

# WHO SHOULD BE TREATED? EMPIRICAL WELFARE MAXIMIZATION METHODS FOR TREATMENT CHOICE

TORU KITAGAWA

Cemmap and Department of Economics, University College London

ALEKSEY TETENOV

Department of Economics, University of Bristol

One of the main objectives of empirical analysis of experiments and quasi-experiments is to inform policy decisions that determine the allocation of treatments to individuals with different observable covariates. We study the properties and implementation of the *Empirical Welfare Maximization (EWM) method*, which estimates a treatment assignment policy by maximizing the sample analog of average social welfare over a class of candidate treatment policies. The EWM approach is attractive in terms of both statistical performance and practical implementation in realistic settings of policy design. Common features of these settings include: (i) feasible treatment assignment rules are constrained exogenously for ethical, legislative, or political reasons, (ii) a policy maker wants a simple treatment assignment rule based on one or more eligibility scores in order to reduce the dimensionality of individual observable characteristics, and/or (iii) the proportion of individuals who can receive the treatment is a priori limited due to a budget or a capacity constraint. We show that when the propensity score is known, the average social welfare attained by EWM rules converges at least at  $n^{-1/2}$  rate to the maximum obtainable welfare uniformly over a minimally constrained class of data distributions, and this uniform convergence rate is minimax optimal. We examine how the uniform convergence rate depends on the richness of the class of candidate decision rules, the distribution of conditional treatment effects, and the lack of knowledge of the propensity score. We offer easily implementable algorithms for computing the EWM rule and an application using experimental data from the National JTPA Study.

**KEYWORDS:** Heterogeneous treatment effects, randomized experiments, program evaluation, individualized treatment rules, empirical risk minimization, risk bounds.

## 1. INTRODUCTION

TREATMENT EFFECTS OFTEN VARY with observable individual characteristics. An important objective of empirical analysis of experimental and quasi-experimental data is to determine the individuals who should be treated based on their observable characteristics. Empirical researchers often use regression estimates of individual treatment effects to infer the set of individuals who benefit or do not benefit from the treatment and to suggest who should be targeted for treatment. This paper advocates the *Empirical Welfare Maximization (EWM) method*, which offers an alternative way to choose optimal treatment

---

Toru Kitagawa: t.kitagawa@ucl.ac.uk

Aleksey Tetenov: a.tetenov@bristol.ac.uk

We would like to thank Susan Athey, Kei Hirano, Guido Imbens, Shin Kanaya, Charles Manski, Joerg Stoye, and seminar participants at Academia Sinica, Bern, Brown, BC, BU, Cambridge, CEMFI, CORE, Duke, GRIPS, JHU, NYU, UPenn, UPF, UVA, and Warwick for beneficial comments. We would also like to thank the participants at 2014 Bristol Econometric Study Group, 2014 SETA Conference, Hakone Microeconometrics Conference, and ICEEE 2015 for beneficial comments and discussions. Financial support from the ESRC through the ESRC Centre for Microdata Methods and Practice (CeMMAP) Grant RES-589-28-0001 and the European Research Council Starting Grant 715940 is gratefully acknowledged.

assignment based on experimental or observational data from program evaluation studies. We study the frequentist properties of the EWM treatment choice rule and show its optimality in terms of welfare convergence rate, which measures how quickly the average welfare attained by practicing the estimated treatment choice rule converges to the maximal welfare attainable with the knowledge of the true data generating process. We also argue that the EWM approach is well-suited for policy design problems, since it easily accommodates many practical policy concerns, including (i) feasible treatment assignment rules being constrained exogenously for ethical, legislative, or political reasons, (ii) the policy maker facing a budget or capacity constraint that limits the proportion of individuals who can receive one of the treatments, or (iii) the policy maker wanting to have a simple treatment assignment rule based on one or more indices (eligibility scores) to reduce the dimensionality of individual characteristics.

Let the data be a size  $n$  random sample of  $Z_i = (Y_i, D_i, X_i)$ , where  $X_i \in \mathcal{X} \subset \mathbb{R}^{d_x}$  refers to observable pre-treatment covariates of individual  $i$ ,  $D_i \in \{0, 1\}$  is a binary indicator of the individual's treatment assignment, and  $Y_i \in \mathbb{R}$  is her/his post-treatment observed outcome. The population from which the sample is drawn is characterized by  $P$ , a joint distribution of  $(Y_{0,i}, Y_{1,i}, D_i, X_i)$ , where  $Y_{0,i}$  and  $Y_{1,i}$  are potential outcomes that would have been observed if  $i$ 's treatment status were  $D_i = 0$  and  $D_i = 1$ , respectively. We assume *unconfoundedness*, meaning that in the data treatments are assigned independently of the potential outcomes  $(Y_{0,i}, Y_{1,i})$  conditionally on observable characteristics  $X_i$ . Based on this data, the policy maker has to choose a conditional treatment rule that determines whether individuals with covariates  $X$  in a target population will be assigned to treatment 0 or to treatment 1. We restrict our analysis to non-randomized treatment rules. The set of treatment rules could then be indexed by their *decision sets*  $G \subset \mathcal{X}$  of covariate values, which determine the group of individuals  $\{X \in G\}$  to whom treatment 1 is assigned. We denote the collection of candidate treatment rules by  $\mathcal{G} = \{G \subset \mathcal{X}\}$ .

The goal of our analysis is to empirically select a treatment assignment rule that gives the highest *welfare* to the target population. We assume that the joint distribution of  $(Y_{0,i}, Y_{1,i}, X_i)$  of the target population is identical to that of the sampled population.<sup>1</sup> We consider the *additive* welfare criterion defined by the average of the individual outcomes in the target population.<sup>2</sup> When treatment rule  $G$  is applied to the target population, the social welfare defined by the sum of individual outcomes in the population is proportional to

$$W(G) \equiv E_P[Y_1 \cdot 1\{X \in G\} + Y_0 \cdot 1\{X \notin G\}], \quad (1.1)$$

where  $E_P(\cdot)$  is the expectation with respect to  $P$ . Our framework could incorporate a broad range of social preferences by suitably redefining the outcome variable. Setting  $Y$  to be a concave transformation of one's measure of wealth leads to an inequality-averse social welfare of Atkinson (1970). When multiple outcome variables enter into the individual utility (e.g., consumption and leisure),  $Y$  can be set to a known function of these outcomes. The cost of treatment can be incorporated into the social welfare by redefining the individual potential outcome  $Y_d$  to be the outcome minus the cost of treatment  $d$ .

Denoting the conditional mean treatment response by  $m_d(x) \equiv E[Y_d|X = x]$  and the conditional average treatment effect by  $\tau(x) \equiv m_1(x) - m_0(x)$ , we could also express the

<sup>1</sup>In Remark 2.2, we consider a setting where the target and the sampled populations have identical conditional treatment effects, but different marginal distributions of  $X$ .

<sup>2</sup>In the econometrics literature of treatment choice, the additive social welfare is often referred to as a utilitarian social welfare.

welfare criterion as

$$W(G) = E_P(m_0(X)) + E_P[\tau(X) \cdot 1\{X \in G\}]. \tag{1.2}$$

Assuming unconfoundedness, equivalence of the distributions of  $(Y_{0,i}, Y_{1,i}, X_i)$  between the target and sampled populations, and the overlap condition for the propensity score  $e(X) = E_P[D|X]$  in the sampled population, the welfare criterion (1.1) can be written equivalently as

$$\begin{aligned} W(G) &= E_P \left[ \frac{YD}{e(X)} \cdot 1\{X \in G\} + \frac{Y(1-D)}{1-e(X)} \cdot 1\{X \notin G\} \right] \\ &= E_P(Y_0) + E_P \left[ \left( \frac{YD}{e(X)} - \frac{Y(1-D)}{1-e(X)} \right) \cdot 1\{X \in G\} \right]. \end{aligned} \tag{1.3}$$

Hence, if the probability distribution of observables  $(Y, D, X)$  is fully known to the decision-maker, an optimal treatment rule from the utilitarian perspective can be written as

$$G^* \in \arg \max_{G \in \mathcal{G}} W(G), \tag{1.4}$$

or, equivalently, as a maximizer of the welfare gain relative to  $E_P(Y_0)$ :

$$G^* \in \arg \max_{G \in \mathcal{G}} E_P[\tau(X) \cdot 1\{X \in G\}], \quad \text{or} \tag{1.5}$$

$$G^* \in \arg \max_{G \in \mathcal{G}} E_P \left[ \left( \frac{YD}{e(X)} - \frac{Y(1-D)}{1-e(X)} \right) \cdot 1\{X \in G\} \right]. \tag{1.6}$$

The main idea of *Empirical Welfare Maximization (EWM)* is to solve a sample analog of the population maximization problem (1.4),

$$\begin{aligned} \hat{G}_{\text{EWM}} &\in \arg \max_{G \in \mathcal{G}} W_n(G), \quad \text{where} \\ W_n(G) &= E_n \left[ \frac{Y_i D_i}{e(X_i)} \cdot 1\{X_i \in G\} + \frac{Y_i(1-D_i)}{1-e(X_i)} \cdot 1\{X_i \notin G\} \right] \end{aligned} \tag{1.7}$$

and  $E_n(\cdot)$  is the sample average. One notable feature of our framework is that the class of candidate treatment rules  $\mathcal{G} = \{G \subset \mathcal{X}\}$  is not as rich as the class of all subsets of  $\mathcal{X}$ , and it may not include the *first-best decision set*

$$G_{\text{FB}}^* \equiv \{x \in \mathcal{X} : \tau(x) \geq 0\}, \tag{1.8}$$

which maximizes the population welfare (1.1) if any assignment rule is feasible to implement. Our framework with a constrained class of feasible assignment rules allows us to incorporate several types of exogenous constraints that generally restrict the complexity of feasible treatment assignment rules. For instance, when assigning treatments to individuals in the target population, it may not be realistic to implement a complex treatment assignment rule due to logistic, legal, ethical, or political restrictions.

The largest welfare that could be obtained by any treatment rule in class  $\mathcal{G}$  is

$$W_{\mathcal{G}}^* \equiv \sup_{G \in \mathcal{G}} W(G), \tag{1.9}$$

which is the *second-best* welfare if  $W_G^* < W(G_{FB}^*)$ . In line with Manski (2004) and the subsequent literature on statistical treatment rules, we evaluate the performance of estimated treatment rules  $\hat{G} \in \mathcal{G}$  in terms of their average welfare loss (regret) relative to the maximum feasible welfare  $W_G^*$

$$W_G^* - E_{P^n}[W(\hat{G})] = E_{P^n}[W_G^* - W(\hat{G})] \geq 0, \tag{1.10}$$

where the expectation  $E_{P^n}$  is taken over different realizations of the random sample. This criterion measures the average difference between the best attainable population welfare and the welfare attained by implementing estimated policy  $\hat{G}$ . Since we assess the statistical performance of  $\hat{G}$  by its welfare value  $W(\hat{G})$ , we do not require  $\arg \max_{G \in \mathcal{G}} W(G)$  to be unique or  $\hat{G}$  to converge to a specific set.

