

Let $\hat{e}(x)$ be an estimator for the propensity score $\Pr(D = 1|X = x)$. The empirical welfare criterion of assignment policy $\{X \in G\}$ with the estimated propensity scores plugged in is given by

$$W(\widehat{F}_G^e) = \int_0^M \Lambda(\widehat{F}_G^e(y) \wedge 1) dy,$$

$$\widehat{F}_G^e(y) \equiv \frac{1}{n} \sum_{i=1}^n \left[\frac{D_i 1\{Y_i \leq y\}}{\hat{e}(X_i)} \cdot 1\{X_i \in G\} + \frac{(1 - D_i) 1\{Y_i \leq y\}}{1 - \hat{e}(X_i)} \cdot 1\{X_i \notin G\} \right].$$

The equality-EWM rule with estimated propensity score is defined accordingly as

$$\widehat{G}^e \in \arg \max_{G \in \mathcal{G}} W(\widehat{F}_G^e).$$

To characterize the uniform convergence rate of the welfare loss of \widehat{G}^e , we first assume that $\hat{e}(\cdot)$ is uniformly consistent to the true propensity score $e(\cdot)$ in the following sense.

Assumption 4.1. For a class of data generating processes \mathcal{P}_e , there exists a sequence

$\phi_n \rightarrow \infty$ such that

$$\begin{aligned} \limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_e} \phi_n E_{P^n} \left[\frac{1}{n} \sum_{i=1}^n \left| 1 - \frac{e(X_i)}{\hat{e}(X_i)} \right| \right] &< \infty \quad \text{and} \\ \limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_e} \phi_n E_{P^n} \left[\frac{1}{n} \sum_{i=1}^n \left| 1 - \frac{1 - e(X_i)}{1 - \hat{e}(X_i)} \right| \right] &< \infty \end{aligned} \quad (4.3)$$

hold.

When the class of data generating processes \mathcal{P}_e constrains the propensity score to a parametric family with compact support of X , parametric estimator $\hat{e}(X_i)$ satisfies this assumption with $\phi_n = n^{1/2}$. When the propensity scores are estimated nonparametrically instead, ϕ_n is generally slower than $n^{1/2}$. The rate of ϕ_n for nonparametrically estimated propensity scores depends on the smoothness of $e(\cdot)$ and the dimension of X , as we discuss further below.

Theorem 4.1. *Suppose Assumptions 2.1, 2.2 and 3.1 hold. For a class of data generating process \mathcal{P}_e , if an estimator for the propensity score satisfies Assumption 4.1, then*

$$\sup_{P \in \mathcal{P}_e \cap \mathcal{P}} E_{P^n} [W_{\hat{G}}^* - W(F_{\hat{G}^e})] \leq O \left(\phi_n^{-1} \vee \sqrt{\frac{v}{n}} \right). \quad (4.4)$$

Proof. See Appendix A.

This theorem extends Theorem 2.5 (e) of Kitagawa and Tetenov (2015) to the rank-dependent social welfare case. The shown uniform convergence rate implies that the parametrically estimated propensity score achieving $\phi_n = n^{1/2}$ does not affect the convergence rate property of the welfare loss. With nonparametrically estimated propensity score, on the other hand, the uniform welfare loss convergence rate can be slower than the one with the known propensity score obtained in Theorem 3.1. For instance, if $\hat{e}(X_i)$ is estimated by local polynomial regression (with proper trimming), then for a suitably defined \mathcal{P}_e , we have $\phi_n = n^{\frac{1}{2+d_x/\beta_e}}$, where $\beta_e \geq 1$ is the parameter constraining smoothness of $e(\cdot)$ in terms of the degree of Hölder class of functions and $d_x \geq 1$ is the dimension of X . Since $\frac{1}{2+d_x/\beta_e} < \frac{1}{2}$, the upper bound of the uniform convergence rate shown in Theorem 4.1 implies

$$\sup_{P \in \mathcal{P}_e \cap \mathcal{P}} E_{P^n} [W_{\hat{G}}^* - W(F_{\hat{G}^e})] \leq O \left(n^{-\frac{1}{2+d_x/\beta_e}} \right), \quad (4.5)$$

as far as the VC-dimension of \mathcal{G} is either constant or does not grow too fast as the sample size. For a formal derivation of (4.5) and the precise construction of the local polynomial estimator for $e(\cdot)$, see Appendix B.

4.3 Cost of Treatment and Capacity Constraint

The analyses in the preceding sections ignore the cost of treatment although most of the programs in reality are not cost-free. In this section, we discuss how to take into account the cost of treatment in the estimation of welfare maximizing treatment assignment policy. For tractability, we restrict the search of an optimal policy to non-randomized ones.

Let $0 \leq c(x) < \infty$, $x \in \mathcal{X}$, be the cost of treatment 1 for a subject whose observable characteristics are x . We assume that treatment 0 is cost-free and $c(\cdot)$ is known. Under the utilitarian welfare, it is simple to construct the social welfare criterion net of the cost of treatment by simply subtracting the per-capita cost of treatment $C(G) \equiv \int_G c(x) dP_X(x)$ from the mean of outcomes. Thanks to the additivity of the utilitarian welfare, the optimal policy subject to the treatment cost is invariant to assumptions about *who pays the cost*. In contrast, for the case of rank-dependent social welfare, this invariance no longer holds; we have to be explicit about who bears the cost in the construction of the social welfare criterion.

To illustrate, consider a scenario where the outcome variable is income and the cost of treatment is self-financed by each recipient of the treatment, i.e., each recipient of treatment pays own cost without a borrowing constraint. In this case, by transforming one's potential outcomes to $\tilde{Y}_{1i} \equiv Y_{1i} + \bar{c} - c(X_i)$ and $\tilde{Y}_{0i} = Y_{0i} + \bar{c}$, where $\bar{c} = \sup_{x \in \mathcal{X}} c(x)$,⁴ we can write the rank-dependent SWF of policy G as

$$W^{sf}(G) = \int_0^\infty \Lambda(F_G^{sf}(y)) dy$$

$$F_G^{sf}(y) \equiv \int_{\mathcal{X}} \left[F_{\tilde{Y}_0|X=x}(y) 1\{x \notin G\} + F_{\tilde{Y}_1|X=x}(y) 1\{x \in G\} \right] dP_X(x), \quad (4.6)$$

⁴To be consistent with our Assumption 3.1 (BO), we add \bar{c} to the potential outcomes to ensure that the support of \tilde{Y}_{1i} to be contained in $[0, M + \bar{c}]$. Note that a uniform addition of a constant does not affect the welfare ranking of the policies.

where $F_{\tilde{Y}_0|X=x}(\cdot)$ and $F_{\tilde{Y}_1|X=x}(\cdot)$ are the cdfs of the transformed potential outcomes. An empirical welfare for $W^{sf}(G)$ can be obtained by replacing $\widehat{F}_G(y)$ in (3.2) by

$$\widehat{F}_G^{sf}(y) \equiv \frac{1}{n} \sum_{i=1}^n \left[\frac{D_i \cdot \mathbf{1}\{Y_i + \bar{c} - c(X_i) \leq y\}}{e(X_i)} \cdot \mathbf{1}\{X_i \in G\} + \frac{(1 - D_i) \cdot \mathbf{1}\{Y_i + \bar{c} \leq y\}}{1 - e(X_i)} \cdot \mathbf{1}\{X_i \notin G\} \right],$$

Since this modification does not affect validity of Assumption 3.1, the EWM rule with the self-financed treatment cost attains the uniform welfare loss upper bounds of Theorem 3.1 with $M + \bar{c}$ in place of M .

Consider an alternative scenario where the treatment cost is financed by all the population members equally in a lump-sum manner, i.e., per-capita treatment cost $C(G)$ is subtracted from every individual's income independent of his assigned treatment. Using the representation of the rank-dependent SWF given in (2.2), the rank-dependent social welfare under the lump-sum treatment cost can be expressed as

$$W^{ls}(G) \equiv \int_0^1 [F_G^{-1}(t) + \bar{c} - C(G)] \omega(t) = W(G) + \bar{c} - C(G), \quad (4.7)$$

where we use $\int_0^1 \omega(t) dt = \Lambda(0) - \Lambda(1) = 1$.⁵ An empirical welfare in this case is simple to construct, $\widehat{W}^{ls}(G) \equiv \widehat{W}(G) + \bar{c} - \frac{1}{n} \sum_{i=1}^n c(X_i) \cdot \mathbf{1}\{X_i \in G\}$, and the EWM rule is obtained by maximizing $\widehat{W}^{ls}(\cdot)$ over $G \in \mathcal{G}$.

As a comparison of $W^{sf}(G)$ and $W^{ls}(G)$ shows, the SWFs with treatment cost can differ depending on the assumption about how the cost of treatment is financed by the population. This implies that with the rank-dependent SWF, not only a treatment assignment policy but also an *allocation of the cost of treatment* can be a vehicle of policy intervention with which the social planner can affect the social welfare. In this paper, we do not consider the joint optimization of the treatment assignment and cost allocation.

