cemmap
centre for microdata methods and practice

Who should be treated? Empirical welfare maximization methods for treatment choice

Toru Kitagawa
Aleksey Tetenov

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP10/15

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL
An ESRC Research Centre

Who should be Treated? Empirical Welfare Maximization Methods for Treatment Choice*

Toru Kitagawa† and Aleksey Tetenov‡

9 March 2015

Abstract

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 propose 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. In comparison with this benchmark rate, we examine how the uniform convergence rate of the average welfare improves or deteriorates depending 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 provide an asymptotically valid inference procedure for the population welfare gain obtained by exercising the EWM rule. We offer easily implementable algorithms for computing the EWM rule and an application using experimental data from the National JTPA Study.

*We would like to thank Shin Kanaya, Charles Manski, Joerg Stoye, and seminar participants at Academia Sinica, NYU, and University of Virginia for beneficial comments. We would also like to thank the participants at 2014 Bristol Econometrics 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 number RES-589-28-0001) is gratefully acknowledged.

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

‡Collegio Carlo Alberto, Email: aleksey.tetenov@carloalberto.org

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 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.[1] We consider the *utilitarian* welfare criterion defined by the average of the individual outcomes in the target population. When treatment rule G is applied to the target population, the utilitarian welfare equals

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

where $E_P(\cdot)$ is the expectation with respect to P. 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 welfare criterion as

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

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

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

Hence, if the probability distribution of observables (Y, D, X) was 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 $E_P(Y_0)$:

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

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

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

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

[1] In Section 4.2, we consider a setting where the target and the sampled populations have identical conditional treatment effects, but different marginal distributions of X.

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,*

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

which maximizes the population welfare (1.1) if any assignment rule were feasible to implement. Our framework with a constrained class of feasible assignment rules allows us to incorporate several 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 legal or ethical restrictions or due to public accountability for the treatment eligibility criterion.

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

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

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 $W^*_{\mathcal{G}}$

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

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 one, as is the case in randomized experiments, we derive a non-asymptotic distribution-free upper bound of $E_{P^n}\left[W^*_{\mathcal{G}} - W(\hat{G}_{EWM})\right]$ 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 $W^*_{\mathcal{G}}$ 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, i.e., 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, i.e., $G^*_{FB} \in \mathcal{G}$, 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

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

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

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

where $\hat{\tau}^e_i \equiv \left[\frac{Y_i D_i}{\hat{e}(X_i)} - \frac{Y_i(1-D_i)}{1-\hat{e}(X_i)}\right] \cdot 1\{\varepsilon_n \leq \hat{e}(X_i) \leq 1 - \varepsilon_n\}$ with a converging positive sequence $\varepsilon_n \to 0$ as $n \to \infty$. 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 $\hat{e}(\cdot)$.

When performing the treatment choice analysis, it could also be of interest to assess the sampling uncertainty of the estimated welfare gain from implementing the treatment rule $\hat{G}$. For this purpose, this paper proposes an inference procedure for $W(\hat{G}) - W(G_0)$, where G_0 is a benchmark treatment assignment rule, such as no treatment ($G_0 = \emptyset$) or the non-individualized implementation of the treatment ($G_0 = \mathcal{X}$).

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 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 et al. (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 (2014). Manski (2004) proposes to assess the welfare properties of statistical treatment rules by their maximum regret and derives finite-sample bounds on the maximum regret of *Conditional Empirical Success* 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 shows 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) shows 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.

The problem of conditional treatment choice has some similarities to the *classification problem* in machine learning and statistics, since it seeks a way to optimally "classify" individuals into those who benefit from the treatment and those who do not. This similarity allows us to draw on recent theoretical results for classification by Devroye et al. (1996), Tsybakov (2004), Massart and Nédélec (2006), Audibert and Tsybakov (2007), and Kerkyacharian et al. (2014), among others, and to adapt them to the treatment choice problem. The minimax rate optimality of the EWM treatment choice rule (proved in Theorems 2.1 and 2.2 below) is analogous to the minimax rate optimality of the Empirical Risk Minimization classifier in the classification problem shown by Devroye and Lugosi (1995). There are, however, substantive differences between treatment choice and classification problems: (1) the observed outcomes are real-valued rather than binary, (2) in treatment choice only one of the potential outcomes is observed for each individual, whereas in classification the correct choice is known for each training sample observation, (3) the EWM criterion depends on the propensity score, which may be unknown (as in observational studies), (4) policy settings often

impose constraints on practicable treatment rules or on the proportion of the population that could be treated. Accommodating these fundamental differences and establishing the convergence rate results for the welfare loss criterion constitute the main theoretical contributions of this paper.

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

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

where $\hat{m}_d(x)$ is a parametric or a nonparametric estimator of $E(Y_d|X = x)$. Hirano and Porter (2009) establish local asymptotic minimax optimality of plug-in rules for parametric and semiparametric models of treatment response. Bhattacharya and Dupas (2012) apply nonparametric plug-in rules with an aggregate budget constraint and derive some of their properties. Armstrong and Shen (2014) consider statistical inference for the first-best decision rule G^*_{FB} from the perspective of inference for conditional moment inequalities. In empirical practice of program evaluation, researchers assess who should be treated by stratifying the population based on the predicted value of the individual outcome in the absence of treatment and estimating the average causal effects for each strata using the experimental data. Abadie et al. (2014) point out that the naive implementation of this idea is subject to a bias due to the endogenous stratification and provide a method to correct the bias.

To assess the 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 the social welfare. By formulating the treatment choice problem as a formal statistical decision problem, a certain treatment assignment rule could be obtained by specifying a prior distribution for P and solving for a Bayes decision rule (see Dehejia (2005), Chamberlain (2011), and Kasy (2014) for Bayesian approaches to the treatment choice problem). In contrast to the Bayesian approach, the EWM approach is purely data-driven and does not require a prior distribution over the data generating processes.

The analysis of optimal *individualized treatment rules* has also received considerable attention in biostatistics. Qian and Murphy (2011) propose a plug-in approach using the treatment response function $\hat{m}_D(X) \equiv D\hat{m}_1(X) + (1 - D)\hat{m}_0(X)$ estimated by penalized least squares. If the square loss $E(Y - \hat{m}_D(X))^2$ converges to the minimum $E(Y - DE(Y_1|X) + (1 - D)E(Y_0|X))^2$, then the welfare of the plug-in treatment rule converges to the maximum $W(G^*_{FB})$. Assuming that

treatment response functions are linear in the chosen basis, they show welfare convergence rate of $n^{-1/2}$ or better (with a margin condition). Zhao et al. (2012) propose another *surrogate loss function* approach that estimates the treatment rule using a Support Vector Machine with a growing basis. This yields 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.