Assuming that the propensity score  $e(X)$  is known and bounded away from zero and 1, as is the case in randomized experiments, we derive a non-asymptotic distribution-free upper bound of  $E_{P^n}[W_G^* - W(\hat{G}_{EWM})]$  as a function of sample size  $n$  and a measure of complexity of  $\mathcal{G}$ . Based on this bound, we show that the average welfare of the EWM treatment rule converges to  $W_G^*$  at rate  $O(n^{-1/2})$  uniformly over a minimally constrained class of probability distributions. We also show that this uniform convergence rate of  $\hat{G}_{EWM}$  is optimal in the sense that no estimated treatment choice rule of any kind can attain a faster uniform convergence rate compared to the EWM rule, that is, minimax rate optimality of  $\hat{G}_{EWM}$ . For further refinement of this theoretical result, we analyze how this uniform convergence rate improves if the first-best decision rule  $G_{FB}^*$  is feasible, that is,  $G_{FB}^* \in \mathcal{G}$ , and if the class of data generating processes is constrained by the *margin assumption*, which restricts the distribution of conditional treatment effects in a neighborhood of zero. We show that  $\hat{G}_{EWM}$  remains minimax rate optimal with these additional restrictions.

When the data are from an observational study, the propensity score is usually unknown, so it is not feasible to implement the EWM rule (1.7). As a feasible version of the EWM rule, we consider *hybrid EWM* approaches that plug in estimators of the regression equations or the propensity score in the sample analogs of (1.5) or (1.6). Specifically, with estimated regression functions  $\hat{m}_d(x) = \hat{E}(Y_d|X = x) = \hat{E}(Y|X = x, D = d)$ , we define the *m-hybrid* rule as

$$\hat{G}_{m\text{-hybrid}} \in \arg \max_{G \in \mathcal{G}} E_n[\hat{\tau}^m(X_i) \cdot 1\{X_i \in G\}], \tag{1.11}$$

where  $\hat{\tau}^m(X_i) \equiv \hat{m}_1(X_i) - \hat{m}_0(X_i)$ . Similarly, with the estimated propensity score  $\hat{e}(x)$ , we define an *e-hybrid* rule as

$$\hat{G}_{e\text{-hybrid}} \in \arg \max_{G \in \mathcal{G}} E_n[\hat{\tau}_i^e \cdot 1\{X_i \in G\}], \tag{1.12}$$

where  $\hat{\tau}_i^e \equiv [\frac{Y_i D_i}{\hat{e}(X_i)} - \frac{Y_i(1-D_i)}{1-\hat{e}(X_i)}] \cdot 1\{\varepsilon_n \leq \hat{e}(X_i) \leq 1 - \varepsilon_n\}$  with a converging positive sequence  $\varepsilon_n \rightarrow 0$  as  $n \rightarrow \infty$ . We investigate the performance of these hybrid approaches in terms of the uniform convergence rate of the welfare loss and clarify how this rate is affected by the estimation uncertainty in  $\hat{m}_d(\cdot)$  and  $\hat{e}(\cdot)$ .

Since the welfare criterion function involves optimization over a class of sets, estimation of the EWM and hybrid treatment rules could present challenging computational problems when  $\mathcal{G}$  is rich, similarly to the maximum score estimation (Manski (1975), Manski

and Thompson (1989)). We argue, however, that exact maximization of the EWM criterion is now practically feasible for many problems in economics using widely-available optimization software and an approach proposed by Florios and Skouras (2008), which we extend and improve upon.

To illustrate EWM in practice, we compare EWM and plug-in treatment rules computed from the experimental data of the National Job Training Partnership Act Study analyzed by Bloom, Orr, Bell, Cave, Doolittle, Lin, and Bos (1997).

### 1.1. *Related Literature*

Our paper contributes to a growing literature on statistical treatment rules in econometrics, including Manski (2004), Dehejia (2005), Hirano and Porter (2009), Stoye (2009, 2012), Chamberlain (2011), Bhattacharya and Dupas (2012), Tetenov (2012), and Kasy (2017). Manski (2004) proposed to assess the welfare properties of statistical treatment rules by their maximum regret and derived finite-sample regret bounds for *Conditional Empirical Success (CES)* rules. CES rules take a finite partition of the covariate space and, separately for each set in this partition, assign the treatment that yields the highest sample average outcome. CES rules can be viewed as a type of EWM rules for which  $\mathcal{G}$  consists of all unions of the sets in the partition and the empirical welfare criterion uses the sample propensity score. Manski showed that with the partition fixed, their welfare regret converges to zero at least at  $n^{-1/2}$  rate. We show that this rate holds for a broader class of EWM rules and that it cannot be improved uniformly without additional restrictions on  $P$ .

Stoye (2009) showed that in the absence of ex ante restrictions on how outcome distributions vary with covariates, finite-sample minimax regret is attained by rules that take the finest partition of the covariate space and operate independently for each covariate value. This important result implies that with continuous covariates, minimax regret does not converge to zero with sample size because the first-best treatment rule may be arbitrarily “wiggly” and difficult to approximate from countable data. Our approach does not give rise to Stoye’s non-convergence result because we restrict the complexity of  $\mathcal{G}$  and define regret relative to the maximum attainable welfare in  $\mathcal{G}$  instead of the unconstrained first-best welfare. However, we do not derive exact finite-sample minimax regret rules in the more complex setting of our paper.

Treatment choice has substantial similarities with *classification*. In a binary classification problem, the researcher observes a random sample  $(Y_i, X_i)$ , where  $Y_i \in \{-1, +1\}$  denotes which of two classes an observation belongs to. A *classifier* aims to predict the correct classification of future observations based on observed covariates  $X$ . A treatment rule similarly seeks to “classify” individuals into those who should and should not be treated based on their covariates. Treatment choice, however, differs from classification in a few significant ways: (1) observed outcomes can be real-valued rather than binary, (2) we only observe one of the two potential outcomes, and not the correct classification of individuals in the sample, (3) policy settings often impose constraints on practicable treatment rules or on the proportion of the population that could be treated.

The earliest works noting the connection between treatment choice and classification are Zdrozny (2003) and Beygelzimer and Langford (2009). They proposed algorithms that transform a sample from a treatment choice problem into a sample from a standard binary classification problem. Treatment rules could be then generated using any existing binary classification algorithm. Beygelzimer and Langford (2009) showed that the welfare loss of any treatment rule  $G$  with respect to the first-best ( $W(G_{\text{FB}}^*) - W(G)$ ) is bounded

above by a classification regret with respect to the first-best classifier. It implies that if the regret of the classification algorithm converges to zero, then consistency of  $W(\hat{G})$  to  $W(G_{\text{FB}}^*)$  holds. They did not consider any restrictions on  $\mathcal{G}$  and did not study the welfare loss convergence rates. Instead, we consider maximizing  $W_n(G)$  over a constrained class of policies without converting it into a classification problem.

The idea of optimizing the sample analog of a population decision problem is known as the *Empirical Risk Minimization (ERM) Principle* in classification (see Vapnik (1998) and references therein). The similarity between treatment choice and classification allows us to draw on recent results by Devroye, Györfi, and Lugosi (1996), Tsybakov (2004), Massart and Nédélec (2006), Audibert and Tsybakov (2007), and Kerkycharian, Tsybakov, Temlyakov, Picard, and Koltchinskii (2014), among others. We extend these convergence rate results for ERM classifiers to the treatment choice problem, accommodating the differences between classification and treatment choice and addressing issues specific to treatment choice. Establishing uniform convergence rates of the welfare regret of the EWM rule and its minimax rate optimality constitute the main theoretical contributions of this paper.

The analysis of *individualized treatment rules* has also received considerable attention in biostatistics. Qian and Murphy (2011) proposed a plug-in approach using  $E(Y_d|X)$  estimated by penalized least squares. They derived welfare convergence rate of  $n^{-1/2}$  or better (with a margin condition), assuming that  $E(Y_d|X)$  is well approximated by a sparse representation. Zhao, Zeng, Rush, and Kosorok (2012) proposed estimation of the treatment rule using a Support Vector Machine. This approach substitutes the EWM treatment choice objective function by a convex surrogate. They derived the welfare convergence rates that depend on the dimension of the covariates, similarly to nonparametric plug-in rules. These approaches are computationally attractive but cannot be used to choose from a constrained set of treatment rules or under a capacity constraint. Dudík, Langford, and Li (2011) and Zhang, Tsiatis, Laber, and Davidian (2012) considered maximizing a doubly-robust estimate of the welfare over a set of policies and showed by simulation that this approach outperforms the  $e$ -hybrid EWM approach in terms of welfare. Athey and Wager (2017) analytically characterized advantages of the doubly-robust approach by showing an improved constant term in the welfare regret upper bounds.

Several works in econometrics consider the *plug-in* approach to treatment choice using estimated regression equations,

$$\hat{G}_{\text{plug-in}} = \{x : \hat{\tau}^m(x) \geq 0\}, \quad \hat{\tau}^m(x) = \hat{m}_1(x) - \hat{m}_0(x), \quad (1.13)$$

where  $\hat{m}_d(x)$  is a parametric or a nonparametric estimator of  $E(Y_d|X = x)$ . Hirano and Porter (2009) established local asymptotic minimax optimality of plug-in rules for parametric and semiparametric models of treatment response. Under an aggregate budget constraint, Bhattacharya and Dupas (2012) considered nonparametric plug-in rules with propensity score weighted estimators of the regression equations and derived some of their properties. Armstrong and Shen (2015) considered statistical inference for the first-best decision rule  $G_{\text{FB}}^*$  from the perspective of inference for conditional moment inequalities. Empirical researchers often assess who should be treated by stratifying the population on an estimated predictor of  $Y_0$ , which leads to biased treatment effect estimates (Abadie, Chingos, and West (2017)). Kasy (2016) considered estimation of a partially-ordered welfare ranking over treatment assignment policies with a set-identified welfare criterion.

To assess treatment effect heterogeneity, estimation and inference for conditional treatment effects based on parametric or nonparametric regressions are often reported, but

the stylized output of statistical inference (e.g., confidence intervals,  $p$ -values) fails to offer the policy maker a direct guidance on what treatment rule to follow. In contrast, our EWM approach offers the policy maker a specific treatment assignment rule designed to maximize social welfare.

A treatment assignment rule could also be obtained by specifying a prior distribution for  $P$  and solving for a Bayes decision rule (see Dehejia (2005) and Chamberlain (2011) for Bayesian approaches to the treatment choice problem). Kasy (2017) proposed a non-parametric Bayesian approach to policy estimation for a range of public policy applications. In contrast to the Bayesian approach, the EWM approach utilizes only the empirical distribution of the data and does not require a prior distribution over the data generating processes.

Elliott and Lieli (2013) and Lieli and White (2010) also proposed maximizing the sample analog of a utilitarian decision criterion similar to EWM. They considered the problem of forecasting binary outcomes based on observations of  $(Y_i, X_i)$ , as in Manski and Thompson (1989), where a forecast leads to a binary decision.