Another practical concern ruled out in the preceding sections is a capacity constraint, which limits the proportion of the population that can be assigned to treatment. Suppose that the proportion of the target population that could receive treatment 1 cannot exceed $K \in (0, 1)$. When P_X is unknown, Kitagawa and Tetenov (2015) propose a capacity constrained EWM procedure for the utilitarian welfare case by assuming that when treatment

⁵Again, adding \bar{c} to all the quantiles is to ensure that the support of the outcome distribution is contained in $[0, M + \bar{c}]$.

rule G violates the capacity constraint, $P_X(G) > K$, the scarce treatment is randomly rationed to a fraction $K/P_X(G)$ of the population of $\{X \in G\}$ independently of (Y_0, Y_1, X) . This random-rationing approach can be straightforwardly extended to the EWM for the rank-dependent social welfare.

With the capacity constraint and random rationing, the cdf of outcomes realized at an assignment policy G can be written as

$$F_G^K(y) = \int_{\mathcal{X}} \left[F_{Y_0|X=x}(y)1\{x \notin G\} + \left(1 - \min\left\{1, \frac{K}{P_X(G)}\right\}\right) F_{Y_0|X=x}(y)1\{x \in G\} \right. \\ \left. + \min\left\{1, \frac{K}{P_X(G)}\right\} F_{Y_1|X=x}(y)1\{x \in G\} \right] dP_X(x).$$

We hence obtain the social welfare under the capacity constraint and random rationing as $W^K(G) \equiv \int_0^1 \Lambda(F_G^K(y))dy$. Its sample analog can be constructed by replacing $\widehat{F}_G^K(y)$ in (3.2) with

$$\widehat{F}_G^K(y) = \frac{1}{n} \sum_{i=1}^n \left[\frac{(1-D_i)1\{Y_i \leq y\}}{1-e(X_i)} \cdot 1\{X_i \notin G\} + \left(1 - \min\left\{1, \frac{K}{P_{X,n}(G)}\right\}\right) \frac{(1-D_i)1\{Y_i \leq y\}}{1-e(X_i)} \cdot 1\{X_i \in G\} \right. \\ \left. + \min\left\{1, \frac{K}{P_{X,n}(G)}\right\} \frac{D_i 1\{Y_i \leq y\}}{e(X_i)} \cdot 1\{X_i \in G\} \right],$$

where $P_{X,n}(G)$ is a sample analog of $P_X(G)$. By combining the proof of Theorem 3.1 and the proof of Theorem 4.1 in Kitagawa and Tetenov (2015), we can show that the resulting EWM rule converges to the maximal attainable welfare at $n^{-1/2}$ rate.

5 Empirical Illustration

To illustrate equality-minded empirical treatment choice, we apply our method to the experimental data from the National Job Training Partnership Act (JTPA) Study.⁶ A detailed description of the study and an assessment of average program effects for five large subgroups of the target population is found in Bloom et al. (1997). The study randomized whether applicants were eligible to receive a mix of training, job-search assistance, and other services provided by the JTPA for a period of 18 months. It collected background information about the applicants prior to random assignment, as well as administrative and survey data on applicants' earnings in the 30-month period following assignment. We use the same sample

⁶We use the data provided in the Joshua Angrist's Data Archive at <http://economics.mit.edu/faculty/angrist/data1/data/abangim02>

Table 1: Gini welfare and average income under selected treatment rules that condition on education and pre-program earnings.

Treatment rule:	Representative income for Gini welfare function	Average income	Share of population assigned to treatment
<i>Simple treatment rules</i>			
Treat nobody	\$6,610	\$14,969	0
Treat everyone	\$7,310	\$16,238	1
<i>Quadrant class conditioning on years of education and pre-program earnings</i>			
Maximize average income	\$7,406	\$16,583	0.8266
Maximize Gini welfare	\$7,440	\$15,953	0.7266
<i>Linear class conditioning on years of education and pre-program earnings</i>			
Maximize average income	\$7,411	\$16,626	0.8197
Maximize Gini welfare	\$7,478	\$16,063	0.7760

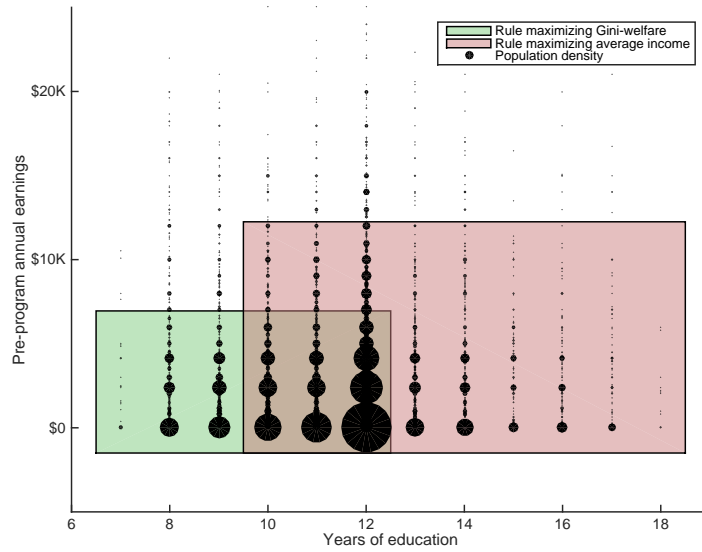


Figure 5: Treatment rules that maximize Gini welfare function and average income within the quadrant class conditioning on years of education and pre-program earnings

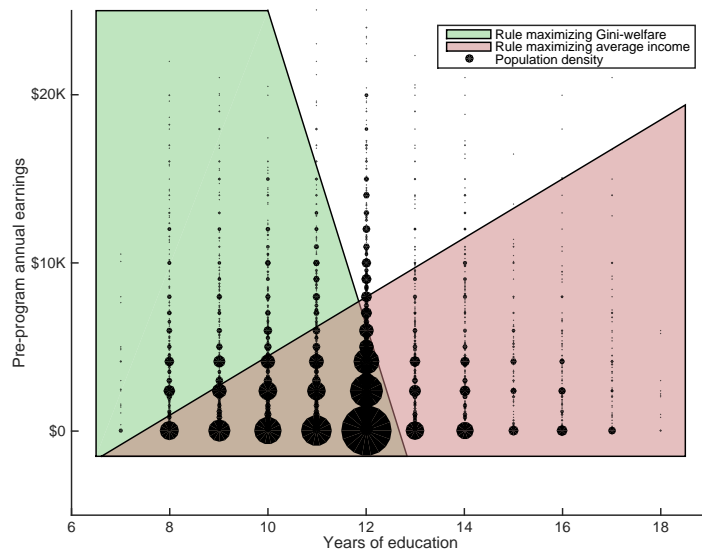


Figure 6: Treatment rules that maximize Gini welfare function and average income within the linear class conditioning on years of education and pre-program earnings

of 11204 adults (22 years and older) used in the original evaluation of the program and many subsequent studies (Bloom et al., 1997, Heckman et al., 1997a, Abadie et al., 2002). The probability of being assigned to treatment was two thirds in this sample.

For this illustration, total individual earnings in the 30-month period following program assignment serve as the measure of income. We evaluate the Gini social welfare function as an example of an equality-minded criterion. We use average income as an example of a utilitarian criterion. For simplicity, we consider only the distribution of earnings in the population sampled for the experiment in the social welfare function. In practice, the targeted population is only part of the total population. The social welfare function should be evaluated on the full income distribution that combines both targeted and non-targeted individuals.

Pre-treatment variables on which we consider conditioning treatment assignment are the individual's years of education and earnings in the year prior to assignment. We do not use race, sex, or age to condition treatment assignment. Though treatment effects may vary with these characteristics, using them to condition treatment assignment is often socially unacceptable and illegal. Education and earnings are verifiable characteristics, which is also important for conditioning treatment assignment. The performance of treatment rules that condition on unverifiable characteristics is hard to evaluate if individuals change their self-reported characteristics to obtain their desired treatment assignment.

Table 1 compares empirical estimates of social welfare measures for a few alternative treatment rules. First, we consider simple treatment rules that either assign nobody or everybody to treatment. Second, we consider empirically optimal rules from the class of *quadrant treatment rules*:

$$\mathcal{G}_Q \equiv \left\{ \begin{array}{l} \{x : s_1(\text{education} - t_1) > 0 \ \& \ s_2(\text{prior earnings} - t_2) > 0\}, \\ s_1, s_2 \in \{-1, 0, 1\}, t_1, t_2 \in \mathbb{R} \end{array} \right\}. \quad (5.1)$$

This class of treatment eligibility rules is easily implementable and is often used in practice. To be assigned to treatment according to such rules, an individual's education and pre-program earnings both have to be above (or below) some specific thresholds. Third, we consider empirically optimal rules from the class of *linear treatment rules*:

$$\mathcal{G}_{LES} \equiv \{\{x : \beta_0 + \beta_1 \cdot \text{education} + \beta_2 \cdot \text{prior earnings} > 0\}, \beta_0, \beta_1, \beta_2 \in \mathbb{R}\}. \quad (5.2)$$

The first column evaluates the Gini social welfare function, expressed in terms of the *representative income* of the policy. The income distribution generated by the policy is valued as much as an equal income distribution with the representative income. The second column lists the average income. The third column lists the proportion of the target population assigned to treatment by each policy.