Elliott and Lieli (2013) and Lieli and White (2010) also proposed maximizing the sample analog of a utilitarian decision criterion similar to EWM. They consider the problem of forecasting binary outcomes based on observations of (Y_i, X_i), 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*[2] $v < \infty$ and is countable.[3]

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

[2]The VC-diemension of $\mathcal{G}$ is defined by the maximal number of points in $\mathcal{X}$ that can be shattered by $\mathcal{G}$. 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, e.g., van der Vaart and Wellner (1996).

[3]Coutability of $\mathcal{G}$ is imposed in order to avoid measurability complications in proving our theoretical results.

to the observational studies in which unconfoundedness is credible. The second assumption (BO) implies boundedness of the treatment effects, i.e.,

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

where P_X is the marginal distribution of X and $\tau(\cdot)$ is the conditional treatment effect $\tau(X) = E(Y_1 - Y_0|X)$. Since the implementation of EWM does not require knowledge of M and unbounded Y is rare in social science, this assumption is innocuous and imposed only 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 σ-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 $x \in \mathbb{R}^{d_x}$, $x_{sub} \in \mathbb{R}^{d_{sub}}$, $d_{sub} \leq d_x$. The class of decision sets generated by Linear Eligibility Scores (LES) is defined as*

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

We accordingly obtain an EWM assignment rule by maximizing

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

in $\beta = \left(\beta_0, \beta_{sub}^T\right) \in \mathbb{R}^{d_{sub}+1}$. It is well known that the class of half-spaces spanned by $\left(\beta_0, \beta_{sub}^T\right) \in \mathbb{R}^{d_{sub}+1}$ has the VC-dimension $v = d_{sub}+1$, so Assumption 2.1 (VC) holds. In Section 5, we discuss

how to compute $\hat{G}_{EWM}$ when the class of decision sets is given by $\mathcal{G}_{LES}$. A plug-in rule based on a parametric linear regression also selects a treatment rule from $\mathcal{G}_{LES}$, but their welfare does not converge to the maximum welfare $W^*_{\mathcal{G}_{LES}}$ if the regression equations are misspecified, whereas the welfare of $\hat{G}_{EWM}$ always does (as shown in Theorem 2.1 below).

Example 2.2. *(Generalized Eligibility Score) Let $f_j(\cdot)$, $j = 1, \ldots, m$, 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),*

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

The class of decision sets $\mathcal{G}_{GES}$ generalizes the linear eligibility score rules (2.1), as it allows for eligibility scores that are nonlinear in x, i.e., $\mathcal{G}_{GES}$ can accommodate decision sets that partition the space of covariates by nonlinear boundaries. It can be shown that $\mathcal{G}_{GES}$ 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 $\mathcal{G}_{GES,l}$, $l = 1, \ldots, L$, 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, $\mathcal{G} \equiv \bigcap_{l=1}^{L} \mathcal{G}_{GES,l} = \left\{ \bigcap_{l=1}^{L} G_l : G_l \in \mathcal{G}_{GES,l}, l = 1, \ldots, L \right\}$. 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)); thus, Assumption 2.1 (VC) holds for this $\mathcal{G}$. 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 $\{\mathcal{G}_{GES,l}, l = 1, \ldots, L\}$, 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 $\sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n} \left[W^*_{\mathcal{G}} - W(\hat{G}_{EWM}) \right]$ 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 $\tilde{G} \in \mathcal{G}$:

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

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

$$W^*_{\mathcal{G}} - W(\hat{G}_{EWM}) \leq 2 \sup_{G \in \mathcal{G}} |W_n(G) - W(G)| . \tag{2.2}$$

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 Appendix A.2.

Theorem 2.1. *Under Assumption 2.1, we have*

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

where C_1 is a universal constant defined in Lemma A.4 in Appendix A.1.

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, $\hat{G}_{EWM}$ tends to overfit the data in the sense that the distribution of regret $W^*_{\mathcal{G}} - W(\hat{G}_{EWM})$ is more and more dispersed, and, with n fixed, this overfitting results in inflating the average welfare regret.[4]

[4]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 $\mathcal{G}$.

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 and the VC-dimension of $\mathcal{G}$ is $v \geq 2$. Then, for any treatment choice rule $\hat{G}$, as a function of $(Z_1, \ldots, Z_n)$, it holds*

$$\sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n}\left[W_{\mathcal{G}}^* - W(\hat{G})\right] \geq 4^{-1} \exp\left\{-2\sqrt{2}\right\} M\sqrt{\frac{v-1}{n}} \quad \textit{for all } n \geq 16\,(v-1)\,.$$

This theorem, combined with Theorem 2.1, implies that $\hat{G}_{EWM}$ is *minimax rate optimal* over the class of data generating process $\mathcal{P}(M,\kappa)$, since the rate of the convergence of the upper bound of $\sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n}\left[W_{\mathcal{G}}^* - W(\hat{G}_{EWM})\right]$ 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 $\hat{G}_{EWM}$ 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 $E_{P^n}\left[W(\hat{G})\right] > W_{\mathcal{G}}^*$ for some pairs of $\hat{G}$ and P, but it can never hold for all distributions in $\mathcal{P}(M,\kappa)$.[5]

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

[5] For example, if $\hat{G}$ is a nonparametric plug-in rule and the first-best decision rule G_{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 $E_{P^n} W(\hat{G}) \leq W_{\mathcal{G}}^* - (n^{-1/2}$ bound) for the same sample size.

of candidate treatment rules $\mathcal{G}$.

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

$$P_X(|\tau(X)| \leq t) \leq \left(\frac{t}{\eta}\right)^{\alpha}, \quad \forall 0 \leq t \leq \eta,$$

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

The assumption of correct specification means that the class of the 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 η and α characterize the size of population with the conditional treatment effect close to the margin $\tau(X) = 0$. *Smaller* η and α imply that *more* individuals can concentrate in a neighborhood of $\tau(X) = 0$. The next examples illustrate this interpretation of η and α.

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 effects to be $\tau(X) = (-X)^3$. In this specification, $\tau(X)$ "flats out" at $\tau(X) = 0$, 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

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

By setting $\eta = h$, the margin condition holds for arbitrarily large α. 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 α.