## 2. THEORETICAL PROPERTIES OF EWM

### 2.1. Setup and Assumptions

Throughout our investigation of theoretical properties of EWM, we maintain the following assumptions.

ASSUMPTION 2.1: (UCF) Unconfoundedness:  $(Y_1, Y_0) \perp D|X$ .

(BO) Bounded Outcomes: *There exists  $M < \infty$  such that the support of outcome variable  $Y$  is contained in  $[-M/2, M/2]$ .*

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

(VC) VC-class: *A class of decision sets  $\mathcal{G}$  has a finite VC-dimension<sup>3</sup>  $v < \infty$  and is countable.<sup>4</sup>*

The assumption of unconfoundedness (selection on observables) holds if data are obtained from an experimental study with a randomized treatment assignment. In observational studies, unconfoundedness is a non-testable and often controversial assumption. Our analysis could be applied to the observational studies in which unconfoundedness is credible. The second assumption (BO) implies boundedness of the treatment effects, that is,

$$P_X(|\tau(X)| \leq M) = 1,$$

<sup>3</sup>Let  $\mathbf{x}^l \equiv \{x_1, \dots, x_l\}$  be a finite set with  $l \geq 1$  points in  $\mathcal{X}$ . Given a class of subsets  $\mathcal{G}$  in  $\mathcal{X}$ , define  $N(\mathbf{x}^l) = |\{\mathbf{x}^l \cap G : G \in \mathcal{G}\}|$  to be the number of different subsets of  $\mathbf{x}^l$  picked out by  $G \in \mathcal{G}$ . The VC-dimension  $v \geq 1$  of  $\mathcal{G}$  is defined by the largest  $l$  such that  $\sup_{\mathbf{x}^l} N(\mathbf{x}^l) = 2^l$  holds (Vapnik (1998)). The VC-dimension is commonly used to measure the complexity of a class of sets in the statistical learning literature (see Vapnik (1998), Dudley (1999, Chapter 4), and van der Vaart and Wellner (1996) for extensive discussions). Note that the VC-dimension is smaller by one compared to the VC-index used to measure the complexity of a class of sets in the empirical process theory; see, for example, van der Vaart and Wellner (1996).

<sup>4</sup>Countability of  $\mathcal{G}$  is imposed to simplify measurability issues in proving our theoretical results. In Examples 2.1–2.3 below, we formulate  $\mathcal{G}$  to be uncountable, whereas any practical implementation will only use a countable subset of  $\mathcal{G}$  in search of the EWM rule.

where  $P_X$  is the marginal distribution of  $X$  and  $\tau(\cdot)$  is the conditional treatment effect  $\tau(X) = E(Y_1 - Y_0|X)$ . Implementing EWM does not require knowledge of  $M$  and this assumption is imposed mainly for analytical convenience. The third assumption (SO) is a standard assumption in the treatment effect literature. It is satisfied in randomized controlled trials by design, but it may be violated in observational studies if almost all the individuals are in the same group (treatment or control) for some values of  $X$ . We let  $\mathcal{P}(M, \kappa)$  denote the class of distributions of  $(Y_0, Y_1, D, X)$  that satisfy Assumption 2.1 (UCF), (BO), and (SO).

The fourth assumption (VC) restricts the complexity of the class of candidate treatment rules  $\mathcal{G}$  in terms of its VC-dimension. If  $X$  has a finite support, then the VC-dimension  $v$  of any class  $\mathcal{G}$  does not exceed the number of support points. If some of  $X$  is continuously distributed, Assumption 2.1 (VC) requires  $\mathcal{G}$  to be smaller than the Borel  $\sigma$ -algebra of  $\mathcal{X}$ . The following examples illustrate several practically relevant classes of the feasible treatment rules satisfying Assumption 2.1 (VC).

**EXAMPLE 2.1—Linear Eligibility Score:** Suppose that a feasible assignment rule is constrained to those that assign the treatment according to an eligibility score. By the eligibility score, we mean a scalar-valued function of the individual’s observed characteristics that determines whether one receives the treatment based on whether the eligibility score exceeds a certain threshold. The main objective of data analysis is therefore to construct an eligibility score that yields a welfare-maximizing treatment rule. Specifically, we assume that the eligibility score is constrained to being linear in a subvector of  $x \in \mathbb{R}^{d_x}$ ,  $x_{\text{sub}} \in \mathbb{R}^{d_{\text{sub}}}$ ,  $d_{\text{sub}} \leq d_x$ . The class of decision sets generated by Linear Eligibility Scores (LES) is defined as

$$\mathcal{G}_{\text{LES}} \equiv \{ \{x \in \mathbb{R}^{d_x} : \beta_0 + x_{\text{sub}}^T \beta_{\text{sub}} \geq 0\} : (\beta_0, \beta_{\text{sub}}^T) \in \mathbb{R}^{d_{\text{sub}}+1} \}. \tag{2.1}$$

We accordingly obtain an EWM assignment rule by maximizing

$$W_n(\beta) \equiv E_n \left[ \frac{Y_i D_i}{e(X_i)} \cdot 1\{\beta_0 + X_{\text{sub},i}^T \beta_{\text{sub}} \geq 0\} + \frac{Y_i(1 - D_i)}{1 - e(X_i)} \cdot 1\{\beta_0 + X_{\text{sub},i}^T \beta_{\text{sub}} < 0\} \right]$$

in  $\beta = (\beta_0, \beta_{\text{sub}}^T) \in \mathbb{R}^{d_{\text{sub}}+1}$ . It is well known that the class of half-spaces spanned by  $(\beta_0, \beta_{\text{sub}}^T) \in \mathbb{R}^{d_{\text{sub}}+1}$  has the VC-dimension  $v = d_{\text{sub}} + 1$ , so the requirement of finite VC-dimension in Assumption 2.1 (VC) holds. In Appendix C of Kitagawa and Tetenov (2018a), we discuss how to compute  $\hat{G}_{\text{EWM}}$  when the class of decision sets is given by  $\mathcal{G}_{\text{LES}}$ . A plug-in rule based on a parametric linear regression also selects a treatment rule from  $\mathcal{G}_{\text{LES}}$ , but their welfare does not converge to the maximum welfare  $W_{\mathcal{G}_{\text{LES}}}^*$  if the regression equations are misspecified, whereas the welfare of  $\hat{G}_{\text{EWM}}$  always does (as shown in Theorem 2.1 below).

**EXAMPLE 2.2—Generalized Eligibility Score:** Let  $f_j(\cdot)$ ,  $j = 1, \dots, m$ , and  $g(\cdot)$  be known functions of  $x \in \mathbb{R}^{d_x}$ . Consider a class of assignment rules generated by Generalized Eligibility Scores (GES),

$$\mathcal{G}_{\text{GES}} \equiv \left\{ \left\{ x \in \mathbb{R}^{d_x} : \sum_{j=1}^m \beta_j f_j(x) \geq g(x) \right\}, (\beta_1, \dots, \beta_m) \in \mathbb{R}^m \right\}.$$

The class of decision sets  $\mathcal{G}_{\text{GES}}$  generalizes the linear eligibility score rules (2.1), as it allows for eligibility scores that are nonlinear in  $x$ , that is,  $\mathcal{G}_{\text{GES}}$  can accommodate decision sets



that partition the space of covariates by nonlinear boundaries. It can be shown that  $\mathcal{G}_{\text{GES}}$  has the VC-dimension  $v = m + 1$  (Theorem 4.2.1 in Dudley (1999)).

**EXAMPLE 2.3—Intersection Rule of Multiple Eligibility Scores:** Consider a situation where there are  $L \geq 2$  eligibility scores. Let  $\mathcal{G}_{\text{GES},l}, l = 1, \dots, L$ , be classes of decision sets such that each of them is generated by contour sets of the  $l$ th eligibility score. Suppose that a feasible decision rule is constrained to those that assign the treatment if the individual has all the  $L$  eligibility scores exceeding thresholds. In this case, the class of decision sets is constructed by the intersections,  $\mathcal{G} \equiv \bigcap_{l=1}^L \mathcal{G}_{\text{GES},l} = \{\bigcap_{l=1}^L G_l : G_l \in \mathcal{G}_{\text{GES},l}, l = 1, \dots, L\}$ . An intersection of a finite number of VC-classes is a VC-class with a finite VC-dimension (Theorem 4.5.4 in Dudley (1999)). We can also consider a class of treatment rules that assigns a treatment if at least one of the  $L$  eligibility scores exceeds a threshold. In this case, instead of intersections, the class of decision sets is formed by the unions of  $\{\mathcal{G}_{\text{GES},l}, l = 1, \dots, L\}$ , which is also known to have a finite VC-dimension (Theorem 4.5.4 in Dudley (1999)).

### 2.2. Uniform Rate Optimality of EWM

To analyze statistical performance of EWM rules, we focus on a non-asymptotic upper bound of the worst-case welfare loss  $\sup_{P \in \mathcal{P}(M, \kappa)} E_{P^n} [W_{\mathcal{G}}^* - W(\hat{G}_{\text{EWM}})]$  and examine how it depends on sample size  $n$  and VC-dimension  $v$ . This finite-sample upper bound allows us to assess the uniform convergence rate of the welfare and to examine how richness (complexity) of the class of candidate decision rules affects the worst-case performance of EWM. The main reason that we focus on the uniform convergence rate rather than a pointwise convergence rate is that the pointwise convergence rate of the welfare loss can vary depending on a feature of the data distribution and fails to provide a guaranteed learning rate of an optimal policy when no additional assumption, other than Assumption 2.1, is available.

For heuristic illustration of the derivation of the uniform convergence rate, consider the following inequality, which holds for any  $\tilde{G} \in \mathcal{G}$ :

$$\begin{aligned} W(\tilde{G}) - W(\hat{G}_{\text{EWM}}) &= W(\tilde{G}) - W_n(\hat{G}_{\text{EWM}}) + W_n(\hat{G}_{\text{EWM}}) - W(\hat{G}_{\text{EWM}}) \\ &\leq W(\tilde{G}) - W_n(\tilde{G}) + \sup_{G \in \mathcal{G}} |W_n(G) - W(G)| \\ &\quad (\because W_n(\hat{G}_{\text{EWM}}) \geq W_n(\tilde{G})) \\ &\leq 2 \sup_{G \in \mathcal{G}} |W_n(G) - W(G)|. \end{aligned} \tag{2.2}$$

Since it applies to  $W(\tilde{G})$  for all  $\tilde{G}$ , it also applies to  $W_{\mathcal{G}}^* = \sup W(\tilde{G})$ :

$$W_{\mathcal{G}}^* - W(\hat{G}_{\text{EWM}}) \leq 2 \sup_{G \in \mathcal{G}} |W_n(G) - W(G)|. \tag{2.3}$$

Therefore, the expected welfare loss can be bounded uniformly in  $P$  by a distribution-free upper bound of  $E_{P^n}(\sup_{G \in \mathcal{G}} |W_n(G) - W(G)|)$ . Since  $W_n(G) - W(G)$  can be seen as the centered empirical process indexed by  $G \in \mathcal{G}$ , an application of the existing moment inequality for the supremum of centered empirical processes indexed by a VC-class yields the following distribution-free upper bound. A proof, which closely follows the proofs of

Theorems 1.16 and 1.17 in Lugosi (2002) in the classification problem, is given in Section A.2 of Kitagawa and Tetenov (2018a).