Figure 5 compares the quadrant treatment rules maximizing the sample analogs of the Gini SWF and of the average income. Figure 6 compares the linear treatment rules maximizing these criteria. The size of black dots shows the number of individuals with different covariate values. Many individuals would be assigned to treatment by both rules, but there are also notable differences. Both quadrant and linear treatment rules maximizing the Gini SWF are focused on assigning low income, low education individuals to treatment. Treatment rules maximizing the average income instead favor low income, but more educated individuals. In both cases, treatment rules maximizing the Gini SWF would assign the program treatment to a lower proportion of the target population.

6 Conclusion

This paper develops an empirical welfare maximization method for treatment choice when the policy maker's objective is to maximize a rank-dependent SWF. We showed that we can learn the optimal feasible assignment policy under a rank-dependent social welfare at the same uniform convergence rate as we can for the optimal utilitarian assignment rule. The key restriction underlying these rate results is the complexity restriction (Assumption 2.2) imposed on the set of feasible policies. This complexity restriction should be seen as an attractive feature rather than a disadvantage of the method, as it offers a flexible and convenient way to incorporate the exogenous constraints that the policy maker faces in realistic settings of policy design. Our analytical results cover a general class of equality-minded rank-dependent SWFs. Computing the EWM rule is more challenging than in the utilitarian welfare case. Efficient computation for the equality-minded EWM rule remains an open question.

A Appendix: Lemmas and Proofs

Proof of Theorem 2.1. Denote an upper level set of $\delta(x)$ at level $u \in [0, 1]$ by $G(u) \equiv \{x \in \mathcal{X} : \delta(x) \geq u\}$. By noting

$$\delta(x) = \int_0^1 1\{x \in G(u)\} du,$$

we can rewrite $F_\delta(y)$ defined in (2.4) as

$$\begin{aligned} F_\delta(y) &= \int_{\mathcal{X}} \left[\int_0^1 1\{x \notin G(u)\} du \cdot F_{Y_0|X=x}(y) + \int_0^1 1\{x \in G(u)\} du \cdot F_{Y_1|X=x}(y) \right] dP_X(x) \\ &= \int_0^1 \left[\int_{\mathcal{X}} (1\{x \notin G(u)\} \cdot F_{Y_0|X=x}(y) + 1\{x \in G(u)\} \cdot F_{Y_1|X=x}(y)) dP_X(x) \right] du \\ &= \int_0^1 F_{G(u)}(y) du, \end{aligned}$$

where $F_{G(u)}(y)$ is the distribution of outcomes induced by treatment rule $\delta_{G(u)} \equiv 1\{x \in G(u)\}$. By convexity of $\Lambda(\cdot)$, we obtain

$$\Lambda(F_\delta(y)) \leq \int_0^1 \Lambda(F_{G(u)}(y)) du,$$

and this leads to

$$W(F_\delta) \leq \int_0^1 W(F_{G(u)}) du \equiv \bar{W}. \quad (\text{A.1})$$

Suppose that $\bar{W} - W(F_{G(u)}) > 0$ for all $u \in [0, 1]$. Then the integral of this function over the set $u \in [0, 1]$ of positive measure must also be strictly positive,

$$0 < \int_0^1 (\bar{W} - W(F_{G(u)})) du = \bar{W} - \bar{W},$$

which is a contradiction. Therefore, there exists $u^* \in [0, 1]$ for which $W(F_{G(u^*)}) \geq \bar{W}$, hence $W(F_{G(u^*)}) \geq W(F_\delta)$. If all upper level sets $G(u)$ of δ belong to \mathcal{G} , then also $G(u^*) \in \mathcal{G}$. □

Proof of Theorem 3.1. Take an arbitrary set $G^* \in \mathcal{G}$, then

$$\begin{aligned} W(G^*) - W(\hat{G}) &= W(G^*) - \widehat{W}(\hat{G}) + \widehat{W}(\hat{G}) - W(\hat{G}) \\ &\leq W(G^*) - \widehat{W}(G^*) + \widehat{W}(\hat{G}) - W(\hat{G}) \\ &\leq 2 \sup_{G \in \mathcal{G}} \left| \widehat{W}(G) - W(G) \right|, \end{aligned}$$

where the second line follows since $\widehat{W}(\widehat{G})$ maximizes $\widehat{W}(G)$ over $G \in \mathcal{G}$. Since $W_{\mathcal{G}}^* = \max_{G \in \mathcal{G}} W(G)$, it follows that

$$W_{\mathcal{G}}^* - W(\widehat{G}) \leq 2 \sup_{G \in \mathcal{G}} \left| \widehat{W}(G) - W(G) \right|. \quad (\text{A.2})$$

Since $\Lambda(\cdot)$ is convex and non-increasing,

$$\begin{aligned} \sup_{G \in \mathcal{G}} \left| \widehat{W}(G) - W(G) \right| &= \sup_{G \in \mathcal{G}} \left| \int_0^M \Lambda(\widehat{F}_G(y)) dy - \int_0^M \Lambda(F_G(y)) dy \right| \\ &\leq \sup_{G \in \mathcal{G}} \int_0^M \left| \Lambda(\widehat{F}_G(y)) - \Lambda(F_G(y)) \right| dy \\ &\leq \int_0^M \sup_{G \in \mathcal{G}} \left| \Lambda(\widehat{F}_G(y)) - \Lambda(F_G(y)) \right| dy \\ &\leq |\Lambda'(0)| \int_0^M \sup_{G \in \mathcal{G}} \left| \widehat{F}_G(y) - F_G(y) \right| dy. \end{aligned} \quad (\text{A.3})$$

Combining (A.2) and (A.3), the average welfare loss of \widehat{G} can be bounded by

$$E_{P^n} \left[W_{\mathcal{G}}^* - W(\widehat{G}) \right] \leq 2 |\Lambda'(0)| \int_0^M E_{P^n} \left[\sup_{G \in \mathcal{G}} \left| \widehat{F}_G(y) - F_G(y) \right| \right] dy. \quad (\text{A.4})$$

Now consider treatment rule $\widehat{G}^R \equiv \arg \max_{G \in \mathcal{G}} \widehat{W}^R(G)$ that maximizes empirical welfare computed from rescaled cdfs $\widehat{W}^R(G) \equiv \int_0^{\widehat{M}} \Lambda(\widehat{F}_G^R(y)) dy$. The argument used to obtain equation (A.2) also applies, thus

$$W_{\mathcal{G}}^* - W(\widehat{G}^R) \leq 2 \sup_{G \in \mathcal{G}} \left| \widehat{W}^R(G) - W(G) \right|. \quad (\text{A.5})$$

An analog of (A.3) also holds for $\widehat{W}^R(G)$ (note that $\widehat{F}_G^R(y) = 1$ for all $y \geq \widehat{M}$ and $\Lambda(1) = 0$):

$$\begin{aligned} \sup_{G \in \mathcal{G}} \left| \widehat{W}^R(G) - W(G) \right| &= \sup_{G \in \mathcal{G}} \left| \int_0^{\widehat{M}} \Lambda(\widehat{F}_G^R(y)) dy - \int_0^M \Lambda(F_G(y)) dy \right| \\ &= \sup_{G \in \mathcal{G}} \left| \int_0^{\widehat{M}} \Lambda(\widehat{F}_G^R(y)) dy + \int_{\widehat{M}}^M \Lambda(\widehat{F}_G^R(y)) dy - \int_0^M \Lambda(F_G(y)) dy \right| \\ &= \sup_{G \in \mathcal{G}} \left| \int_0^M \Lambda(\widehat{F}_G^R(y)) dy - \int_0^M \Lambda(F_G(y)) dy \right| \\ &\leq \sup_{G \in \mathcal{G}} \int_0^M \left| \Lambda(\widehat{F}_G^R(y)) - \Lambda(F_G(y)) \right| dy \\ &\leq |\Lambda'(0)| \int_0^M \sup_{G \in \mathcal{G}} \left| \widehat{F}_G^R(y) - F_G(y) \right| dy. \end{aligned} \quad (\text{A.6})$$

Combining (A.5) and (A.6), the average welfare loss of \widehat{G}^R can be bounded by

$$E_{P^n} \left[W_G^* - W(\widehat{G}^R) \right] \leq 2|\Lambda'(0)| \int_0^M E_{P^n} \left[\sup_{G \in \mathcal{G}} \left| \widehat{F}_G^R(y) - F_G(y) \right| \right] dy. \quad (\text{A.7})$$