From now on, we denote the class of P satisfying Assumptions 2.1 and 2.2 by $\mathcal{P}_{FB}(M, \kappa, \eta, \alpha)$.[6] The next theorem provides the upper bound of the welfare loss of the EWM rule when a class of data distributions is constrained to $\mathcal{P}_{FB}(M, \kappa, \eta, \alpha)$.

Theorem 2.3. *Under Assumptions 2.1 and 2.2,*

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

holds for all n, where c is a positive constant that depends only on M, κ, η, and α.

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).[7] 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 $\hat{G}_{EWM}$) 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.

[6] Note that $P_{FB}(M, \kappa, \eta, \alpha)$ depends on the set of feasible treatment rules $\mathcal{G}$ via Assumption 2.2 (FB).

[7] Tsybakov (2004) defines the complexity of the decision sets $\mathcal{G}$ in terms of the growth coefficient ρ 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 ρ being arbitrarily close to zero.

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 $(Z_1, \ldots, Z_n)$, it holds*

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

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

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

Remark 2.1. *The convergence rates of the worst-case EWM welfare loss obtained by Theorems 2.1 and 2.3 highlight how margin coefficient α influences the uniform performance of the EWM rule. Higher α 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 α 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 2.2. *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 $\{\mathcal{G}_n : n = 1, 2, \ldots\}$, for which the VC-dimension grows with the sample size at rate n^λ, $0 < \lambda < 1$, 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. 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 2.3. *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 $\hat{G}_{plug-in}$ defined in (1.13) is subject to the same lower bound. In Section 4.3, we further discuss the welfare loss uniform convergence rate of the nonparametric plug-in rule.*

Remark 2.4. *Let $\mathcal{P}_{FB}(M, \kappa)$ 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*

Appendix A.2 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 $\mathcal{P}_{FB}(M,\kappa)$. 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 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, $\epsilon_n \leq O(n^{-a})$, $a > 0$.[8]

Let $W_n^\tau(G)$ be the sample analogue of the welfare criterion (1.2) that one would construct if the true regression equations were known, $W_n^\tau(G) \equiv E_n(m_0(X_i)) + E_n(\tau(X_i)\cdot 1\{X_i \in G\})$, and $\hat{W}_n^\tau(G)$ be the empirical welfare with the conditional treatment effect estimators $\hat{\tau}^m(\cdot)$ plugged in,

$$\hat{W}_n^\tau(G) \equiv E_n\left[m_0(X_i) + \hat{\tau}^m(X_i)\, 1\{X_i \in G\}\right]. \tag{2.3}$$

Since the m-hybrid rule maximizes $\hat{W}_n^\tau(\cdot)$, it holds $\hat{W}_n^\tau(\hat{G}_{m-hybrid}) - \hat{W}_n^\tau(\tilde{G}) \geq 0$ for any $\tilde{G} \in \mathcal{G}$. The following inequalities therefore follow:

$$\begin{aligned} W(\tilde{G}) - W(\hat{G}_{m-hybrid}) &\leq W_n^\tau(\tilde{G}) - \hat{W}_n^\tau\left(\tilde{G}\right) - W_n^\tau(\hat{G}_{m-hybrid}) + \hat{W}_n^\tau\left(\hat{G}_{m-hybrid}\right) \\ &\quad + W(\tilde{G}) - W(\hat{G}_{m-hybrid}) - W_n^\tau(\tilde{G}) + W_n^\tau(\hat{G}_{m-hybrid}) \\ &= \frac{1}{n}\sum_{i=1}^n \left[\tau(X_i) - \hat{\tau}^m(X_i)\right]\left[1\left\{X_i \in \tilde{G}\right\} - 1\left\{X_i \in \hat{G}_{m-hybrid}\right\}\right] \\ &\quad + W(\tilde{G}) - W_n^\tau(\tilde{G}) + W_n^\tau(\hat{G}_{m-hybrid}) - W(\hat{G}_{m-hybrid}) \\ &\leq \frac{1}{n}\sum_{i=1}^n |\hat{\tau}^m(X_i) - \tau(X_i)| + 2\sup_{G\in\mathcal{G}} |W_n^\tau(G) - W(G)|. \end{aligned} \tag{2.4}$$

This implies that the average welfare loss of the m-hybrid rule can be bounded by

$$E_{P^n}\left[W_{\mathcal{G}}^* - W(\hat{G}_{m-hybrid})\right] \leq E_{P^n}\left[\frac{1}{n}\sum_{i=1}^n |\hat{\tau}^m(X_i) - \tau(X_i)|\right] + 2E_{P^n}\left[\sup_{G\in\mathcal{G}} |W_n^\tau(G) - W(G)|\right].$$

[8] The trimming sequence ε_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.

(2.5)

For the e-hybrid rule, replacing $W_n^\tau(\cdot)$ and $\hat{W}_n^\tau(\cdot)$ in (2.4) with the empirical welfare $W_n(\cdot)$ defined in (1.7) and $\hat{W}_n(G) \equiv E_n\left[\frac{Y_i(1-D_i)}{1-e(X_i)} + \hat{\tau}_i^e \cdot 1\{X_i \in G\}\right]$, respectively, yields a similar upper bound

$$E_{P^n}\left[W_{\mathcal{G}}^* - W(\hat{G}_{e-hybrid})\right] \leq E_{P^n}\left[\frac{1}{n}\sum_{i=1}^n |\hat{\tau}_i^e - \tau_i|\right] + 2E_{P^n}\left[\sup_{G\in\mathcal{G}} |W_n(G) - W(G)|\right], \tag{2.6}$$

where $\tau_i = \frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)}$. Since the uniform convergence rate of $E_{P^n}\left[\sup_{G\in\mathcal{G}} |W_n^\tau(G) - W(G)|\right]$ is the same as that of $E_{P^n}\left[\sup_{G\in\mathcal{G}} |W_n(G) - W(G)|\right]$,[9] these upper bounds imply that the lack of knowledge of the propensity score may harm the welfare loss convergence rate if the average estimation error of the conditional treatment effect converges slower than does $E_{P^n}\left[\sup_{G\in\mathcal{G}} |W_n(G) - W(G)|\right]$. It is therefore convenient to first state the condition regarding 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 $\psi_n \to \infty$ such that

$$\limsup_{n\to\infty} \sup_{P\in\mathcal{P}_m} \psi_n E_{P^n}\left[\frac{1}{n}\sum_{i=1}^n |\hat{\tau}^m(X_i) - \tau(X_i)|\right] < \infty \tag{2.7}$$

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 $\phi_n \to \infty$ such that

$$\limsup_{n\to\infty} \sup_{P\in\mathcal{P}_e} \phi_n E_{P^n}\left[\frac{1}{n}\sum_{i=1}^n |\hat{\tau}_i^e - \tau_i|\right] < \infty. \tag{2.8}$$

In Appendix B, we show 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

[9] This claim follows by applying the proof of Theorem 2.1 with the following class of functions:

$$\mathcal{F}^\tau \equiv \{f(X_i; G) \equiv m_0(X_i) + \tau(X_i)\cdot 1\{X_i \in G\} : G \in \mathcal{G}\}.$$

$\mathcal{F}^\tau$ is the VC-subgraph class with the VC-dimension at most v by Lemma A.1 in Appendix A.1.