**THEOREM 2.1:** *Under Assumption 2.1, we have*

$$\sup_{P \in \mathcal{P}(M, \kappa)} E_{P^n} [W_{\mathcal{G}}^* - W(\hat{G}_{\text{EWM}})] \leq C_1 \frac{M}{\kappa} \sqrt{\frac{v}{n}},$$

where  $C_1$  is a universal constant defined in Lemma A.4 in Kitagawa and Tetenov (2018a).

This theorem shows that the convergence rate of the worst-case welfare loss for the EWM rule is no slower than  $n^{-1/2}$ . The upper bound is increasing in the VC-dimension of  $\mathcal{G}$ , implying that, as the candidate treatment assignment rules become more complex in terms of VC-dimension,  $\hat{G}_{\text{EWM}}$  tends to overfit the data in the sense that the distribution of regret  $W_{\mathcal{G}}^* - W(\hat{G}_{\text{EWM}})$  is more and more dispersed, and, with  $n$  fixed, this overfitting results in inflating the average welfare regret.<sup>5</sup>

The next theorem concerns a universal lower bound of the worst-case average welfare loss. It shows that no data-based treatment choice rule can have a uniform convergence rate faster than  $n^{-1/2}$ .

**THEOREM 2.2:** *Suppose that Assumption 2.1 holds. For any treatment choice rule  $\hat{G}$  as a function of  $(Z_1, \dots, Z_n)$ , it holds that*

$$\sup_{P \in \mathcal{P}(M, \kappa)} E_{P^n} [W_{\mathcal{G}}^* - W(\hat{G})] \geq 2^{-1} \exp\{-2\sqrt{2}\} M \sqrt{\frac{v}{n}} \quad \text{for all } n \geq 16v.$$

We derive this lower bound by bounding below the worst-case welfare regret by the risk of a parametric Bayes decision problem (i.e., a prior that only supports a parametric subclass  $\mathcal{P}^* \subset \mathcal{P}(M, \kappa)$ ) and maximizing the Bayes risk over  $P \in \mathcal{P}^*$ . A similar proof technique appears in Devroye and Lugosi (1995) in their regret lower bound analysis of ERM classifiers.

This theorem, combined with Theorem 2.1, implies that  $\hat{G}_{\text{EWM}}$  is *minimax rate optimal* over the class of data generating processes  $\mathcal{P}(M, \kappa)$ , since the convergence rate of the upper bound of  $\sup_{P \in \mathcal{P}(M, \kappa)} E_{P^n} [W_{\mathcal{G}}^* - W(\hat{G}_{\text{EWM}})]$  agrees with the convergence rate of the universal lower bound. Accordingly, we can conclude that no other data-driven procedure for obtaining a treatment choice rule can outperform  $\hat{G}_{\text{EWM}}$  in terms of the uniform convergence rate over  $\mathcal{P}(M, \kappa)$ . It is worth noting that the rate lower bound is uniform in  $P$  and does not apply pointwise. Theorem 2.3 shows that EWM rules have faster convergence rates for *some* distributions. It is also possible that  $E_{P^n} [W(\hat{G})] > W_{\mathcal{G}}^*$  for some pairs of  $\hat{G}$  and  $P$ , but it can never hold for all distributions in  $\mathcal{P}(M, \kappa)$ .<sup>6</sup>

<sup>5</sup>Note that  $W_{\mathcal{G}}^*$  weakly increases if a more complex class  $\mathcal{G}$  is chosen. Our welfare loss criterion is defined for a specific class  $\mathcal{G}$  and does not capture the potential gain in the maximal welfare from the choice of a more complex  $\mathcal{G}$ .

<sup>6</sup>For example, if  $\hat{G}$  is a nonparametric plug-in rule and the first-best decision rule  $G_{\text{FB}}^*$  for distribution  $P$  does not belong to  $\mathcal{G}$ , then the welfare of  $\hat{G}$  will exceed  $W_{\mathcal{G}}^*$  in sufficiently large samples. However, the uniform lower bound still applies because there exist other distributions for which  $E_{P^n} W(\hat{G}) \leq W_{\mathcal{G}}^* - (n^{-1/2} \text{ bound})$  for the same sample size.

REMARK 2.1: Capacity or budget constraints that restrict the proportion of the target population that could be assigned to treatment exist in various treatment choice problems. An attractive feature of the empirical welfare maximization method is the simplicity of incorporating these constraints in the estimation of a treatment choice rule.

Assume that the proportion of the target population that could receive treatment 1 cannot exceed  $K \in (0, 1)$ . If the population distribution of covariates  $P_X$  were known, maximization of the empirical welfare criterion could be simply restricted to sets in class  $\mathcal{G}$  that satisfy the capacity constraint  $\mathcal{G}^K \equiv \{G \in \mathcal{G} : P_X(G) \leq K\}$ . Being a subset of  $\mathcal{G}$ , the class of sets  $\mathcal{G}^K$  has the same complexity as  $\mathcal{G}$  (or lower), and Theorem 2.1 could be applied simply by replacing  $\mathcal{G}$  with  $\mathcal{G}^K$ .

When  $P_X$  is unknown, it is not guaranteed with certainty that estimated treatment rule  $\hat{G}$  satisfies the capacity constraint. To evaluate the welfare in this setting, we assume that if the treatment rule  $G$  violates the capacity constraint,  $P_X(G) > K$ , then the scarce treatment is randomly allocated (“rationed”) to a fraction  $\frac{K}{P_X(G)}$  of the assigned recipients with  $X \in G$  independently of  $(X, Y_0, Y_1)$ .<sup>7</sup> If  $G$  does not violate the capacity constraint, then there is no rationing and all recipients with covariates  $X \in G$  receive treatment 1. This allows us to clearly define the capacity-constrained welfare of the treatment rule indexed by any subset  $G \subset \mathcal{X}$  of the covariate space as

$$W^K(G) \equiv E_P \left[ \left[ Y_1 \cdot \min \left\{ 1, \frac{K}{P_X(G)} \right\} + Y_0 \cdot \left( 1 - \min \left\{ 1, \frac{K}{P_X(G)} \right\} \right) \right] \cdot 1\{X \in G\} + Y_0 \cdot 1\{X \notin G\} \right].$$

Then the capacity-constrained welfare gain of the treatment rule  $G$  relative to the no-treatment policy is given by

$$V^K(G) \equiv W^K(G) - W^K(\emptyset) = \min \left\{ 1, \frac{K}{P_X(G)} \right\} \cdot E_P[\tau(X)1\{X \in G\}].$$

Observe that rationing dilutes the effect of treatment rules that violate the capacity constraint and we take into account this effect on welfare. We hence propose a treatment rule that maximizes the empirical analog of the capacity-constrained welfare gain  $V^K(G)$  (and, hence, welfare):

$$\hat{G}^K \equiv \arg \max_{G \in \mathcal{G}} V_n^K(G), \quad (2.4)$$

where

$$V_n^K(G) \equiv \min \left\{ 1, \frac{K}{P_{X,n}(G)} \right\} \cdot E_n \left[ \left( \frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1 - D_i)}{1 - e(X_i)} \right) \cdot 1\{X_i \in G\} \right],$$

and  $P_{X,n}$  is the empirical probability distribution of  $(X_1, \dots, X_n)$ . Theorem D.1 of Kitagawa and Tetenov (2018b) shows that similarly to Theorem 2.1, the expected welfare of  $\hat{G}^K$  converges to the maximum at least at  $n^{-1/2}$  rate.

REMARK 2.2: Empirical Welfare Maximization method can be adapted to situations where a target population shares the conditional treatment effect with the sampled population, but differs in the distribution of covariates  $X$ . Let  $\rho(x) \equiv p_X^T(x)/p_X(x)$  be the

<sup>7</sup>In comparison, nonparametric plug-in treatment rules proposed by Bhattacharya and Dupas (2012) are only required to satisfy the capacity constraint on average over repeated data samples.

density ratio of the marginal distributions of  $X$ , where  $p_X$  and  $p_X^T$  are those of the sampled population and the target population. Assume  $\rho(x) \leq \bar{\rho} < \infty$  for all  $x$ . The welfare gain of treatment rule  $G$  on the target population can be written as

$$V^T(G) \equiv \int_X \tau(x) 1\{x \in G\} \rho(x) dP_X(x).$$

Since the first-best treatment rule  $G_{FB}^* = 1\{x : \tau(x) \geq 0\}$  is the same in the sampled and the target populations, if  $G_{FB}^* \in \mathcal{G}$ , we could directly apply the EWM rule computed for the sampled population to the target population. In contrast, if the first-best policy is not feasible ( $G_{FB}^* \notin \mathcal{G}$ ), the second-best policies for the sampled and the target populations are generally different, and the welfare of treatment rules proposed in the previous sections does not generally converge to the second-best in the target population  $\sup_{G \in \mathcal{G}} V^T(G)$ .

The second-best in the target population could be obtained by reweighting the argument of the EWM problem by the density ratio  $\rho(X_i)$ :

$$\hat{G}_{EWM}^T \in \arg \max_{G \in \mathcal{G}} E_n \left[ \left( \frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1 - D_i)}{1 - e(X_i)} \right) \cdot \rho(X_i) \cdot 1\{X_i \in G\} \right]. \tag{2.5}$$

As an extension of Theorem 2.1, the welfare loss of the reweighted EWM rule in the target population can be shown to converge to zero at least at  $n^{-1/2}$  rate.

REMARK 2.3: The EWM rule (1.7) is invariant to multiplying all outcomes  $Y_i$  by a positive constant, but is not invariant to adding a constant. If all  $Y_i$ 's are replaced by  $Y_i + c$ , the welfare estimate  $W_n(G)$  changes by  $c \cdot E_n[\frac{D_i}{e(X_i)} \cdot 1\{X_i \in G\} + \frac{1-D_i}{1-e(X_i)} \cdot 1\{X_i \notin G\}]$ . This difference converges to  $c$  for every  $G$ , but its value varies with  $G$  in finite samples. This could pose problems for applied work because the researcher has some room to change the treatment assignment rule by changing the coding of the outcome variable. We propose a simple modification of the EWM rule invariant to positive affine transformations of outcomes. Denote by  $Y_i^{dm} \equiv Y_i - E_n[Y_i]$  the outcomes demeaned by their sample mean; they are invariant to changing  $Y_i$ 's by a constant. Then the demeaned EWM rule

$$\hat{G}_{EWM}^{dm} \in \arg \max_{G \in \mathcal{G}} W_n^{dm}(G), \quad \text{where} \tag{2.6}$$

$$W_n^{dm}(G) \equiv E_n \left[ \frac{Y_i^{dm} D_i}{e(X_i)} \cdot 1\{X_i \in G\} + \frac{Y_i^{dm}(1 - D_i)}{1 - e(X_i)} \cdot 1\{X_i \notin G\} \right], \tag{2.7}$$

is invariant to positive affine transformations of  $Y_i$ .