We will show that for any G and $y \leq M$,

$$\left| \widehat{F}_G^R(y) - F_G(y) \right| \leq \left| \widehat{F}_G(y) - F_G(y) \right| + \left| \widehat{F}_G(M) - F_G(M) \right|. \quad (\text{A.8})$$

First, we consider the case when $\widehat{F}_G(\widehat{M}) > 0$. It follows from the definition (3.4) that

$$\left| \widehat{F}_G(y) - \widehat{F}_G^R(y) \right| = \frac{\widehat{F}_G(y)}{\widehat{F}_G(\widehat{M})} \left| \widehat{F}_G(\widehat{M}) - 1 \right| \leq \left| \widehat{F}_G(M) - F_G(M) \right|$$

since $\widehat{F}_G(y) \leq \widehat{F}_G(\widehat{M}) = \widehat{F}_G(M)$ and $F_G(M) = 1$. Then

$$\begin{aligned} \left| \widehat{F}_G^R(y) - F_G(y) \right| &\leq \left| \widehat{F}_G(y) - F_G(y) \right| + \left| \widehat{F}_G(y) - \widehat{F}_G^R(y) \right| \\ &\leq \left| \widehat{F}_G(y) - F_G(y) \right| + \left| \widehat{F}_G(M) - F_G(M) \right|. \end{aligned}$$

If, instead, $\widehat{F}_G(\widehat{M}) = 0$, it implies that $\widehat{F}_G(y) = 0$ for all y , including $\widehat{F}_G(M) = 0$. Since $\widehat{F}_G^R(y) \in [0, 1]$ and $F_G(y) \in [0, 1]$, and $F_G(M) = 1$ by Assumption 3.1(BO),

$$\left| \widehat{F}_G^R(y) - F_G(y) \right| \leq 1 = \left| \widehat{F}_G(M) - F_G(M) \right|.$$

Hence inequality (A.8) also holds when $\widehat{F}_G(\widehat{M}) = 0$.

By combining (A.7), and (A.8), the average welfare loss of the rescaled EWM treatment rule \widehat{G}^R can be bounded by

$$\begin{aligned} E_{P^n} \left[W_G^* - W(\widehat{G}^R) \right] &\leq 2|\Lambda'(0)| \int_0^M E_{P^n} \left[\sup_{G \in \mathcal{G}} \left| \widehat{F}_G(y) - F_G(y) \right| \right] dy \\ &\quad + 2|\Lambda'(0)| \cdot M \cdot E_{P^n} \left[\sup_{G \in \mathcal{G}} \left| \widehat{F}_G(M) - F_G(M) \right| \right]. \end{aligned} \quad (\text{A.9})$$

To bound (A.4) and (A.9), we derive a distribution-free upper bound for

$$E_{P^n} \left[\sup_{G \in \mathcal{G}} \left| \widehat{F}_G(y) - F_G(y) \right| \right].$$

Denote the observations by $Z_i = (Y_i, D_i, X_i)$, $i = 1, \dots, n$ and the sample average with respect to them by $E_n(\cdot)$. To apply a maximal inequality for the centered empirical process

to $E_{P^n} \left[\sup_{G \in \mathcal{G}} \left| \widehat{F}_G(y) - F_G(y) \right| \right]$, note that with y fixed, it can be written as

$$E_{P^n} \left[\sup_{f \in \mathcal{F}} |E_n(f) - E_P(f)| \right],$$

where

$$\mathcal{F} \equiv \left\{ f(Z_i) = \frac{D_i \cdot 1\{Y_i \leq y\}}{e(X_i)} \cdot 1\{X_i \in G\} + \frac{(1 - D_i) \cdot 1\{Y_i \leq y\}}{1 - e(X_i)} \cdot 1\{X_i \notin G\}, G \in \mathcal{G} \right\}.$$

\mathcal{F} is a VC-subgraph class of functions with envelope κ^{-1} and VC-dimension less than or equal to v by Lemma A.1 of Kitagawa and Tetenov (2015). Accordingly, by applying the maximal inequality for the centered empirical processes over the VC-subgraph class of functions in the form given in Lemma A.4 of Kitagawa and Tetenov (2015), we obtain

$$E_{P^n} \left[\sup_{G \in \mathcal{G}} \left| \widehat{F}_G(y) - F_G(y) \right| \right] \leq C_1 \kappa^{-1} \sqrt{\frac{v}{n}}, \quad (\text{A.10})$$

where C_1 is a universal constant.

Plugging inequality (A.10) into (A.4) yields

$$E_{P^n} \left[W_{\widehat{G}}^* - W(\widehat{G}) \right] \leq 2C_1 |\Lambda'(0)| \frac{M}{\kappa} \sqrt{\frac{v}{n}}.$$

Since this upper bound does not depend on $P \in \mathcal{P}$, this completes the proof of (3.6).

Similarly, plugging (A.10) into (A.9) completes the proof of (3.7)

$$E_{P^n} \left[W_{\widehat{G}^R}^* - W(\widehat{G}^R) \right] \leq 4C_1 |\Lambda'(0)| \frac{M}{\kappa} \sqrt{\frac{v}{n}}.$$

□

Proof of Theorem 3.2. We consider a suitable subclass $\mathcal{P}^* \subset \mathcal{P}$, for which the worst case welfare loss can be bounded from below by a distribution-free term that converges at rate $n^{-1/2}$. To simplify the proof, we restrict the range of outcomes to $Y \in [0, 1]$. To rescale to $Y \in [0, M]$, we multiply M to the lower bound as the rank-dependent welfare is linear in multiplicative constant M .

The construction of \mathcal{P}^* proceeds as follows. By the definition of VC-dimension, there exists a set of v points in \mathcal{X} , denoted $x_1, \dots, x_v \in \mathcal{X}$ that are shattered by \mathcal{G} . We constrain the marginal distribution of X to being supported only on (x_1, \dots, x_v) . Let $t^* \in (0, 1]$ stated in the current theorem be given. We put mass $p \equiv \frac{t^*}{v-1}$ at $x_i, i < v$, and mass $1 - t^*$ at x_v . Thus-constructed marginal distribution of X is common in \mathcal{P}^* . Let the distribution of treatment indicator D be independent of (Y_0, Y_1, X) , and let D follow the Bernoulli distribution with $\Pr(D = 1) = 1/2$. Let $\mathbf{b} = (b_1, \dots, b_{v-1}) \in \{0, 1\}^{v-1}$ be a bit vector used to index a

member of \mathcal{P}^* , i.e., $\mathcal{P}^* = \{P_{\mathbf{b}} : \mathbf{b} \in \{0, 1\}^{v-1}\}$ consists of a finite number of DGPs. For each $j = 1, \dots, (v - 1)$, and depending on \mathbf{b} , construct the following conditional distributions of potential outcomes given $X = x_j$; if $b_j = 1$,

$$Y_0|X = x_j \sim Ber\left(\frac{1 - \gamma}{2}\right), \quad Y_1|X = x_j \sim Ber\left(\frac{1 + \gamma}{2}\right), \quad (\text{A.11})$$

and, if $b_j = 0$,

$$Y_0|X = x_j \sim Ber\left(\frac{1 + \gamma}{2}\right), \quad Y_1|X = x_j \sim Ber\left(\frac{1 - \gamma}{2}\right), \quad (\text{A.12})$$

where $Ber(m)$ denotes the Bernoulli distribution with mean m and $\gamma \in (0, 1)$ is chosen properly in a later step of the proof. For $j = v$, we set the distribution of potential outcomes to be degenerate at the maximum value of Y , $P(Y_0 = Y_1 = 1|X = x_v) = 1$. Clearly, $P_{\mathbf{b}} \in \mathcal{P}$ for every $\mathbf{b} \in \{0, 1\}^{v-1}$. We accordingly define $\mathcal{P}^* = \{P_{\mathbf{b}} : \mathbf{b} \in \{0, 1\}^{v-1}\} \subset \mathcal{P}$.

Note that when the outcome distribution is Bernoulli with mean μ , the equality-minded welfare function equals $W(F) = \Lambda(1 - \mu)$, which is a non-decreasing function of μ . Hence, given knowledge of $P_{\mathbf{b}}$, an optimal treatment assignment rule for the equality-minded welfare coincides with that for the utilitarian welfare case,

$$G_{\mathbf{b}}^* = \{x_j : j < v, b_j = 1\},$$

which is feasible, since $G_{\mathbf{b}}^* \in \mathcal{G}$ by the construction of the support points of X . The maximized social welfare is accordingly obtained as

$$W(G_{\mathbf{b}}^*) = \Lambda(1 - \mu^*),$$

$$\mu^* \equiv p(v - 1) \left(\frac{1 + \gamma}{2}\right) + (1 - t^*) = t^* \left(\frac{1 + \gamma}{2}\right) + (1 - t^*),$$

which does not depend on \mathbf{b} .