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}. Theorem 2.5 follows as a corollary of Theorem 2.1 and inequalities (2.5) and (2.6), so we omit a proof.

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,

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

(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,

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

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 of the EWM with known propensity score. Note that if some nonparametric estimator is used to estimate $\tau(\cdot)$ or $e(\cdot)$, ψ_n or ϕ_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 ψ_n^{-1} or ϕ_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, i.e., ψ_n or ϕ_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^*_{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 (2.7) and (2.8) are replaced by*

$$\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_m} E_{P^n}\left[\left(\tilde{\psi}_n \max_{1\leq i\leq n}|\hat{\tau}^m(X_i)-\tau(X_i)|\right)^2\right] < \infty \quad \textit{and} \tag{2.9}$$

$$\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_e} E_{P^n}\left[\left(\tilde{\phi}_n \max_{1\leq i\leq n}|\hat{\tau}^e_i-\tau_i|\right)^2\right] < \infty, \tag{2.10}$$

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

$$\begin{aligned}
\sup_{P \in \mathcal{P}_m \cap \mathcal{P}_{FB}(M,\kappa,\alpha,\eta)} E_{P^n}\left[W(G^*_{FB}) - W(\hat{G}_{m-hybrid})\right] &\leq O\left(\tilde{\psi}_n^{-(1+\alpha)} \vee n^{-\frac{1+\alpha}{2+\alpha}} \log \tilde{\psi}_n\right), \\
\sup_{P \in \mathcal{P}_e \cap \mathcal{P}_{FB}(M,\kappa,\alpha,\eta)} E_{P^n}\left[W(G^*_{FB}) - W(\hat{G}_{e-hybrid})\right] &\leq O\left(\tilde{\phi}_n^{-(1+\alpha)} \vee n^{-\frac{1+\alpha}{2+\alpha}} \log \tilde{\phi}_n\right).
\end{aligned}$$

Theorem 2.6 shows that even when $\tau(\cdot)$ or $e(\cdot)$ have to be estimated, the margin coefficient α influences the rate upper bound of the welfare loss. A higher α 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 $O(\cdot)$, $n^{-\frac{1+\alpha}{2+\alpha}}$. On the other hand, when $\tau(\cdot)$ or $e(\cdot)$ is estimated nonparametrically, which of the two terms in $O(\cdot)$ converges slower depends on the dimension of X and the degree of smoothness of the underlying nonparametric function. See Corollaries 2.1 and 2.2 below for specific expressions of $\tilde{\psi}_n$ and $\tilde{\phi}_n$ when local polynomial regressions are used to estimate $\tau(\cdot)$ or $e(\cdot)$.

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.

2.4.1 Hybrid EWM with Local Polynomial Estimators

In this subsection, we focus on local polynomial estimators for $\tau(x)$ and $e(x)$, and spell out classes of data generating processes $\mathcal{P}_m$ and $\mathcal{P}_e$ as well as ψ_n, $\tilde{\psi}_n$, ϕ_n, and $\tilde{\phi}_n$ that satisfy Condition 2.1 and the assumption of Theorem 2.6.

Consider the m-hybrid approach in which the leave-one-out local polynomial estimators are used to estimate $m_1(X_i)$ and $m_0(X_i)$, i.e., $\hat{m}_1(X_i)$ and $\hat{m}_0(X_i)$ are constructed by fitting the local polynomials excluding the i-th observation.[10] For any multi-index $s = (s_1, \ldots, s_{d_x}) \in \mathbb{N}^{d_x}$ and any $(x_1, \ldots, x_{d_x}) \in \mathbb{R}^{d_x}$, 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 $\|x\| \equiv \left(x_1^2 + \cdots + x_{d_x}^2\right)$. Let $K(\cdot) : \mathbb{R}^{d_x} \to \mathbb{R}$ be a kernel function and $h > 0$ be a bandwidth. At each X_i, $i = 1, \ldots, n$, we define the leave-one-out local polynomial coefficient estimators with

[10] The reason to consider the leave-one-out fitted values is to simplify analytical verification of Condition 2.1. We believe that the welfare loss convergence rates of the hybrid approaches will not be affected even when the i-th observation is included in estimating $\hat{m}_1(X_i)$ and $\hat{m}_0(X_i)$.

degree $l \geq 0$ as

$$\hat{\theta}_1(X_i) = \arg\min_{\theta} \sum_{j \neq i, D_j = 1} \left[Y_j - \theta^T U\left(\frac{X_j - X_i}{h}\right) \right]^2 K\left(\frac{X_j - X_i}{h}\right),$$

$$\hat{\theta}_0(X_i) = \arg\min_{\theta} \sum_{j \neq i, D_j = 0} \left[Y_j - \theta^T U\left(\frac{X_j - X_i}{h}\right) \right]^2 K\left(\frac{X_j - X_i}{h}\right),$$

where $U\left(\frac{X_j - X_i}{h}\right)$ is the vector with elements indexed by the multi-index s, i.e., $U\left(\frac{X_j - X_i}{h}\right) \equiv \left(\left(\frac{X_j - X_i}{h}\right)^s\right)_{|s| \leq l}$.[11] With a slight abuse of notation, we define $U(0) = (1, 0, \ldots, 0)^T$. Let $\lambda_{n,1}(X_i)$ be the smallest eigenvalue of $B_1(X_i) \equiv \left(nh^{d_x}\right)^{-1} \sum_{j \neq i, D_j = 1} U\left(\frac{X_j - X_i}{h}\right) U^T\left(\frac{X_j - X_i}{h}\right) K\left(\frac{X_i - X_j}{h}\right)$ and $\lambda_{n,0}(X_i)$ be the smallest eigenvalue of $B_0(X_i) \equiv \left(nh^{d_x}\right)^{-1} \sum_{j \neq i, D_j = 0} U\left(\frac{X_j - X_i}{h}\right) U^T\left(\frac{X_j - X_i}{h}\right) K\left(\frac{X_i - X_j}{h}\right)$. Accordingly, we construct leave-one-out local polynomial fits for $m_1(X_i)$ and $m_0(X_i)$ by

$$\hat{m}_1(X_i) = U^T(0)\hat{\theta}_1(X_i) \cdot 1\{\lambda_{n,1}(X_i) \geq t_n\},$$

$$\hat{m}_0(X_i) = U^T(0)\hat{\theta}_0(X_i) \cdot 1\{\lambda_{n,0}(X_i) \geq t_n\},$$

where t_n is a positive sequence that slowly converges to zero, such as $t_n \propto (\log n)^{-1}$. These trimming rules regularize the regressor matrices of the local polynomial regressions and simplify the proof of the uniform consistency of the local polynomial estimators.