The rate result of Theorem 2.1 also holds for demeaned EWM rule  $\hat{G}_{EWM}^{dm}$ , as we show in Section D.2 of Kitagawa and Tetenov (2018b).<sup>8</sup> In our empirical application, we compute the demeaned EWM treatment rules.

REMARK 2.4: In Section D.3 of Kitagawa and Tetenov (2018b), we discuss how the EWM approach could be applied with more than two treatments. The rate result of Theorem 2.1 holds for multiple treatments if treatment assignment covariate subsets for each treatment belong to a VC-class.

<sup>8</sup>In our simulations, demeaned EWM never performed much worse than standard EWM in terms of welfare. Demeaned EWM performed much better in cases where  $E[Y]$  was very far from zero.

2.3. Rate Improvement by Margin Assumption

The welfare loss upper bounds obtained in Theorem 2.1 can indeed tighten up and the uniform convergence rate can improve, as we further constrain the class of data generating processes. In this section, we investigate (i) what feature of data generating processes can affect the upper bound on the welfare loss of the EWM rule, and (ii) whether or not the EWM rule remains minimax rate optimal even under the additional constraints. For this goal, we consider imposing the following two assumptions.

ASSUMPTION 2.2: (FB) Correct Specification: *The first-best treatment rule  $G_{FB}^*$  defined in (1.8) belongs to the class of candidate treatment rules  $\mathcal{G}$ .*

(MA) Margin Assumption: *There exist constants  $0 < \eta \leq M$  and  $0 < \alpha < \infty$  such that*

$$P_X(|\tau(X)| \leq t) \leq (t/\eta)^\alpha, \quad \forall 0 \leq t \leq \eta,$$

where  $M < \infty$  is the constant as defined in Assumption 2.1 (BO).

The assumption of correct specification means that the class of feasible policy rules specified by  $\mathcal{G}$  contains an unconstrained first-best treatment rule  $G_{FB}^*$ . This assumption is plausible if, for instance, the policy maker’s specification of  $\mathcal{G}$  is based on a credible assumption about the shape of the contour set  $\{x : \tau(x) \geq 0\}$ . This assumption can be, on the other hand, restrictive if the specification of  $\mathcal{G}$  comes from some exogenous constraints for feasible policy rules, as in the case of Example 2.1.

The second assumption (MA) concerns the way in which the distribution of conditional treatment effect  $\tau(X)$  behaves in the neighborhood of  $\tau(X) = 0$ . A similar assumption has been considered in the literature on classification analysis (Mammen and Tsybakov (1999), Tsybakov (2004), among others), and we borrow the term “margin assumption” from Tsybakov (2004). Parameters  $\eta$  and  $\alpha$  characterize the size of population with the conditional treatment effect close to the margin  $\tau(X) = 0$ . *Smaller  $\eta$  and  $\alpha$  imply that more individuals can concentrate in a neighborhood of  $\tau(X) = 0$ .* The next examples illustrate this interpretation of  $\eta$  and  $\alpha$ .

EXAMPLE 2.4: Suppose that  $X$  contains a continuously distributed covariate and that the conditional treatment effect  $\tau(X)$  is continuously distributed. If the probability density function of  $\tau(X)$  is bounded from above by  $p_\tau < \infty$ , then the margin assumption holds with  $\alpha = 1$  and  $\eta = (2p_\tau)^{-1}$ .

EXAMPLE 2.5: Suppose that  $X$  is a scalar and follows the uniform distribution on  $[-1/2, 1/2]$ . Specify the conditional treatment effect to be  $\tau(X) = (-X)^3$ . In this specification,  $\tau(X)$  “flats out” at  $X = 0$ , and accordingly, the density function of  $\tau(X)$  is unbounded in the neighborhood of  $\tau(X) = 0$ . This specification leads to  $P_X(|\tau(X)| \leq t) = 2t^{1/3}$ , so the margin assumption holds with  $\alpha = 1/3$  and  $\eta = 1/8$ .

EXAMPLE 2.6: Suppose that the distribution of  $X$  is the same as in Example 2.5. Let  $h > 0$  and specify  $\tau(X)$  as

$$\tau(X) = \begin{cases} X - h & \text{for } X \leq 0, \\ X + h & \text{for } X > 0. \end{cases}$$

This  $\tau(X)$  is discontinuous at  $X = 0$ , and the distribution of  $\tau(X)$  has zero probability around the margin of  $\tau(X) = 0$ . It holds that

$$P_X(|\tau(X)| \leq t) = \begin{cases} 0 & \text{for } t \leq h, \\ 2(t - h) & \text{for } h < t \leq \frac{1}{2} + h. \end{cases}$$

By setting  $\eta = h$ , the margin condition holds for arbitrarily large  $\alpha$ . In general, if the distribution of  $\tau(X)$  has a gap around the margin of  $\tau(X) = 0$ , the margin condition holds with arbitrarily large  $\alpha$ .

From now on, we denote the class of  $P$  satisfying Assumptions 2.1 and 2.2 by  $\mathcal{P}_{\text{FB}}(M, \kappa, \eta, \alpha)$ .<sup>9</sup> The next theorem provides the upper bound of the welfare loss of the EWM rule when a class of data distributions is constrained to  $\mathcal{P}_{\text{FB}}(M, \kappa, \eta, \alpha)$ .

**THEOREM 2.3:** *Under Assumptions 2.1 and 2.2,*

$$\sup_{P \in \mathcal{P}_{\text{FB}}(M, \kappa, \eta, \alpha)} E_{P^n} [W(G_{\text{FB}}^*) - W(\hat{G}_{\text{EWM}})] \leq c \left( \frac{v}{n} \right)^{\frac{1+\alpha}{2+\alpha}}$$

*holds for all  $n$ , where  $c$  is a positive constant that depends only on  $M, \kappa, \eta$ , and  $\alpha$ .*

Similarly to Theorem 2.1, the presented welfare loss upper bound is non-asymptotic, and it is valid for every sample size. Our derivation of this theorem can be seen as an extension of the finite-sample risk bound for the classification error shown in Theorem 2 of Massart and Nédélec (2006). Our rate upper bound is consistent with the uniform convergence rate of the classification risk of the empirical risk-minimizing classifier shown in Theorem 1 of Tsybakov (2004).<sup>10</sup> This coincidence is somewhat expected, given that the empirical welfare criterion that the EWM rule maximizes resembles the empirical classification risk in the classification problem.

The next theorem shows that the uniform convergence rate of  $n^{-\frac{1+\alpha}{2+\alpha}}$  obtained in Theorem 2.3 attains the minimax rate lower bound, implying that any treatment choice rule  $\hat{G}$  based on data (including  $\hat{G}_{\text{EWM}}$ ) cannot attain a uniform convergence rate faster than  $n^{-\frac{1+\alpha}{2+\alpha}}$ . This means that the EWM rule remains rate optimal even when the class of data generating processes is constrained additionally by Assumption 2.2.<sup>11</sup>

**THEOREM 2.4:** *Suppose Assumptions 2.1 and 2.2 hold. Assume that the VC-dimension of  $\mathcal{G}$  satisfies  $v \geq 2$ . Then, for any treatment choice rule  $\hat{G}$  as a function of  $(Z_1, \dots, Z_n)$ , it holds that*

$$\sup_{P \in \mathcal{P}_{\text{FB}}(M, \kappa, \eta, \alpha)} E_{P^n} [W(G_{\text{FB}}^*) - W(\hat{G})] \geq 2^{-1} \exp\{-2\sqrt{2}\} M^{\frac{2(1+\alpha)}{2+\alpha}} \eta^{-\frac{\alpha}{2+\alpha}} \left( \frac{v-1}{n} \right)^{\frac{1+\alpha}{2+\alpha}}$$

*for all  $n \geq \max\{(M/\eta)^2, 4^{2+\alpha}\}(v-1)$ .*

<sup>9</sup>Note that  $P_{\text{FB}}(M, \kappa, \eta, \alpha)$  depends on the set of feasible treatment rules  $\mathcal{G}$  via Assumption 2.2 (FB).

<sup>10</sup>Tsybakov (2004) defined the complexity of the decision sets  $\mathcal{G}$  in terms of the growth coefficient  $\rho$  of the bracketing number of  $\mathcal{G}$ . We control complexity of  $\mathcal{G}$  in terms of the VC-dimension, which corresponds to Tsybakov's growth coefficient  $\rho$  being arbitrarily close to zero.

<sup>11</sup>Assumption 2.2 rules out data generating processes with  $P_X(\tau(X) = 0) > 0$ , which can constitute focal null hypotheses (often  $P_X(\tau(X) = 0) = 1$ ) in program evaluation studies. A practical implication of the refined minimax rate result shown in this section is that the EWM rule remains a recommended choice even when we know ex ante that there is substantial effect heterogeneity in  $x$  and  $\tau(x) \neq 0$  for most  $x$ .

The following remarks summarize some analytical insights associated with Theorems 2.1–2.4.

REMARK 2.5: The convergence rates of the worst-case EWM welfare loss obtained by Theorems 2.1 and 2.3 highlight how margin coefficient  $\alpha$  influences the uniform performance of the EWM rule. Higher  $\alpha$  improves the welfare loss convergence rate of EWM, and the convergence rate approaches  $n^{-1}$  in an extreme case, where the distribution of  $\tau(X)$  has a gap around  $\tau(X) = 0$ . As fewer individuals are around the margin of  $\tau(X) = 0$ , we can attain the maximal welfare quicker. Conversely, as  $\alpha$  approaches zero (more individuals around the margin), the welfare loss convergence rate of EWM approaches  $n^{-1/2}$ , and it corresponds to the uniform convergence rate of Theorem 2.1.

REMARK 2.6: The upper bounds of welfare loss convergence rate shown in Theorems 2.1 and 2.3 are increasing in the VC-dimension of  $\mathcal{G}$ . Since they are valid at every  $n$ , we can allow the VC-dimension of the candidate treatment rules to grow with the sample size. For instance, if we consider a sequence of candidate decision sets  $\{\mathcal{G}_n : n = 1, 2, \dots\}$ , for which the VC-dimension grows with the sample size at rate  $n^\lambda$ ,  $0 < \lambda < 1$ , Theorems 2.1 and 2.3 imply that the welfare loss uniform convergence rate of the EWM rule slows down to  $n^{-\frac{1-\lambda}{2}}$  for the case without Assumption 2.2 and to  $n^{-(1-\lambda)\frac{(1+\alpha)}{2+\alpha}}$  for the case with Assumption 2.2.<sup>12</sup> Note that the welfare loss lower bounds shown in Theorems 2.2 and 2.4 have the VC-dimensions of the same order as in the corresponding upper bounds, so we can conclude that the EWM rule is also minimax rate optimal even in the situations where the complexity of  $\mathcal{G}$  grows with the sample size.