Let \widehat{G} be an arbitrary treatment choice rule as a function of observations $Z_i \equiv (Y_i, D_i, X_i)$, $i = 1, \dots, n$, and $\widehat{\mathbf{b}} \in \{0, 1\}^{(v-1)}$ be a binary vector whose j -th element is $\widehat{b}_j = 1\{x_j \in \widehat{G}\}$. Let $\mu_{\widehat{G}}$ be the mean of outcome Y when treatment assignment rule \widehat{G} is implemented for a given realization of the sample. For $P \in \mathcal{P}^*$ outcomes are binary, hence

$$\mu_{\widehat{G}} \equiv \int_{\widehat{G}} \Pr(Y_1 = 1|X = x)dP_X(x) + \int_{\widehat{G}^c} \Pr(Y_0 = 1|X = x)dP_X(x).$$

Consider $\pi(\mathbf{b})$, a prior distribution for \mathbf{b} , such that b_1, \dots, b_{v-1} are iid and $b_1 \sim \text{Ber}(1/2)$. The welfare loss satisfies the following inequalities:

$$\begin{aligned}
\sup_{P \in \mathcal{P}} E_{P^n} [W_G^* - W(\widehat{G})] &\geq \sup_{P_{\mathbf{b}} \in \mathcal{P}^*} E_{P_{\mathbf{b}}^n} [W(G_{\mathbf{b}}^*) - W(\widehat{G})] \\
&\geq \int_{\mathbf{b}} E_{P_{\mathbf{b}}^n} [W(G_{\mathbf{b}}^*) - W(\widehat{G})] d\pi(\mathbf{b}) \\
&= \int_{\mathbf{b}} E_{P_{\mathbf{b}}^n} [\Lambda(1 - \mu^*) - \Lambda(1 - \mu_{\widehat{G}})] d\pi(\mathbf{b}) \\
&\geq \int_{\mathbf{b}} E_{P_{\mathbf{b}}^n} [|\Lambda'(1 - \mu_{\widehat{G}})|(\mu^* - \mu_{\widehat{G}})] d\pi(\mathbf{b}) \\
&\geq |\Lambda'(t^*)| \int_{\mathbf{b}} E_{P_{\mathbf{b}}^n} [\mu^* - \mu_{\widehat{G}}] d\pi(\mathbf{b}), \tag{A.13}
\end{aligned}$$

where the fourth line follows since $\Lambda(\cdot)$ is convex and non-increasing. The fifth line follows from the observation that for all $P \in \mathcal{P}^*$, $\mu_G \geq 1 - t^*$ for any treatment rule G , therefore $|\Lambda'(1 - \mu_{\widehat{G}})| \geq |\Lambda'(t^*)|$.

Consider now bounding $\int_{\mathbf{b}} E_{P_{\mathbf{b}}^n} [\mu^* - \mu_{\widehat{G}}] d\pi(\mathbf{b})$ from below. Building on the lower bound calculation for the classification risk of the empirical risk minimizing classifier in Lugosi (2002), the proof of Theorem 2.2 in Kitagawa and Tetenov (2015) considers bounding a similar quantity, though the current construction of \mathcal{P}^* is different from theirs. Therefore, in what follows, we reproduce the proof of Theorem 2.2 in Kitagawa and Tetenov (2015) with some necessary modifications.

Consider

$$\begin{aligned}
\int_{\mathbf{b}} E_{P_{\mathbf{b}}^n} [\mu^* - \mu_{\widehat{G}}] d\pi(\mathbf{b}) &\geq \gamma \int_{\mathbf{b}} E_{P_{\mathbf{b}}^n} [P_X(G_{\mathbf{b}}^* \Delta \widehat{G})] d\pi(\mathbf{b}) \\
&= \gamma \int_{\mathbf{b}} \int_{Z_1, \dots, Z_n} P_X(\{b(X) \neq \widehat{b}(X)\}) dP^n(Z_1, \dots, Z_n | \mathbf{b}) d\pi(\mathbf{b}) \\
&\geq \inf_{\widehat{G}} \gamma \int_{\mathbf{b}} \int_{Z_1, \dots, Z_n} P_X(\{b(X) \neq \widehat{b}(X)\}) dP^n(Z_1, \dots, Z_n | \mathbf{b}) d\pi(\mathbf{b})
\end{aligned}$$

where each $b(X)$ and $\widehat{b}(X)$ is an element of \mathbf{b} and $\widehat{\mathbf{b}}$ such that $b(x_j) = b_j$, $\widehat{b}(x_j) = \widehat{b}_j$, and $b(x_v) = \widehat{b}(x_v) = 0$. Note that the last expression can be seen as the minimized Bayes risk with the loss function corresponding to the classification error for predicting binary unknown random variable $b(X)$. Hence, the minimizer of the Bayes risk is attained by the

Bayes classifier,

$$\hat{G}^* = \left\{ x_j : \pi(b_j = 1 | Z_1, \dots, Z_n) \geq \frac{1}{2}, j < v \right\},$$

where $\pi(b_j | Z_1, \dots, Z_n)$ is the posterior of b_j . The minimized Bayes risk is given by

$$\begin{aligned} & \gamma \int_{Z_1, \dots, Z_n} E_X [\min \{ \pi(b(X) = 1 | Z_1, \dots, Z_n), 1 - \pi(b(X) = 1 | Z_1, \dots, Z_n) \}] d\tilde{P}^n \\ &= \gamma \int_{Z_1, \dots, Z_n} \sum_{j=1}^{v-1} p [\min \{ \pi(b_j = 1 | Z_1, \dots, Z_n), 1 - \pi(b_j = 1 | Z_1, \dots, Z_n) \}] d\tilde{P}^n \end{aligned} \quad (\text{A.14})$$

where \tilde{P}^n is the marginal likelihood of $\{(Y_i, D_i, X_i) : i = 1, \dots, n\}$ corresponding to prior $\pi(\mathbf{b})$. For each $j = 1, \dots, (v-1)$ let

$$\begin{aligned} k_j^+ &= \# \{i : X_i = x_j, Y_i D_i = 1 \text{ or } (1 - Y_i)(1 - D_i) = 1\}, \\ k_j^- &= \# \{i : X_i = x_j, (1 - Y_i) D_i = 1 \text{ or } Y_i(1 - D_i) = 1\}. \end{aligned}$$

The posterior for $b_j = 1$ can be written as

$$\pi(b_j = 1 | Z_1, \dots, Z_n) = \begin{cases} \frac{1}{2} & \text{if } \#\{i : X_i = x_j\} = 0, \\ \frac{\left(\frac{1+\gamma}{2}\right)^{k_j^+} \left(\frac{1-\gamma}{2}\right)^{k_j^-}}{\left(\frac{1+\gamma}{2}\right)^{k_j^+} \left(\frac{1-\gamma}{2}\right)^{k_j^-} + \left(\frac{1+\gamma}{2}\right)^{k_j^-} \left(\frac{1-\gamma}{2}\right)^{k_j^+}} & \text{otherwise.} \end{cases}$$

Hence,

$$\begin{aligned} & \min \{ \pi(b_j = 1 | Z_1, \dots, Z_n), 1 - \pi(b_j = 1 | Z_1, \dots, Z_n) \} \\ &= \frac{\min \left\{ \left(\frac{1+\gamma}{2}\right)^{k_j^+} \left(\frac{1-\gamma}{2}\right)^{k_j^-}, \left(\frac{1+\gamma}{2}\right)^{k_j^-} \left(\frac{1-\gamma}{2}\right)^{k_j^+} \right\}}{\left(\frac{1+\gamma}{2}\right)^{k_j^+} \left(\frac{1-\gamma}{2}\right)^{k_j^-} + \left(\frac{1+\gamma}{2}\right)^{k_j^-} \left(\frac{1-\gamma}{2}\right)^{k_j^+}} \\ &= \frac{\min \left\{ 1, \left(\frac{1+\gamma}{1-\gamma}\right)^{k_j^+ - k_j^-} \right\}}{1 + \left(\frac{1+\gamma}{1-\gamma}\right)^{k_j^+ - k_j^-}} \\ &= \frac{1}{1 + a^{|k_j^+ - k_j^-|}}, \text{ where } a = \frac{1+\gamma}{1-\gamma} > 1. \end{aligned} \quad (\text{A.15})$$

Coarsen an observation of (Y_i, D_i) into \tilde{Y}_i defined as

$$\tilde{Y}_i = \begin{cases} 1 & \text{if } Y_i D_i + (1 - Y_i)(1 - D_i) = 1, \\ -1 & \text{otherwise.} \end{cases} \quad (\text{A.16})$$