To characterize $\mathcal{P}_m$ in Condition 2.1, we impose the following restrictions.

Assumption 2.3.

(Smooth-m) Smoothness of the Regressions: The regression equations $m_1(\cdot)$ and $m_0(\cdot)$ belong to a Hölder class of functions with degree $\beta_m \geq 1$ and constant $L_m < \infty$.[12]

(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

$$Leb\left(\mathcal{X} \cap B(x,r)\right) \geq \underline{c} Leb(B(x,r)) \quad \forall 0 < r \leq r_0, \forall x \in \mathcal{X}, \tag{2.11}$$

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 $x \in \mathcal{X}$.

[11] We specify the same bandwidth for these two local polynomial regressions only to suppress notational burden.

[12] Let D^s denote the differential operator $D^s \equiv \frac{\partial^{s_1 + \cdots + s_{d_x}}}{\partial x_1^{s_1} \cdots x_{d_x}^{s_{d_x}}}$. Let $\beta \geq 1$ be an integer. For any $x \in \mathbb{R}^{d_x}$ and any $(\beta - 1)$ times continuously differentiable function $f : \mathbb{R}^{d_x} \to \mathbb{R}$, 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 $f : \mathbb{R}^{d_x} \to \mathbb{R}$ that are $(\beta - 1)$ times continuously differentiable and satisfy, for any x and $x' \in \mathbb{R}^{d_x}$, the inequality $|f_x(x') - f(x)| \leq L \|x - x'\|^{\beta}$.

(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 $\sup_u K(u) \leq K_{\max} < \infty$.

Smoothness of the regression equations, Assumption 2.3 (Smooth-m), is a standard assumption in the context of nonparametric regressions. Assumption 2.3 (PX) is borrowed from Audibert and Tsybakov (2007), and it provides regularity conditions on the marginal distribution of X. Inequality condition (2.11) 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 $r \to 0$ faster than the rate of r^2 for some x in the boundary of $\mathcal{X}$.

Lemma B.4 in Appendix B shows that when $\mathcal{P}_m$ consists of the data generating processes satisfying Assumption 2.3 (Smooth-m) and (PX), Condition 2.1 (m) holds with $\psi_n = n^{\frac{1}{2+d_x/\beta_m}}$, and (2.9) holds with $\tilde{\psi}_n = n^{\frac{1}{2+d_x/\beta_m}} (\log n)^{-\frac{1}{2+d_x/\beta_m}-2}$. The following corollary therefore follows.

Corollary 2.1. *Let $\mathcal{P}_m$ consist of data generating processes that satisfy Assumption 2.3 (Smooth-m) and (PX). Let $\hat{m}_1(X_i)$ and $\hat{m}_0(X_i)$ be the leave-one-out local polynomial estimators with degree $l = (\beta_m - 1)$, whose kernels satisfy Assumption 2.3 (Ker).*

(i) Suppose Assumption 2.1 holds. Then, it holds

$$\sup_{P \in \mathcal{P}_m \cap \mathcal{P}(M,\kappa)} E_{P^n}\left[W^*_{\mathcal{G}} - W(\hat{G}_{m-hybrid})\right] \leq O\left(n^{-\frac{1}{2+d_x/\beta_m}}\right).$$

(ii) Suppose Assumptions 2.1 and 2.2 hold with margin coefficient $\alpha \in (0,1]$. Then, it holds

$$\begin{aligned} &\sup_{P \in \mathcal{P}_m \cap \mathcal{P}_{FB}(M,\kappa,\alpha,\eta)} E_{P^n}\left[W(G^*_{FB}) - W(\hat{G}_{m-hybrid})\right] \\ \leq\ & O\left(n^{-\frac{1+\alpha}{2+d_x/\beta_m}} (\log n)^{\left(\frac{1}{2+d_x/\beta_m}+2\right)(1+\alpha)} \vee n^{-\frac{1+\alpha}{2+\alpha}} \log n\right). \end{aligned}$$

Next, consider the e-hybrid approach. For each $i = 1, \ldots, n$, define a leave-one-out local polynomial fit for propensity score as

$$\begin{aligned} \hat{e}(X_i) &= U^T(0)\hat{\theta}_e(X_i) \cdot 1\{\lambda_n(X_i) \geq t_n\}, \\ \hat{\theta}_e(X_i) &= \arg\min_{\theta} \sum_{j \neq i} \left[D_j - \theta^T U\left(\frac{X_j - X_i}{h}\right)\right]^2 K\left(\frac{X_j - X_i}{h}\right). \end{aligned}$$

We then construct an estimate of individual treatment effect as

$$\hat{\tau}_i = \left[\frac{Y_i D_i}{\hat{e}(X_i)} - \frac{Y_i(1-D_i)}{1-\hat{e}(X_i)}\right] \cdot 1\{\varepsilon_n \leq \hat{e}(X_i) \leq 1-\varepsilon_n\}, \quad 0 < \epsilon_n \leq O\left(n^{-a}\right),\ a > 0,$$

To ensure Condition 2.1 (ii), we now assume smoothness of the propensity score function $e(\cdot)$.

Assumption 2.4. This assumption is the same as Assumption 2.3 except that 2.3 *(Smooth-m)* is replaced by
(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 $L_e < \infty$.

Again, Lemma B.4 in Appendix B shows that $\mathcal{P}_e$ formed by the data generating processes satisfying Assumption 2.4, Condition 2.1 (e) holds with $\phi_n = n^{-\frac{1}{2+d_x/\beta_e}}$ and (2.10) with $\tilde{\phi}_n = n^{\frac{1}{2+d_x/\beta_e}} (\log n)^{-\frac{1}{2+d_x/\beta_e}-2}$.