REMARK 2.7: Note that the welfare loss lower bounds of Theorems 2.2 and 2.4 are valid for any estimated treatment choice rule  $\hat{G}$  irrespective of whether  $\hat{G}$  is constrained to  $\mathcal{G}$  or not. Therefore, the nonparametric plug-in rule  $\hat{G}_{\text{plug-in}}$  defined in (1.13) is subject to the same lower bound.<sup>13</sup>

REMARK 2.8: Let  $\mathcal{P}_{\text{FB}}(M, \kappa)$  be the class of data generating processes that satisfy Assumption 2.1 and Assumption 2.2 (FB). A close inspection of the proofs of Theorems 2.1 and 2.2 given in Section A.2 of Kitagawa and Tetenov (2018a) shows that the same lower and upper bounds of Theorems 2.1 and 2.2 can be obtained even when  $\mathcal{P}(M, \kappa)$  is replaced with  $\mathcal{P}_{\text{FB}}(M, \kappa)$ . In this sense, Assumption 2.2 (MA) plays the main role in improving the welfare loss convergence rate.

## 2.4. Unknown Propensity Score

We have so far considered situations where the true propensity score is known. This would not be the case if the data were obtained from an observational study in which the assignment of treatment is not generally under the control of the experimenter. To cope with the unknown propensity score, this section considers two hybrids of the EWM approach and the parametric/nonparametric plug-in approach: the  $m$ -hybrid rule defined

<sup>12</sup>Note that for the case without Assumption 2.2 (FB), the maximal attainable welfare  $W_{\mathcal{G}}^*$  increases weakly as the complexity of  $\mathcal{G}$  grows. On the other hand, with Assumption 2.2 (FB), the set of data generating processes  $\mathcal{P}_{\text{FB}}(M, \kappa, \eta, \alpha)$  expands as the complexity of  $\mathcal{G}$  grows.

<sup>13</sup>Section D.4 of Kitagawa and Tetenov (2018b) discusses the welfare loss uniform convergence rate of the nonparametric plug-in rule.

in (1.11) and the  $e$ -hybrid rule defined in (1.12). The  $e$ -hybrid rule employs the trimming rule  $1\{\varepsilon_n \leq \hat{e}(X_i) \leq 1 - \varepsilon_n\}$  with a deterministic sequence  $\{\varepsilon_n : n = 1, 2, \dots\}$ , which we assume to converge to zero faster than some polynomial rate,  $\varepsilon_n \leq O(n^{-a})$ ,  $a > 0$ .<sup>14</sup>

The next condition concerns the convergence rate of the average estimation error of the conditional treatment effect estimators.

CONDITION 2.1: (m) ( $m$ -hybrid case): Let  $\hat{\tau}^m(x) = \hat{m}_1(x) - \hat{m}_0(x)$  be an estimator for the conditional treatment effect  $\tau(x) = m_1(x) - m_0(x)$ . For a class of data generating processes  $\mathcal{P}_m$ , there exists a sequence  $\psi_n \rightarrow \infty$  such that

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_m} \psi_n E_{P^n} \left[ \frac{1}{n} \sum_{i=1}^n |\hat{\tau}^m(X_i) - \tau(X_i)| \right] < \infty \tag{2.8}$$

holds.

(e) ( $e$ -hybrid case): Let  $\hat{\tau}_i^e = \left[ \frac{Y_i D_i}{\hat{e}(X_i)} - \frac{Y_i(1-D_i)}{1-\hat{e}(X_i)} \right] \cdot 1\{\varepsilon_n \leq \hat{e}(X_i) \leq 1 - \varepsilon_n\}$  be an estimator for  $\tau_i = \frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)}$ , where  $\hat{e}(\cdot)$  is an estimated propensity score. For a class of data generating processes  $\mathcal{P}_e$ , there exists a sequence  $\phi_n \rightarrow \infty$  such that

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_e} \phi_n E_{P^n} \left[ \frac{1}{n} \sum_{i=1}^n |\hat{\tau}_i^e - \tau_i| \right] < \infty. \tag{2.9}$$

In Section E of Kitagawa and Tetenov (2018b), we showed that the estimators  $\hat{\tau}^m(\cdot)$  and  $\hat{\tau}_i^e$  constructed via local polynomial regressions satisfy this condition for a certain class of data generating processes. Theorems 2.5 and 2.6 below derive the uniform convergence rate bounds of the hybrid rules in two different scenarios. In Theorem 2.5, we constrain the class of data generating processes only by Assumption 2.1 and Condition 2.1, and, importantly, we allow the class of decision rules  $\mathcal{G}$  to exclude the first-best rule  $G_{FB}^*$ . See Kitagawa and Tetenov (2018a) for proofs of these theorems.

THEOREM 2.5: *Suppose Assumption 2.1 holds.*

(m) ( $m$ -hybrid case): *Given a class of data generating processes  $\mathcal{P}_m$ , if an estimator for the conditional treatment effect  $\hat{\tau}^m(\cdot)$  satisfies Condition 2.1 (m), then*

$$\sup_{P \in \mathcal{P}_m \cap \mathcal{P}(M, \kappa)} E_{P^n} [W_{\mathcal{G}}^* - W(\hat{G}_{m\text{-hybrid}})] \leq O(\psi_n^{-1} \vee n^{-1/2}).$$

(e) ( $e$ -hybrid case): *Given a class of data generating processes  $\mathcal{P}_e$ , if an estimator for the propensity score  $\hat{e}(\cdot)$  satisfies Condition 2.1 (e), then*

$$\sup_{P \in \mathcal{P}_e \cap \mathcal{P}(M, \kappa)} E_{P^n} [W_{\mathcal{G}}^* - W(\hat{G}_{e\text{-hybrid}})] \leq O(\phi_n^{-1} \vee n^{-1/2}).$$

A comparison of Theorem 2.5 with Theorem 2.1 shows that the uniform rate upper bounds for the hybrid EWM rules are no faster than the welfare loss convergence rate

<sup>14</sup>The trimming sequence  $\varepsilon_n$  is introduced only to simplify the derivation of the rate upper bound of the welfare loss. In practical terms, if the overlap condition is well satisfied in the given data, the trimming is not necessary for computing the  $e$ -hybrid rule.



of the EWM with known propensity score. Note that if some nonparametric estimator is used to estimate  $\tau(\cdot)$  or  $e(\cdot)$ ,  $\psi_n$  or  $\phi_n$  specified in Condition 2.1 is generally slower than  $n^{1/2}$ . Hence, the welfare loss upper bounds of the hybrid rules are determined by the nonparametric rate  $\psi_n^{-1}$  or  $\phi_n^{-1}$ . A special case where the estimation of  $\tau(\cdot)$  or  $e(\cdot)$  does not affect the uniform convergence rate is when  $\tau(\cdot)$  or  $e(\cdot)$  is assumed to belong to a parametric family and it is estimated parametrically, that is,  $\psi_n$  or  $\phi_n$  is equal to  $n^{1/2}$ .

In the second scenario, we consider the case where  $\mathcal{G}$  contains the first-best decision rule  $G_{\text{FB}}^*$  and the data generating processes are constrained further by the margin assumption (Assumption 2.2) with margin coefficient  $\alpha \in (0, 1]$ .

**THEOREM 2.6:** *Suppose Assumptions 2.1 and 2.2 hold with a margin coefficient  $\alpha \in (0, 1]$ . Assume that a stronger version of Condition 2.1 holds, where (2.8) and (2.9) are replaced by*

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_m} E_{P^n} \left[ \left( \tilde{\psi}_n \max_{1 \leq i \leq n} |\hat{\tau}^m(X_i) - \tau(X_i)| \right)^2 \right] < \infty \quad \text{and} \quad (2.10)$$

$$\limsup_{n \rightarrow \infty} \sup_{P \in \mathcal{P}_e} E_{P^n} \left[ \left( \tilde{\phi}_n \max_{1 \leq i \leq n} |\hat{\tau}_i^e - \tau_i| \right)^2 \right] < \infty \quad (2.11)$$

for sequences  $\tilde{\psi}_n \rightarrow \infty$  and  $\tilde{\phi}_n \rightarrow \infty$ , respectively. Then, we have

$$\sup_{P \in \mathcal{P}_m \cap \mathcal{P}_{\text{FB}}(M, \kappa, \alpha, \eta)} E_{P^n} [W(G_{\text{FB}}^*) - W(\hat{G}_{m\text{-hybrid}})] \leq O(\tilde{\psi}_n^{-(1+\alpha)} \vee n^{-\frac{1+\alpha}{2+\alpha}} \log \tilde{\psi}_n),$$

$$\sup_{P \in \mathcal{P}_e \cap \mathcal{P}_{\text{FB}}(M, \kappa, \alpha, \eta)} E_{P^n} [W(G_{\text{FB}}^*) - W(\hat{G}_{e\text{-hybrid}})] \leq O(\tilde{\phi}_n^{-(1+\alpha)} \vee n^{-\frac{1+\alpha}{2+\alpha}} \log \tilde{\phi}_n).$$

Theorem 2.6 shows that even when  $\tau(\cdot)$  or  $e(\cdot)$  has to be estimated, the margin coefficient  $\alpha$  influences the rate upper bound of the welfare loss. A higher  $\alpha$  leads to a faster rate of the welfare loss convergence regardless of whether  $\tau(\cdot)$  and  $e(\cdot)$  are estimated parametrically or nonparametrically. In the situation where  $\tau(\cdot)$  or  $e(\cdot)$  is estimated parametrically (with a compact support of  $X$ ),  $\tilde{\psi}_n$  or  $\tilde{\phi}_n$  is equal to  $n^{1/2}$ ; thus, the uniform welfare loss convergence rate is given by the second argument in  $O(\cdot)$ ,  $n^{-\frac{1+\alpha}{2+\alpha}} \log n$ . On the other hand, when  $\tau(\cdot)$  or  $e(\cdot)$  is estimated nonparametrically, which of the two terms in  $O(\cdot)$  converges more slowly depends on the dimension of  $X$  and the degree of smoothness of the underlying nonparametric function. For instance, Section E of Kitagawa and Tetenov (2018b) shows, for suitably constructed  $\mathcal{P}_m$ , local polynomial estimation for  $\tau(x)$  can attain  $\tilde{\psi}_n = n^{\frac{1}{2+d_x/\beta_m}} (\log n)^{-\frac{1}{2+d_x/\beta_m}-2}$ , where  $d_x$  is the dimension of  $X$  and  $\beta_m$  is the degree of Hölder smoothness of  $m_1(x)$  and  $m_0(x)$ . Hence, if  $d_x/\beta_m > \alpha$ , the rate upper bound of Theorem 2.6 implies that the welfare loss convergence rate of the  $m$ -hybrid rule is guaranteed to be  $n^{-\frac{1+\alpha}{2+d_x/\beta_m}} (\log n)^{\frac{1}{(2+d_x/\beta_m)+2}(1+\alpha)}$ .