Since $k_j^+ - k_j^- = \sum_{i: X_i = x_j} \tilde{Y}_i$, plugging (A.15) into (A.14) yields

$$\gamma \sum_{j=1}^{v-1} p E_{\tilde{P}^n} \left[\frac{1}{1 + a^{|\sum_{i: X_i = x_j} \tilde{Y}_i|}} \right] \geq \frac{\gamma}{2} \sum_{j=1}^{v-1} p E_{\tilde{P}^n} \left[\frac{1}{a^{|\sum_{i: X_i = x_j} \tilde{Y}_i|}} \right] \geq \frac{\gamma}{2} p \sum_{i=1}^{v-1} a^{-E_{\tilde{P}^n} |\sum_{i: X_i = x_j} \tilde{Y}_i|},$$

where $E_{\tilde{P}^n}(\cdot)$ is the expectation with respect to the marginal likelihood of $\{(Y_i, D_i, X_i), i = 1, \dots, n\}$. The second line follows by $a > 1$, and the third line follows by Jensen's inequality. Given our prior specification for \mathbf{b} , the marginal distribution of Y_i is $\Pr(\tilde{Y}_i = 1) = \Pr(\tilde{Y}_i = -1) = 1/2$. Hence,

$$E_{\tilde{P}^n} \left| \sum_{i: X_i = x_j} \tilde{Y}_i \right| = \sum_{k=0}^n \binom{n}{k} p^k (1-p)^{n-k} E \left| 2B(k, \frac{1}{2}) - k \right|$$

holds, where $B(k, \frac{1}{2})$ is a random variable following the binomial distribution with parameters k and $\frac{1}{2}$. By noting

$$\begin{aligned} E \left| B(k, \frac{1}{2}) - \frac{k}{2} \right| &\leq \sqrt{E \left(B(k, \frac{1}{2}) - \frac{k}{2} \right)^2} \quad (\because \text{Cauchy-Schwartz inequality}) \\ &= \sqrt{\frac{k}{4}}, \end{aligned}$$

we obtain

$$\begin{aligned} E_{\tilde{P}^n} \left| \sum_{i: X_i = x_j} \tilde{Y}_i \right| &\leq \sum_{k=0}^n \binom{n}{k} p^k (1-p)^{n-k} \sqrt{k} \\ &= E \sqrt{B(n, p)} \\ &\leq \sqrt{np}. \quad (\because \text{Jensen's inequality}). \end{aligned}$$

Hence, the Bayes risk (A.14) is bounded from below by

$$\begin{aligned} &\frac{\gamma}{2} p (v-1) a^{-\sqrt{np}} \\ &\geq \frac{\gamma}{2} p (v-1) e^{-(a-1)\sqrt{np}} \quad (\because 1+x \leq e^x \forall x) \\ &= \frac{p\gamma}{2} (v-1) e^{-\frac{2\gamma}{1-\gamma}\sqrt{np}}, \end{aligned} \tag{A.17}$$

therefore

$$\int_{\mathbf{b}} E_{P_{\mathbf{b}}^n} [\mu^* - \mu_{\hat{G}}] d\pi(\mathbf{b}) \geq \frac{p\gamma}{2} (v-1) e^{-\frac{2\gamma}{1-\gamma}\sqrt{np}}. \tag{A.18}$$

This lower bound of the Bayes risk has the slowest convergence rate when γ is set to be proportional to $n^{-1/2}$. Specifically, let $\gamma = \sqrt{\frac{v-1}{nt^*}}$. Then for all $n \geq 4(v-1)/t^*$, $\gamma \leq 1/2$ and

since $p = \frac{t^*}{v-1}$,

$$-\frac{2\gamma}{1-\gamma}\sqrt{np} = -\frac{2}{1-\gamma}\sqrt{\frac{v-1}{nt^*}}\sqrt{\frac{nt^*}{v-1}} = -\frac{2}{1-\gamma} \geq -4.$$

Then

$$\frac{p\gamma}{2}(v-1)e^{-\frac{2\gamma}{1-\gamma}\sqrt{np}} \geq \frac{p\gamma}{2}(v-1)e^{-4} = \frac{t^*}{2}\sqrt{\frac{v-1}{nt^*}}e^{-4} = \frac{e^{-4}}{2}\sqrt{t^*}\sqrt{\frac{v-1}{n}}$$

Inserting this bound into (A.18) and multiplying by M provides a lower bound for (A.13).

This completes the proof. \square

Proof of Theorem 4.1. For any $G \in \mathcal{G}$, it holds

$$\begin{aligned} W(F_G) - W(F_{\hat{G}^e}) &\leq W(\hat{F}_G) - W(\hat{F}_G^e) - W(\hat{F}_{\hat{G}^e}) + W(\hat{F}_{\hat{G}^e}^e) \\ &\quad + W(F_G) - W(F_{\hat{G}^e}) - W(\hat{F}_G) + W(\hat{F}_{\hat{G}^e}) \\ &\leq 2 \sup_{G \in \mathcal{G}} |W(\hat{F}_G) - W(\hat{F}_G^e)| + 2 \sup_{G \in \mathcal{G}} |W(\hat{F}_G) - W(F_G)|, \end{aligned} \quad (\text{A.19})$$

where the first inequality uses $W(\hat{F}_{\hat{G}^e}^e) - W(\hat{F}_G^e) \geq 0$. By convexity and monotonicity of $\Lambda(\cdot)$, we have

$$|W(\hat{F}_G) - W(\hat{F}_G^e)| \leq |\Lambda'(0)| \int_0^M |\hat{F}_G(y) - \hat{F}_G^e(y)| dy. \quad (\text{A.20})$$

For every y , $|\hat{F}_G(y) - \hat{F}_G^e(y)|$ can be bounded by

$$\begin{aligned} |\hat{F}_G(y) - \hat{F}_G^e(y)| &\leq \frac{1}{n} \sum_{i=1}^n \left| \frac{1}{e(X_i)} - \frac{1}{\hat{e}(X_i)} \right| D_i 1\{Y_i \leq y\} 1\{X_i \in G\} \\ &\quad + \frac{1}{n} \sum_{i=1}^n \left| \frac{1}{1-e(X_i)} - \frac{1}{1-\hat{e}(X_i)} \right| (1-D_i) 1\{Y_i \leq y\} 1\{X_i \notin G\} \\ &= \frac{1}{n} \sum_{i=1}^n \left| \frac{1}{e(X_i)} - \frac{1}{\hat{e}(X_i)} \right| D_i + \frac{1}{n} \sum_{i=1}^n \left| \frac{1}{1-e(X_i)} - \frac{1}{1-\hat{e}(X_i)} \right| (1-D_i) \end{aligned} \quad (\text{A.21})$$

Combining (A.20) and (A.21) gives

$$E_{P^n} \left[\sup_{G \in \mathcal{G}} |W(\hat{F}_G) - W(\hat{F}_G^e)| \right] \leq E_{P^n} \left[\frac{1}{n} \sum_{i=1}^n \left| 1 - \frac{e(X_i)}{\hat{e}(X_i)} \right| \right] + E_{P^n} \left[\frac{1}{n} \sum_{i=1}^n \left| 1 - \frac{1-e(X_i)}{1-\hat{e}(X_i)} \right| \right]. \quad (\text{A.22})$$

Hence, from (A.19) and (A.22), we obtain

$$\begin{aligned} & \sup_{P \in \mathcal{P}_e \cap \mathcal{P}} E_{P^n} [W_G^* - W(F_{\hat{G}_e})] \\ \leq & 2 \sup_{P \in \mathcal{P}_e} E_{P^n} \left[\frac{1}{n} \sum_{i=1}^n \left| 1 - \frac{e(X_i)}{\hat{e}(X_i)} \right| \right] + \sup_{P \in \mathcal{P}_e} E_{P^n} \left[\frac{1}{n} \sum_{i=1}^n \left| 1 - \frac{1 - e(X_i)}{1 - \hat{e}(X_i)} \right| \right] \\ & + 2 \sup_{P \in \mathcal{P}} E_{P^n} \left[\sup_{G \in \mathcal{G}} \left| W(\hat{F}_G) - W(F_G) \right| \right] \end{aligned}$$

Assumption 4.1 and $\sup_{P \in \mathcal{P}} E_{P^n} \left[\sup_{G \in \mathcal{G}} \left| W(\hat{F}_G) - W(F_G) \right| \right] = O(\sqrt{\frac{v}{n}})$ shown in the proof of Theorem 3.1 lead to the conclusion. \square

B Equality-minded EWM with Nonparametrically Estimated Propensity Score

In this appendix, we consider the equality-minded EWM approach with unknown propensity score estimated nonparametrically by local polynomial regressions. We provide regularity conditions under which the nonparametric estimator of the propensity score satisfies Assumption 4.1 with certain ϕ_n .