Corollary 2.2. *Let $\mathcal{P}_e$ consist of data generating processes that satisfy Assumption 2.4 (Smooth-e) and (PX). Let $\hat{e}(X_i)$ be the leave-one-out local polynomial estimator with degree $l = (\beta_e - 1)$, whose kernel satisfy Assumption 2.3 (Ker).*

(i) Suppose Assumption 2.1 holds. Then, it holds

$$\sup_{P \in \mathcal{P}_e \cap \mathcal{P}(M,\kappa)} E_{P^n}\left[W^*_{\mathcal{G}} - W(\hat{G}_{e-hybrid})\right] \leq O\left(n^{-\frac{1}{2+d_x/\beta_e}}\right).$$

(ii) Suppose Assumptions 2.1 and 2.2 hold with margin coefficient $\alpha \in (0,1]$. Then, it holds

$$\begin{aligned} &\sup_{P \in \mathcal{P}_e \cap \mathcal{P}_{FB}(M,\kappa,\alpha,\eta)} E_{P^n}\left[W(G^*_{FB}) - W(\hat{G}_{e-hybrid})\right] \\ \leq\ & O\left(n^{-\frac{1+\alpha}{2+d_x/\beta_e}} (\log n)^{\left(\frac{1}{2+d_x/\beta_e}+2\right)(1+\alpha)} \vee n^{-\frac{1+\alpha}{2+\alpha}} \log n\right). \end{aligned}$$

A comparison of Corollaries 2.1 and 2.2 shows that the rate upper bound of welfare loss differs between the m-hybrid EWM and the e-hybrid EWM approaches when the degree of Hölder smoothness of the regression equations β_m and that of the propensity score β_e are different. For instance, if the propensity score $e(\cdot)$ is smoother than the regression equations of outcome $m_1(\cdot)$ and $m_0(\cdot)$ in the sense of $\beta_e > \beta_m$ and the degree of local polynomial regressions is chosen accordingly, then the rate upper bound of the e-hybrid EWM rule converges faster than that of the m-hybrid EWM rule.

3 Inference for Welfare

In the proposed EWM procedure, the maximized empirical welfare $W_n(\hat{G}_{EWM})$ can be seen as an estimate of $W(\hat{G}_{EWM})$, the welfare level attained by implementing the estimated treatment rule.[13] In this section, we provide a procedure for constructing asymptotically valid confidence intervals for the population welfare gain of implementing the estimated rule.

Let $\hat{G} \in \mathcal{G}$ be an estimated treatment rule such as $\hat{G}_{EWM}$, $\hat{G}_{m-hybrid}$, and $\hat{G}_{e-hybrid}$. Define the welfare gain of implementing an estimated treatment rule $\hat{G} \in \mathcal{G}$ by

$$V(\hat{G}) = W(\hat{G}) - W(G_0),$$

where G_0 is a benchmark treatment assignment rule with which the estimated treatment rule $\hat{G}$ is compared in terms of the social welfare. For instance, if the estimated treatment rule $\hat{G}$ is compared with the "no treatment" case, G_0 is the empty set $\emptyset$. Alternatively, if a benchmark policy is the non-individualized uniform adoption of the treatment, G_0 is set at $G_0 = \mathcal{X}$, and $V(\hat{G})$ is interpreted as the welfare gain of implementing individualized treatment assignment instead of the non-individualized implementation of the treatment.

A construction of the confidence intervals for $V(\hat{G})$ proceeds as follows. Let $\nu_n(G) = \sqrt{n}(V_n(G) - V(G))$, where $V_n(G) \equiv W_n(G) - W_n(G_0)$. If there is a random variable $\bar{\nu}_n$ such that $\nu_n\left(\hat{G}\right) \leq \bar{\nu}_n$ holds P^n-almost surely, and if $\bar{\nu}_n$ converges in distribution to a non-degenerate random variable $\bar{\nu}$, then, with $c_{1-\tilde{\alpha}}$, the $(1-\tilde{\alpha})$-th quantile of $\bar{\nu}$, it holds

$$P^n(\nu_n\left(\hat{G}\right) \leq c_{1-\tilde{\alpha}}) \geq P^n(\bar{\nu}_n \leq c_{1-\tilde{\alpha}}) \to \Pr(\bar{\nu} \leq c_{1-\tilde{\alpha}}) = 1-\tilde{\alpha}, \text{ as } n \to \infty.$$

Hence, if $\hat{c}_{1-\tilde{\alpha}}$, a consistent estimator of $c_{1-\tilde{\alpha}}$, is available, an asymptotically valid one-sided confidence intervals for $V(\hat{G})$ with coverage probability $(1-\tilde{\alpha})$ can be given by

$$[V_n\left(\hat{G}\right) - \frac{\hat{c}_{1-\tilde{\alpha}}}{\sqrt{n}}, \infty). \tag{3.1}$$

In the algorithm summarized below, we specify $\bar{\nu}_n$ to be $\bar{\nu}_n = \sqrt{n}\sup_{G\in\mathcal{G}}(V_n(G) - V(G))$ and estimate $\hat{c}_{1-\tilde{\alpha}}$ by bootstrapping the supremum of the centered empirical processes.[14]

[13] It is important to note that in finite samples, $W_n\left(\hat{G}_{EWM}\right)$ estimates $W(\hat{G}_{EWM})$ with an upward bias. With fixed n, the size of the bias becomes bigger as $\mathcal{G}$ becomes more complex.

[14] The current choice of $\bar{\nu}_n$ is likely to yield conservative confidence intervals. Keeping the same nominal coverage probability, it is feasible to tighten up the confidence intervals with a more sophisticated choice of $\bar{\nu}_n$, such as, $\bar{\nu}_n = \sqrt{n}\sup_{G\in\hat{\mathcal{G}}}(V_n(G) - V(G))$, where $\hat{\mathcal{G}}$ is a data-dependent subclass of $\mathcal{G}$ that contains $\hat{G}$ with probability approaching one.

Algorithm 3.1. *1. Let $\hat{G} \in \mathcal{G}$ be an estimated treatment assignment rule (e.g., EWM rule, the hybrid-rules, etc.), and $V_n(\cdot) = W_n(\cdot) - W_n(G_0)$ be the empirical welfare gain obtained from the original sample.*

2. Resample n-observations of $Z_i = (Y_i, D_i, X_i)$ randomly with replacement from the original sample and construct the bootstrap analogue of the welfare gain, $V_n^*(\cdot) = W_n^*(\cdot) - W_n^*(G_0)$, where $W_n^*(\cdot)$ is the empirical welfare of the bootstrap sample.