Note that Theorems 2.5 and 2.6 concern the upper bound of the convergence rate. We do not have the universal rate lower bound results for these constrained classes of data generating processes. We leave the investigation of the sharp rate bound of the hybrid-EWM welfare loss for future research.

### 3. EMPIRICAL APPLICATION

We illustrate the Empirical Welfare Maximization method by applying it to experimental data from the National Job Training Partnership Act (JTPA) Study. A detailed descrip-

TABLE I  
ESTIMATED WELFARE GAIN OF ALTERNATIVE TREATMENT ASSIGNMENT RULES THAT CONDITION ON  
EDUCATION AND PRE-PROGRAM EARNINGS<sup>a</sup>

Outcome Variable: Treatment Rule:	30-Month Post-Program Earnings, No Treatment Cost		30-Month Post-Program Earnings, \$774 Cost for Each Assigned Treatment	
	Share of Population to Be Treated	Est. Welfare Gain per Population Member	Share of Population to Be Treated	Est. Welfare Gain per Population Member
Treat everyone	1	\$1,180 (\$464, \$1,896)	1	\$404 (-\$313, \$1,121)
EWM quadrant rule	0.95	\$1,340 (\$441, \$2,239)	0.8	\$643 (-\$258, \$1,544)
EWM linear rule	0.96	\$1,364 (\$398, \$2,330)	0.69	\$792 (-\$177, \$1,761)
EWM linear rule (with (education) <sup>2</sup> and (education) <sup>3</sup> )	0.88	\$1,489 (\$374, \$2,603)	0.75	\$897 (-\$214, \$2,008)
Linear regression plug-in rule	0.98	\$1,152	0.86	\$527
Linear regression plug-in rule (with (education) <sup>2</sup> and (education) <sup>3</sup> )	0.95	\$1,263	0.91	\$547
Nonparametric plug-in rule	0.91	\$1,693	0.78	\$996

<sup>a</sup>Two-sided 95% confidence intervals in parentheses. See Appendix B in Kitagawa and Tetenov (2018a) for their construction and asymptotic validity.

tion of the study and an assessment of average program effects for five large subgroups of the target population are found in Bloom et al. (1997). The study randomized whether applicants would be 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 on the applicants prior to random assignment, as well as administrative and survey data on applicants' earnings in the 30-month period following the assignment. Our sample consists of 9,223 observations with available data on years of education and pre-program earnings from the sample of adults (22 years and older) used in the original evaluation of the program and in subsequent studies (Bloom et al. (1997), Heckman, Ichimura, and Todd (1997), Abadie, Angrist, and Imbens (2002)). The probability of being assigned to the treatment was two thirds in this sample.

We use two welfare outcome measures for our illustration. The first is the total individual earnings in the 30 months after program assignment. The second is the 30-month earnings minus \$774 for individuals who were assigned to treatment. This is the average cost of services per treatment *assignment* (estimated from Table 5 in Bloom et al. (1997)), which takes into account varying take-up of program services. The first outcome measure reflects social preferences that put no weight on the costs of the program incurred by the government. The second outcome measure weighs participants' gains and the government's losses equally.

We consider these outcomes (and costs) from an *intention-to-treat* perspective. We view the policy maker's problem as a choice of the eligibility criteria and not as a choice of the take-up rate (decided by individuals); hence, we are not interested in the treatment effect on compliers. Since we have to compare welfare effects of policies that assign different




FIGURE 1.—Empirical welfare-maximizing treatment rules from the quadrant class conditioning on years of education and pre-program earnings.

proportions of the population to the treatment, we report estimates of the *average effect per population member*  $E[(Y_1 - Y_0) \cdot 1\{X \in G\}]$ , which is proportional to the total welfare effect of the treatment rule  $G$ .

We consider conditioning treatment assignment on two pre-treatment variables: the individual's years of education and earnings in the year prior to the assignment. Both variables may plausibly affect how much effect the individual gets from the program services. We do not use race, gender, or age. Though treatment effects may vary with these characteristics, policy makers usually cannot use them to determine treatment assignment. Education and earnings are generally verifiable characteristics. This is an important feature for implementing the proposed treatment assignment because the empirical welfare estimates are not valid for the target population if the individuals could manipulate their characteristics to obtain the desired treatment.

Table I reports the estimated welfare gains  $W_n^{\text{dm}}(G) - W_n^{\text{dm}}(\emptyset)$  of alternative treatment rules relative to the benchmark of assigning no-one to treatment, as well as the estimated proportion of individuals assigned to treatment 1 by each rule.<sup>15</sup>

<sup>15</sup>We report welfare gain estimates using equation (2.7) with demeaned outcome variable (see Remark 2.3). These estimates are invariant to translation of the outcome variable by a constant.




FIGURE 2.—Empirical welfare-maximizing treatment rules from the linear class conditioning on years of education and pre-program earnings.

We consider three candidate classes of treatment rules for EWM.<sup>16</sup> The first is the class of *quadrant treatment rules*:

$$\mathcal{G}_Q \equiv \left\{ \left\{ x : s_1(\text{education} - t_1) > 0 \ \& \ s_2(\text{prior earnings} - t_2) > 0 \right\}, \right. \\ \left. s_1, s_2 \in \{-1, 1\}, t_1, t_2 \in \mathbb{R} \right\}. \tag{3.1}$$

This class of treatment rules is easily implementable and is often used in practice. To be assigned to treatment according to such a rule, an individual’s education and pre-program earnings have to be above (or below) some specific thresholds. The EWM method searches over all possible thresholds and directions. Figure 1 demonstrates the quadrant treatment rules selected by the EWM criterion. The full shaded area indicates individuals who would be assigned to treatment if it were costless (education  $\leq 15$ , prior earnings  $\leq \$19,670$ ). The EWM treatment rule that takes into account \$774 treatment assignment cost has the same earnings threshold, but lowers the education threshold to 12. The size of black dots indicates the number of individuals with different covariate values.

<sup>16</sup>Specifically, we implement demeaned EWM described in Remark 2.3 with known constant propensity score  $e(X_i) = 2/3$ . Further details on computing EWM rules are found in Appendix C of Kitagawa and Tetenov (2018a).




FIGURE 3.—Parametric plug-in treatment rules based on the linear regressions of treatment outcomes on years of education and pre-program earnings.

Second, we consider two classes of *linear treatment rules*:

$$\begin{aligned}
 \mathcal{G}_{LES} &\equiv \{ \{x : \beta_0 + \beta_1 \cdot \text{prior earnings} + \beta_2 \cdot \text{education} > 0\}, \beta_0, \beta_1, \beta_2 \in \mathbb{R} \}, \\
 \mathcal{G}_{LES}^3 &\equiv \left\{ \left\{ x : \begin{pmatrix} \beta_0 + \beta_1 \cdot \text{prior earnings} + \beta_2 \cdot \text{education} \\ + \beta_3 \cdot (\text{education})^2 + \beta_4 \cdot (\text{education})^3 \end{pmatrix} > 0 \right\}, \right. \\
 &\quad \left. \beta_0, \beta_1, \beta_2, \beta_3, \beta_4 \in \mathbb{R} \right\}. \tag{3.2}
 \end{aligned}$$

Linear treatment rules that maximize empirical welfare could markedly differ from plug-in rules derived from linear regressions. When treatment costs are not considered, the direction of treatment assignment as a function of prior earnings differs between the EWM rule (Figure 2) and the linear regression plug-in rule (Figure 3). When treatment costs are considered, the direction of treatment assignment is similar, but the two treatment rules substantially differ in the proportion of population assigned (69% by the EWM rule vs. 86% by the plug-in rule, Table I). If a regression equation is correctly specified, the regression plug-in and EWM rules have identical large sample limits. If the regression equation is misspecified, however, only the linear EWM treatment rule converges with sample size to the welfare-maximizing limit, and the welfare level attained by the regression plug-in rule can be suboptimal even in large samples.

When quadratic and cubic terms for education are included, the EWM rule (Figure 4) uses the additional flexibility of this functional form to change the shape of the treat-

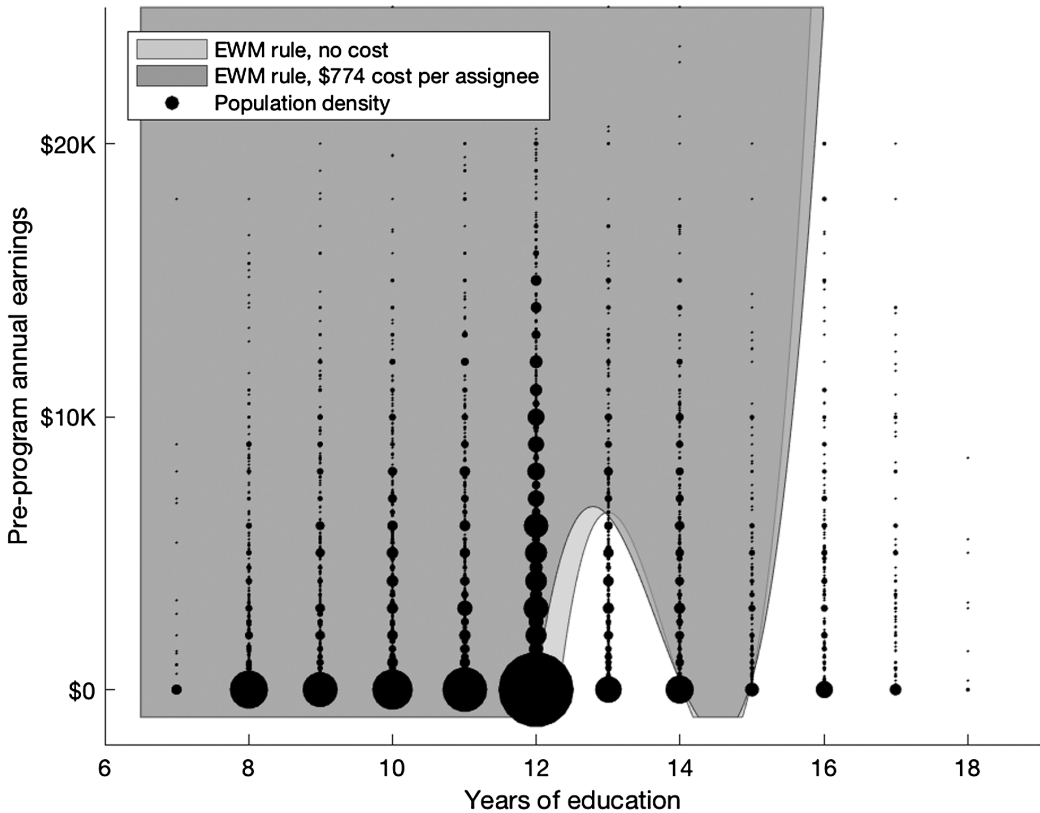


FIGURE 4.—Empirical welfare-maximizing treatment rules from the linear class conditioning on years of education,  $(\text{education})^2$ ,  $(\text{education})^3$ , and pre-program earnings.

ment assignment boundary, setting a higher prior earnings threshold for individuals with 13 years of education. It can be seen from Figure 6 that treatment effects estimated by nonparametric regression for these excluded individuals are also low. In comparison, the linear regression (Figure 5) uses the additional terms to improve the global fit of the earnings equation and does not exclude this subset of the population for whom the treatment effect seems to be negative.