We consider the leave-one-out local polynomial estimator for $e(\cdot)$, i.e., $\hat{e}(X_i)$ is constructed by fitting the local polynomials excluding the i -th observation. For any multi-index $s = (s_1, \dots, s_{d_x}) \in \mathbb{N}^{d_x}$ and any $(x_1, \dots, x_{d_x}) \in \mathbb{R}^{d_x}$, we define $|s| \equiv \sum_{i=1}^{d_x} s_i$, $s! \equiv s_1! \cdots s_{d_x}!$, $x^s \equiv x_1^{s_1} \cdots x_{d_x}^{s_{d_x}}$, and $\|x\| \equiv (x_1^2 + \cdots + x_{d_x}^2)$. Let $K(\cdot) : \mathbb{R}^{d_x} \rightarrow \mathbb{R}$ be a kernel function and $h > 0$ be a bandwidth, whose dependence on the sample size is implicit in the notation. At each X_i , $i = 1, \dots, n$, we define the leave-one-out local polynomial coefficient estimators with degree $l \geq 0$ as

$$\hat{\theta}(X_i) = \arg \min_{\theta} \sum_{j \neq i} \left[D_j - \theta^T U \left(\frac{X_j - X_i}{h} \right) \right]^2 K \left(\frac{X_j - X_i}{h} \right),$$

where $U \left(\frac{X_j - X_i}{h} \right)$ is the vector with elements indexed by the multi-index s , i.e., $U \left(\frac{X_j - X_i}{h} \right) \equiv \left(\left(\frac{X_j - X_i}{h} \right)^s \right)_{|s| \leq l}$. With a slight abuse of notation, we define $U(0) = (1, 0, \dots, 0)^T$. Let $\lambda_n(X_i)$ be the smallest eigenvalue of $B(X_i) \equiv (nh^{d_x})^{-1} \sum_{j \neq i} U \left(\frac{X_j - X_i}{h} \right) U^T \left(\frac{X_j - X_i}{h} \right) K \left(\frac{X_i - X_j}{h} \right)$.

Accordingly, we construct leave-one-out local polynomial fit for $e(X_i)$ by

$$\tilde{e}(X_i) = U^T(0)\hat{\theta}(X_i) \cdot 1\{\lambda_n(X_i) \geq t_n\} \quad (\text{B.1})$$

where t_n is a positive sequence that slowly converges to zero, such as $t_n \propto (\log n)^{-1}$. This trimming constant regularizes the regressor matrix of the local polynomial regression and simplifies the proof of the uniform consistency of the local polynomial estimator.

To characterize \mathcal{P}_e in Assumption 4.1, we impose the following restrictions, which are identical to Assumption 2.4 in Kitagawa and Tetenov (2015).

Assumption B.1. (*Smooth-e*) *Smoothness of the propensity score:* The propensity score $e(\cdot)$ belongs to a Hölder class of functions with degree $\beta_e \geq 1$ and constant $L_e < \infty$.⁷

(*PX*) *Support and Density Restrictions on P_X :* Let $\mathcal{X} \subset \mathbb{R}^{d_x}$ be the support of P_X . Let $Leb(\cdot)$ be the Lebesgue measure on \mathbb{R}^{d_x} . There exist constants \underline{c} and r_0 such that

$$Leb(\mathcal{X} \cap B(x, r)) \geq \underline{c}Leb(B(x, r)) \quad \forall 0 < r \leq r_0, \forall x \in \mathcal{X}, \quad (\text{B.2})$$

and P_X has the density function $\frac{dP_X}{dx}(\cdot)$ with respect to the Lebesgue measure of \mathbb{R}^{d_x} that is bounded from above and bounded away from zero, $0 < \underline{p}_X \leq \frac{dP_X}{dx}(x) \leq \bar{p}_X < \infty$ for all $x \in \mathcal{X}$.

(*Ker*) *Bounded Kernel with Compact Support:* The kernel function $K(\cdot)$ have support $[-1, 1]^{d_x}$, $\int_{\mathbb{R}^{d_x}} K(u)du = 1$, and $\sup_u K(u) \leq K_{\max} < \infty$.

Assumption B.1 (PX) is borrowed from Audibert and Tsybakov (2007), and it provides regularity conditions on the marginal distribution of X . Inequality condition (B.2) constrains the shape of the support of X , and it essentially rules out the case where \mathcal{X} has “sharp” spikes, i.e., $\mathcal{X} \cap B(x, r)$ has an empty interior or $Leb(\mathcal{X} \cap B(x, r))$ converges to zero as $r \rightarrow 0$ faster than the rate of r^2 for some x in the boundary of \mathcal{X} .

⁷Let D^s denote the differential operator $D^s \equiv \frac{\partial^{s_1 + \dots + s_{d_x}}}{\partial x_1^{s_1} \dots \partial x_{d_x}^{s_{d_x}}}$. Let $\beta \geq 1$ be an integer. For any $x \in \mathbb{R}^{d_x}$ and any $(\beta - 1)$ times continuously differentiable function $f : \mathbb{R}^{d_x} \rightarrow \mathbb{R}$, we denote the Taylor expansion polynomial of degree $(\beta - 1)$ at point x by $f_x(x') \equiv \sum_{|s| \leq \beta - 1} \frac{(x' - x)^s}{s!} D^s f(x)$. Let $L > 0$. The Hölder class of functions in \mathbb{R}^{d_x} with degree β and constant $0 < L < \infty$ is defined as the set of function $f : \mathbb{R}^{d_x} \rightarrow \mathbb{R}$ that are $(\beta - 1)$ times continuously differentiable and satisfy, for any x and $x' \in \mathbb{R}^{d_x}$, the inequality $|f_x(x') - f(x)| \leq L \|x - x'\|^\beta$.

The next lemma collects several properties of the local polynomial estimators that are useful to prove the bound shown in (4.5). These claims are borrowed from Theorem 3.2 in Audibert and Tsybakov (2007) and Lemma B.4 in Kitagawa and Tetenov (2015).

Lemma B.1. *Let \mathcal{P}_e consist of the data generating processes satisfying Assumption B.1 (Smooth-e) and (PX). Let $\tilde{e}(X_i)$ be the leave-one-out estimator for the propensity score defined in (B.1) whose kernel function satisfies B.1 (Ker).*

(i) *There exist positive constants c_1 , c_2 , and c_3 that depend only on β_e , d_x , L_e , \underline{c} , r_0 , \underline{p}_X , and \bar{p}_X , such that, for any $0 < h < r_0/\underline{c}$, any $c_1 h^\beta < \delta$, and any $n \geq 2$,*

$$P^{n-1} (|\tilde{e}(x) - e(x)| > \delta) \leq c_2 \exp(-c_3 n h^{d_x} \delta^2),$$

holds for almost all x with respect to P_X , where $P^{n-1}(\cdot)$ is the distribution of $\{(Y_i, D_i, X_i)_{i=1}^{n-1}\}$.

(ii)

$$\sup_{P \in \mathcal{P}_e} \int_{\mathcal{X}} E_{P^{n-1}} [|\tilde{e}(x) - e(x)|] dP_X(x) \leq O(h^{\beta_e}) + O\left(\frac{1}{\sqrt{nh^{d_x}}}\right) \quad (\text{B.3})$$

holds. Hence, an optimal choice of bandwidth that leads to the fastest convergence rate of the uniform upper bound is $h \propto n^{-\frac{1}{2\beta_e + d_x}}$ and the resulting uniform convergence rate is

$$\sup_{P \in \mathcal{P}_e} \int_{\mathcal{X}} E_{P^{n-1}} [|\tilde{e}(x) - e(x)|] dP_X(x) \leq O\left(n^{-\frac{1}{2+d_x/\beta_e}}\right).$$

Making use of Lemma B.1, the next proposition shows a propensity score estimator constructed by suitably trimming $\tilde{e}(X_i)$ satisfies Assumption 4.1.

Proposition B.1. *Let \mathcal{P}_e consist of data generating processes that satisfy Assumption B.1 (Smooth-e) and (PX). Let $\tilde{e}(X_i)$ be the leave-one-out local polynomial estimator with degree $l = (\beta_e - 1)$ whose kernel satisfies Assumption B.1 (Ker). Let*

$$\hat{e}(X_i) \equiv \min\{1 - \epsilon_n, \max\{\epsilon_n, \tilde{e}(X_i)\}\} \in [\epsilon_n, 1 - \epsilon_n] \quad (\text{B.4})$$

with a sequence of trimming constants ϵ_n satisfies $\epsilon_n = O(n^{-a})$ for some $a > 0$. Then, $\hat{e}(X_i)$ satisfies Assumption 4.1 with $\phi_n = n^{\frac{1}{2+d_x/\beta_e}}$.