3. Compute $\bar{\nu}_n^* = \sqrt{n} \sup_{G \in \mathcal{G}} (V_n^*(G) - V_n(G))$.

4. Let $\tilde{\alpha} \in (0, 1/2)$. Repeat step 2 and 3 many times and estimate $\hat{c}_{1-\tilde{\alpha}}$ by the empirical $(1-\tilde{\alpha})$-th quantile of the bootstrap realizations of $\bar{\nu}_n^*$.

Given Assumption 2.1, the uniform central limit theorem for empirical processes assures that $\bar{\nu}_n$ converges in distribution to $\bar{\nu}$ the supremum of a mean zero Brownian bridge process. Furthermore, by the well-known result on the asymptotic validity of the bootstrap empirical processes (see, e.g., Section 3.6 of van der Vaart and Wellner (1996)), the bootstrap critical value $\hat{c}_{1-\tilde{\alpha}}$ consistently estimates the corresponding quantile of $\bar{\nu}$. We can therefore conclude that the confidence intervals constructed in (3.1) has the desired asymptotic coverage probability.

4 Extensions

4.1 Empirical Welfare Maximization with a Capacity Constraint

Empirical welfare maximization could also be used to select treatment rules when one of the treatments is scarce and the planner faces a capacity constraint on the proportion of the target population that could be assigned to it. Capacity constraints exist in various treatment choice problems. In medicine, stocks of new drugs and vaccines could be smaller than the number of patients who may benefit from them. In education, limited number of slots is available in "magnet" schools. Training programs for the unemployed are sometimes capacity-constrained. The limited capacity of a prison system could make it infeasible to assign incarceration as a treatment for all convicted criminals. In these cases, treatment assignment rules which propose to assign too many individuals to the capacity-constrained treatment cannot be fully implemented.

We assume that the availability of treatment 1 is constrained.

Assumption 4.1.
(CC) Capacity Constraint: 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 all previous results could be applied simply by replacing $\mathcal{G}$ with $\mathcal{G}_K$.

The population distribution of covariates is often not known precisely. In these cases, it is impossible to guarantee with certainty that estimated treatment rule $\hat{G}$ will satisfy the capacity constraint with probability one. To evaluate the welfare of any treatment assignment method in this setting, we first need to make an assumption about what happens when the estimated treatment set $\hat{G}$ violates the capacity constraint.

We assume that if the treatment rule G violates the capacity constraint, i.e., $P_X(G) > K$, 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). 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 assumption holds if individuals arrive sequentially for treatment assignment in the order that is independent of (X, Y_0, Y_1) and receive treatment 1 on a first-come first-serve basis if $X \in G$. It could also correspond to settings in which treatment 1 is allocated by lottery to individuals with $X \in G$.

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

$$W_K(G) \equiv E_P \left[\begin{array}{c} \left[m_1(X) \cdot \min\left\{1, \frac{K}{P_X(G)}\right\} + m_0(X) \cdot \left(1 - \min\left\{1, \frac{K}{P_X(G)}\right\}\right)\right] \cdot 1\{X \in G\} \\ + m_0(X) \cdot 1\{X \notin G\} \end{array} \right].$$

Then the capacity-constrained welfare gain of the treatment rule G equals

$$\begin{aligned} V_K(G) &\equiv W_K(G) - W_K(\emptyset) \\ &= E_P\left[\min\left\{1, \frac{K}{P_X(G)}\right\} \cdot \tau(X) \cdot 1\{X \in G\}\right] \\ &= \min\left\{1, \frac{K}{P_X(G)}\right\} \cdot V(G). \end{aligned}$$

Rationing dilutes the effect of treatment rules that violate the capacity constraint and we take into account this effect on welfare. 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.

The maximum welfare gain attainable by treatment rules in $\mathcal{G}$ in the presence of a capacity constraint equals $V_K^* \equiv \sup_{G\in\mathcal{G}} V_K(G)$. Even with full knowledge of the outcome distribution, it is feasible that the optimal policy would assign the scarce treatment to a subset of the population with $P_X(G) > K$, requiring rationing within that subpopulation. This is unlikely to happen in practice when the distribution of covariates does not have atoms and the collection of treatment rules $\mathcal{G}$ is sufficiently rich. The following condition, for example, guarantees that the optimal capacity-constrained policy does not require rationing.

Remark 4.1. *If the collection of treatment rules $\mathcal{G}$ contains the upper contour sets of $\tau(x)$,*

$$G_t \equiv \{x \in \mathcal{X} : \tau(x) \geq t\}, t \in \mathbb{R},$$

and $P_X(G_t)$ is continuous in t, then the optimal policy $G_K = \arg\max_{G\in\mathcal{G}} V_K(G)$ belongs to the set $\{G_t, t \in \mathbb{R}\}$ and satisfies the capacity constraint without rationing: $P_X(\hat{G}_K) \leq K$.

We propose a treatment rule that maximizes the empirical analog of the capacity-constrained welfare gain $V_K(G)$ (and, hence, welfare):

$$\hat{G}_K \equiv \arg\max_{G\in\mathcal{G}} V_{K,n}(G), \tag{4.1}$$

where

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

and $P_{X,n}$ is the empirical probability distribution of $(X_1, \ldots, X_n)$. The following theorem shows that the expected welfare of $\hat{G}_K$ converges to the maximum at least at $n^{-1/2}$ rate. The result is analogous to Theorem 2.1, with the additional term corresponding to potential welfare losses due to misestimation of $P_X(G)$.

Theorem 4.1. *Under Assumptions 2.1 and 4.1,*

$$\sup_{P\in\mathcal{P}(M,\kappa)} E_{P^n}\left[\sup_{G\in\mathcal{G}} W_K(G) - W_K(\hat{G}_K)\right] \leq C_1 \frac{M}{\kappa}\sqrt{\frac{v}{n}} + C_1 \frac{M}{K}\sqrt{\frac{v}{n}},$$

where C_1 is the universal constant in Lemma A.4.

4.2 Target Population that Differs from the Sampled Population

Empirical Welfare Maximization method can be adapted to select treatment rules for a target population that differs from the sampled population in the distribution of covariates X, but has the same conditional treatment effect function $\tau(x)$.

As before, P denotes the probability distribution of $(Y_{0,i}, Y_{1,i}, D_i, X_i)$ in the sampled population. We denote by P^T the probability distribution of $(Y_{0,i}, Y_{1,i}, X_i)$ in the target population and by E^T the expectations with respect to that probability. The welfare of implementing treatment rule G in the target population is $W^T(G)$. We assume that the sampled population has the same conditional treatment effect as the target population.