Figure 6 shows plug-in treatment rules based on kernel regressions of treatment and control outcomes on the covariates.<sup>17</sup> The class of nonparametric plug-in rules is much richer than the quadrant or the linear class of treatment rules, and it may obtain higher welfare in large samples. It is clear from the figure, however, that this class of patchy decision rules may be difficult to implement in public policy, where clear and transparent treatment rules are required.

#### 4. CONCLUSION

The EWM approach considered in this paper directly maximizes a sample analog of the welfare criterion of a policy maker. This welfare-function-based statistical procedure for

<sup>17</sup>The bandwidths were chosen by Silverman's rule of thumb.




FIGURE 5.—Parametric plug-in treatment rules based on the regressions of treatment outcomes on years of education,  $(\text{education})^2$ ,  $(\text{education})^3$ , and pre-program earnings.

treatment choice differs fundamentally from parametric and nonparametric plug-in approaches, which do not integrate statistical inference and the decision problem at hand. We investigated the statistical performances of the EWM rule in terms of the uniform convergence rate of the welfare loss and demonstrated that with known propensity scores, the EWM rule attains minimax optimal rates over various classes of feasible data distributions. The EWM approach offers a useful framework for the individualized policy assignment problems, as the EWM approach can easily accommodate the constraints that policy makers commonly face in reality. We also presented methods to compute the EWM rule for many practically important classes of treatment assignment rules and demonstrated them using experimental data from the JTPA program.

Several extensions and open questions remain to be answered. First, this paper assumed that the class of candidate policies  $\mathcal{G}$  is given exogenously to the policy maker. We did not consider how to select the class  $\mathcal{G}$  when the policy maker is free to do so. See Swaminathan and Joachims (2015) and Mbakop and Tabord-Meehan (2017) for recent developments. Second, while EWM attains minimax rate-optimality, it is unclear whether the EWM rule has stronger decision-theoretic optimality properties for the nonparametric class of data generating processes we considered. It remains to be seen whether EWM obtains the lowest asymptotic maximum regret within the class of minimax rate-optimal rules, whether it is admissible, and whether it is Bayes-optimal for some prior over  $P$ . It is an open question whether modifications of EWM or other rate-optimal rules could perform better.




FIGURE 6.—Nonparametric plug-in treatment rules based on the kernel regressions of treatment outcomes on years of education and pre-program earnings.

Third, we ruled out the case in which the data are subject to selection on unobservables or the overlap conditions fails. In these situations, the welfare criterion could be only set-identified, and it is not clear how to extend the EWM idea to this case. Fourth, we restricted our analysis to the additive social welfare criterion, but in some contexts, policy makers have a non-additive social welfare criterion. See Kitagawa and Tetenov (2017) for an extension of the EWM approach to a class of generalized Gini social welfare functions.

#### REFERENCES

- 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 (1), 91–117. [608]
- ABADIE, A., M. M. CHINGOS, AND M. R. WEST (2017): “Endogenous Stratification in Randomized Experiments,” Unpublished Manuscript. [596]
- ARMSTRONG, T. B., AND S. SHEN (2015): “Inference on Optimal Treatment Assignments,” Unpublished Manuscript. [596]
- ATHEY, S., AND S. WAGER (2017): “Efficient Policy Learning,” Preprint. Available at arXiv:1702.0289. [596]
- ATKINSON, A. B. (1970): “On the Measurement of Inequality,” *Journal of Economic Theory*, 2 (3), 244–263. [592]
- AUDIBERT, J.-Y., AND A. B. TSYBAKOV (2007): “Fast Learning Rates for Plug-in Classifiers,” *The Annals of Statistics*, 35 (2), 608–633. [596]



- BEYGEZLIMER, A., AND J. LANGFORD (2009): “The Offset Tree for Learning With Partial Labels,” in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York: ACM, 129–137. [595]
- BHATTACHARYA, D., AND P. DUPAS (2012): “Inferring Welfare Maximizing Treatment Assignment Under Budget Constraints,” *Journal of Econometrics*, 167 (1), 168–196. [595,596,601]
- 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 (3), 549–576. [595,608]
- 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: Oxford University Press, 11–39. [595,597]
- DEHEJIA (2005): “Program Evaluation as a Decision Problem,” *Journal of Econometrics*, 125, 141–173. [595, 597]
- DEVROYE, L., AND G. LUGOSI (1995): “Lower Bounds in Pattern Recognition and Learning,” *Pattern Recognition*, 28, 1011–1018. [600]
- DEVROYE, L., L. GYÖRFI, AND G. LUGOSI (1996): *A Probabilistic Theory of Pattern Recognition*. New York: Springer. [596]
- DUDÍK, M., J. LANGFORD, AND L. LI (2011): “Doubly Robust Policy Evaluation and Learning,” in *Proceedings of the 28th International Conference on Machine Learning*. Madison, WI: Omnipress, 1097–1104. [596]
- DUDLEY, R. M. (1999): *Uniform Central Limit Theorems*. Cambridge: Cambridge University Press. [597,599]
- ELLIOTT, G., AND R. P. LIELI (2013): “Predicting Binary Outcomes,” *Journal of Econometrics*, 174 (1), 15–26. [597]
- FLORIOS, K., AND S. SKOURAS (2008): “Exact Computation of Max Weighted Score Estimators,” *Journal of Econometrics*, 146 (1), 86–91. [595]
- HECKMAN, J. J., H. ICHIMURA, AND P. E. TODD (1997): “Matching as an Econometric Evaluation Estimator: Evidence From Evaluating a Job Training Programme,” *Review of Economic Studies*, 64 (4), 605–654. [608]
- HIRANO, K., AND J. R. PORTER (2009): “Asymptotics for Statistical Treatment Rules,” *Econometrica*, 77 (5), 1683–1701. [595,596]
- KASY, M. (2016): “Partial Identification, Distributional Preferences, and the Welfare Ranking of Policies,” *Review of Economics and Statistics*, 98 (1), 111–131. [596]
- (2017): “Optimal Taxation and Insurance Using Machine Learning,” Unpublished Manuscript. [595, 597]
- KERKYACHARIAN, G., A. B. TSYBAKOV, V. TEMLYAKOV, D. PICARD, AND V. KOLTCHINSKII (2014): “Optimal Exponential Bounds on the Accuracy of Classification,” *Constructive Approximation*, 39, 421–444. [596]
- KITAGAWA, T., AND A. TETENOV (2017): “Equality-Minded Treatment Choice,” Cemmap Working Paper 10/17. [614]
- (2018a): “Supplement to ‘Who Should Be Treated? Empirical Welfare Maximization Method for Treatment Choice,’” *Econometrica Supplemental Material*, 86, <http://dx.doi.org/10.3982/ECTA13288>. [598, 600,605,606,608,610]
- (2018b): “Online Appendix to ‘Who Should Be Treated? Empirical Welfare Maximization Method for Treatment Choice,’” *Econometrica Supplemental Material*, 86, <http://dx.doi.org/10.3982/ECTA13288>. [601, 602,605-607]
- LIELI, R. P., AND H. WHITE (2010): “The Construction of Empirical Credit Scoring Rules Based on Maximization Principles,” *Journal of Econometrics*, 157 (1), 110–119. [597]
- LUGOSI, G. (2002): “Pattern Classification and Learning Theory,” in *Principles of Nonparametric Learning*, ed. by L. Györfi. Vienna: Springer, 1–56. [600]
- MAMMEN, E., AND A. B. TSYBAKOV (1999): “Smooth Discrimination Analysis,” *The Annals of Statistics*, 27, 1808–1829. [603]
- MANSKI, C. F. (1975): “Maximum Score Estimation of the Stochastic Utility Model of Choice,” *Journal of Econometrics*, 3 (3), 205–228. [594]
- (2004): “Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica*, 72 (4), 1221–1246. [594,595]
- MANSKI, C. F., AND T. S. THOMPSON (1989): “Estimation of Best Predictors of Binary Response,” *Journal of Econometrics*, 40 (1), 97–123. [594,595,597]
- MASSART, P., AND É. NÉDÉLEC (2006): “Risk Bounds for Statistical Learning,” *The Annals of Statistics*, 34 (5), 2326–2366. [596,604]
- MBAKOP, E., AND M. TABORD-MEEHAN (2017): “Model Selection for Treatment Choice: Penalized Welfare Maximization,” Preprint. Available at arXiv:1609.03167. [613]
- QIAN, M., AND S. A. MURPHY (2011): “Performance Guarantees for Individualized Treatment Rules,” *The Annals of Statistics*, 39 (2), 1180–1210. [596]

- STOYE, J. (2009): “Minimax Regret Treatment Choice With Finite Samples,” *Journal of Econometrics*, 151 (1), 70–81. [595]
- (2012): “Minimax Regret Treatment Choice With Covariates or With Limited Validity of Experiments,” *Journal of Econometrics*, 166 (1), 138–156. [595]
- SWAMINATHAN, A., AND T. JOACHIMS (2015): “Batch Learning From Logged Bandit Feedback Through Counterfactual Risk Minimization,” *Journal of Machine Learning Research*, 16, 1731–1755. [613]
- TETENOV, A. (2012): “Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria,” *Journal of Econometrics*, 166 (1), 157–165. [595]
- TSYBAKOV, A. B. (2004): “Optimal Aggregation of Classifiers in Statistical Learning,” *The Annals of Statistics*, 32 (1), 135–166. [596,603,604]
- VAN DER VAART, A. W., AND J. A. WELLNER (1996): *Weak Convergence and Empirical Processes*. New York: Springer. [597]
- VAPNIK, V. N. (1998): *Statistical Learning Theory*. New York: Wiley. [596,597]
- ZADROZNY, B. (2003): “Policy Mining: Learning Decision Policies From Fixed Sets of Data,” Ph.D. Thesis, University of California, San Diego. [595]
- ZHANG, B., A. A. TSIATIS, E. B. LABER, AND M. DAVIDIAN (2012): “A Robust Method for Estimating Optimal Treatment Regimes,” *Biometrics*, 68 (4), 1010–1018. [596]
- ZHAO, Y., D. ZENG, A. J. RUSH, AND M. R. KOSOROK (2012): “Estimating Individualized Treatment Rules Using Outcome Weighted Learning,” *Journal of the American Statistical Association*, 107 (499), 1106–1118. [596]

---

*Co-editor Liran Einav handled this manuscript.*

*Manuscript received 9 March, 2015; final version accepted 29 November, 2017; available online 29 November, 2017.*