Proof of Proposition B.1. Assume that n is large enough so that $\varepsilon_n \leq \kappa/2$ holds. Since $\hat{e}(X_i) = \tilde{e}(X_i)$ whenever $\tilde{e}(X_i) \in [\frac{\kappa}{2}, 1 - \frac{\kappa}{2}] \subset [\varepsilon_n, 1 - \varepsilon_n]$, the following bounds are valid

$$\left| 1 - \frac{e(X_i)}{\hat{e}(X_i)} \right| \leq \begin{cases} \frac{2}{\kappa} |\tilde{e}(X_i) - e(X_i)| & \text{if } \tilde{e}(X_i) \in [\frac{\kappa}{2}, 1 - \frac{\kappa}{2}] \\ \varepsilon_n^{-1} & \text{if } \tilde{e}(X_i) \notin [\frac{\kappa}{2}, 1 - \frac{\kappa}{2}]. \end{cases}$$

Hence,

$$\begin{aligned} E_{P^n} \left[\frac{1}{n} \sum_{i=1}^n \left| 1 - \frac{e(X_i)}{\hat{e}(X_i)} \right| \right] &= E_{P^n} \left[\left| 1 - \frac{e(X_n)}{\hat{e}(X_n)} \right| \right] \\ &\leq \frac{2}{\kappa} E_{P^n} |\tilde{e}(X_n) - e(X_n)| + \varepsilon_n^{-1} P^n \left(\tilde{e}(X_n) \notin \left[\frac{\kappa}{2}, 1 - \frac{\kappa}{2} \right] \right) \end{aligned} \quad (\text{B.5})$$

By Lemma B.1 (ii),

$$\sup_{P \in \mathcal{P}_e} E_{P^n} |\tilde{e}(X_n) - e(X_n)| = \sup_{P \in \mathcal{P}_e} \int_{\mathcal{X}} E_{P^{n-1}} \left[\left| 1 - \frac{e(x)}{\hat{e}(x)} \right| \right] dP_X(x) \leq O \left(n^{-\frac{1}{2+d_x/\beta_e}} \right)$$

Furthermore, by Lemma B.1 (i),

$$\begin{aligned} P^n \left(\tilde{e}(X_n) \notin \left[\frac{\kappa}{2}, 1 - \frac{\kappa}{2} \right] \right) &= \int_{\mathcal{X}} P^{n-1} \left(\tilde{e}(x) \notin \left[\frac{\kappa}{2}, 1 - \frac{\kappa}{2} \right] \right) dP_X(x) \\ &\leq \int_{\mathcal{X}} P^{n-1} \left(|\tilde{e}(x) - e(x)| \geq \frac{\kappa}{2} \right) dP_X(x) \\ &\leq c_2 \exp \left(-\frac{c_3 \kappa^2}{4} n h^{d_x} \right) \end{aligned}$$

holds for all n satisfying $c_1 h^\beta < \kappa/2$, where the c_1 , c_2 , and c_3 are the constants defined in

Lemma B.1 (i). Since ε_n is assumed to converge at a polynomial rate, $\varepsilon_n^{-1} P^n \left(\hat{e}(X_n) \notin \left[\frac{\kappa}{2}, 1 - \frac{\kappa}{2} \right] \right)$ converges faster than $O(n^{-\frac{1}{2+d_x/\beta_e}})$. Thus, from (B.5), we conclude $\sup_{P \in \mathcal{P}_e} E_{P^n} \left[\frac{1}{n} \sum_{i=1}^n \left| 1 - \frac{e(X_i)}{\hat{e}(X_i)} \right| \right] \leq O \left(n^{-\frac{1}{2+d_x/\beta_e}} \right)$. The other bound of Assumption 4.1 can be shown similarly. \square

Combining Proposition B.1 with Theorem 4.1 proves equation (4.5) in the main text.

References

- AABERGE, R., T. HAVNES, AND M. MOGSTAD (2016): “A Partial Identification Approach to Welfare Rankings of Distributions,” *unpublished manuscript*.
- ABADIE, A., J. ANGRIST, AND G. IMBENS (2002): “Instrumental Variables Estimates of the Effect of Subsidized Training on the Quantiles of Trainee Earnings,” *Econometrica*, 70, 91–117.
- ATKINSON, A. B. (1970): “On the measurement of inequality,” *Journal of Economic Theory*, 2, 244 – 263.
- AUDIBERT, J.-Y. AND A. B. TSYBAKOV (2007): “Fast Learning Rates for Plug-in Classifiers,” *The Annals of Statistics*, 35, 608–633.
- BHATTACHARYA, D. AND P. DUPAS (2012): “Inferring Welfare Maximizing Treatment Assignment under Budget Constraints,” *Journal of Econometrics*, 167, 168 – 196.
- BITLER, M. P., J. B. GELBACH, AND H. W. HOYNES (2006): “What Mean Impacts Miss: Distributional Effects of Welfare Reform Experiments,” *American Economic Review*, 96, 988–1012.
- BLACKORBY, C. AND D. DONALDSON (1978): “Measures of Relative Equality and Their Meaning in Terms of Social Welfare,” *Journal of Economic Theory*, 18, 59–80.
- 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,” *Journal of Human Resources*, 32, 549–576.
- CHAMBERLAIN, G. (2011): “Bayesian Aspects of Treatment Choice,” in *The Oxford Handbook of Bayesian Econometrics*, ed. by J. Geweke, G. Koop, and H. van Dijk, Oxford University Press, 11–39.
- COWELL, F. (1995): *Measuring Inequality*, Hemel Hempstead: Harvester Wheatsheaf, 2 ed.

- (2000): “Measurement of Inequality,” in *Handbook of Income Distribution*, ed. by A. Atkinson and F. Bourguignon, Elsevier, vol. 1.
- DEHEJIA (2005): “Program Evaluation as a Decision Problem,” *Journal of Econometrics*, 125, 141–173.
- DONALDSON, D. AND J. A. WEYMARK (1980): “A Single-Parameter Generalization of the Gini Indices of Inequality,” *Journal of Economic Theory*, 22, 67–86.
- (1983): “Ethically Flexible Gini Indices for Income Distributions in the Continuum,” *Journal of Economic Theory*, 29, 353–358.
- FAN, Y. AND S. S. PARK (2010): “Sharp Bounds on the Distribution of Treatment Effects and Their Statistical Inference,” *Econometric Theory*, 26, 931–951.
- FIRPO, S. AND C. PINTO (2015): “Identification and Estimation of Distributional Impacts of Interventions Using Changes in Inequality Measures,” *Journal of Applied Econometrics*, forthcoming.
- FIRPO, S. AND G. RIDDER (2008): “Bounds on functionals of the distribution of treatment effects,” *unpublished manuscript*.
- HECKMAN, J. J., H. ICHIMURA, AND P. E. TODD (1997a): “Matching As An Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme,” *The Review of Economic Studies*, 64, 605–654.
- HECKMAN, J. J., J. SMITH, AND N. CLEMENTS (1997b): “Making The Most Out Of Programme Evaluations and Social Experiments: Accounting For Heterogeneity in Programme Impacts,” *The Review of Economic Studies*, 64, 487–535.
- HIRANO, K. AND J. R. PORTER (2009): “Asymptotics for Statistical Treatment Rules,” *Econometrica*, 77, 1683–1701.
- KASY, M. (2014): “Using Data to Inform Policy,” *unpublished manuscript*.
- (2016): “Partial Identification, Distributional Preferences, and the Welfare Ranking of Policies,” *Review of Economics and Statistics*, 98, 111–131.

- KITAGAWA, T. AND A. TETENOV (2015): “Who should be Treated? Empirical Welfare Maximization Methods for Treatment Choice,” *Cemmap working paper*, <http://dx.medra.org/10.1920/wp.cem.2015.1015>.
- LAMBERT, P. J. (2001): *The Distribution and Redistribution of Income*, Manchester University Press, 3 ed.
- LUGOSI, G. (2002): “Pattern Classification and Learning Theory,” in *Principles of Non-parametric Learning*, ed. by L. Györfi, Vienna: Springer, 1–56.
- MACHINA, M. (1982): ““Expected Utility” Analysis without the Independence Axiom,” *Econometrica*, 50, 277–323.
- MANSKI, C. F. (2004): “Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica*, 72, 1221–1246.
- MEHRAN, F. (1976): “Linear Measures of Income Inequality,” *Econometrica*, 44, 805–09.
- QUIGGIN, J. (1982): “A Theory of Anticipated Utility,” *Journal of Economic Behavior and Organization*, 3, 323–343.
- STOYE, J. (2009): “Minimax regret treatment choice with finite samples,” *Journal of Econometrics*, 151, 70 – 81.
- (2012): “Minimax regret treatment choice with covariates or with limited validity of experiments,” *Journal of Econometrics*, 166, 138 – 156.
- TETENOV, A. (2012): “Statistical treatment choice based on asymmetric minimax regret criteria,” *Journal of Econometrics*, 166, 157–165.
- VAPNIK, V. N. (1998): *Statistical Learning Theory*, John Wiley & Sons.
- WEYMARK, J. A. (1981): “Generalized gini inequality indices,” *Mathematical Social Sciences*, 1, 409 – 430.
- YAARI, M. E. (1988): “A controversial proposal concerning inequality measurement,” *Journal of Economic Theory*, 44, 381 – 397.