Assumption 4.2.
(ID) Identical Treatment Effects: $E^T(Y_1 - Y_0|X) = E_P(Y_1 - Y_0|X)$.
(BDR) Bounded Density Ratio: Probability distributions P_X^T and P_X have densities p_X^T and p_X with respect to a common dominating measure on $\mathcal{X}$ and for some $\rho(x) \leq \bar{\rho} < \infty$:

$$p_X^T(x) = \rho(x) \cdot p_X(x).$$

The welfare gain of treatment rule G on the target population equals

$$V^T(G) \equiv \int_{\mathcal{X}} \tau(x) 1\{x \in G\} \, dP_X^T(x) = \int_{\mathcal{X}} \tau(x) 1\{x \in G\} \rho(x) \, dP_X(x).$$

Assumption 4.2 (ID) implies that the first-best treatment rule $G_{FB}^* = 1\{x : \tau(x) \geq 0\}$ is the same in the sampled and the target populations. Therefore, when the first-best policy is feasible, i.e. $G_{FB}^* \in \mathcal{G}$, we could directly apply the EWM treatment rule $\hat{G}_{EWM}$ computed for the sampled population to the target population. Note that from the definition of G_{FB}^*, it follows that for any treatment rule G and any x

$$\tau(x) 1\{x \in G_{FB}^*\} - \tau(x) 1\{x \in G\} \geq 0.$$

In conjunction with Assumption 4.2 (BDR), this yields an upper bound on the welfare loss of applying any treatment rule in the target population expressed in terms of the welfare loss of

applying the same rule in the sampled population.

$$\begin{aligned} W^T(G^*_{FB}) - W^T(G) &= V^T(G^*_{FB}) - V^T(G) \\ &= \int_{\mathcal{X}} \left[\tau(x)1\{x \in G^*_{FB}\} - \tau(x)1\{x \in G\}\right] \rho(x) dP_X(x) \\ &\leq \bar{\rho} \int_{\mathcal{X}} \left[\tau(x)1\{x \in G^*_{FB}\} - \tau(x)1\{x \in G\}\right] dP_X(x) \\ &= \bar{\rho} \left[W(G^*_{FB}) - W(G)\right]. \end{aligned}$$

All of the welfare convergence rate results derived for EWM and hybrid treatment rules in previous sections also hold when these treatment rules are applied to the target population, as long as Assumptions 2.2 (FB) and 4.2 hold.

When the first-best policy is not feasible ($G^*_{FB} \notin \mathcal{G}$), the second-best policies for the sampled and the target populations, $G^* \in \arg\max_{G\in\mathcal{G}} W(G)$ and $G^{T*} \in \arg\max_{G\in\mathcal{G}} W^T(G)$, are generally different. The welfare of treatment rules proposed in the previous sections does not generally converge to the second-best in the target population $\sup_{G\in\mathcal{G}} W^T(G)$.

The second-best in the target population could be obtained by reweighting the argument of the EWM problem by the density ratio $\rho(X_i)$. This method works even when the first-best treatment rule is not feasible and could be combined with a treatment capacity constraint. The reweighted EWM problem is

$$\hat{G}^T_{EWM} \in \arg\max_{G\in\mathcal{G}} E_n \left[\left(\frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)}\right) \cdot \rho(X_i) \cdot 1\{X_i \in G\}\right]. \tag{4.2}$$

Applying the law of iterated expectations with respect to X, we obtain

$$E_P \left[\left(\frac{YD}{e(X)} - \frac{Y(1-D)}{1-e(X)}\right) \cdot \rho(X) \cdot 1\{X \in G\}\right] = \int_{\mathcal{X}} \tau(x)1\{x \in G\}\rho(x)\ dP_X(x) = V^T(G). \tag{4.3}$$

Theorem 4.2 shows that the welfare loss of the reweighted EWM treatment rule in the target population converges to zero at least at $n^{-1/2}$ rate.

Theorem 4.2. *If the distribution P in the sampled population satisfies Assumption 2.1 and the distribution P^T in the target population satisfies Assumption 4.2, then*

$$\sup_{P\in\mathcal{P}(M,\kappa)} E_{P^n}\left[\sup_{G\in\mathcal{G}} W^T(G) - W^T(\hat{G}^T_{EWM})\right] \leq C_1 \frac{M\bar{\rho}}{\kappa}\sqrt{\frac{v}{n}},$$

where C_1 is the universal constant in Lemma A.4.

4.3 Comparison with the Nonparametric Plug-in Rule

The plug-in treatment choice rule (1.13) with parametrically or nonparametrically estimated $m_1(x)$ and $m_0(x)$ is intuitive and simple to implement. In situations where flexible treatment assignment rules are allowed and the dimension of conditioning covariates is small, the nonparametric plug-in rule would be a competing alternative to the EWM approach. In this section, we review the welfare loss convergence rate results of the nonparametric plug-in rule and discuss potential advantages and disadvantages of these two approaches.

We denote the class of data generating processes that satisfy Assumptions 2.1 (UCF), (BO), (SO), Assumption 2.2 (MA), and Assumption 2.3 by $\mathcal{P}_{smooth}(M, \kappa, \alpha, \eta, \beta_m)$. Given the smoothness assumption of the regression equations, we consider estimating m_1 and m_0 by local polynomial estimators of degree $(\beta - 1)$. The convergence rate results of the nonparametric plug-in classifiers shown in Theorem 3.3 of Audibert and Tsybakov (2007) can be straightforwardly extended to the treatment choice context, resulting in

$$\sup_{P \in \mathcal{P}_{smooth}(M,\kappa,\alpha,\eta,\beta_m)} E_{P^n}\left[W(G^*_{FB}) - W(\hat{G}_{plug-in})\right] \leq O\left(n^{-\frac{1+\alpha}{2+d_x/\beta_m}}\right). \tag{4.4}$$

Furthermore, if $\alpha\beta \leq d_x$, Theorem 3.5 of Audibert and Tsybakov (2007) applied to the current treatment choice setup shows that the nonparametric plug-in rule attains the rate lower bound i.e., for any treatment rule $\hat{G}$,

$$\sup_{P \in \mathcal{P}_{smooth}(M,\kappa,\alpha,\eta,\beta_m)} E_{P^n}\left[W(G^*_{FB}) - W(\hat{G})\right] \geq O\left(n^{-\frac{1+\alpha}{2+d_x/\beta_m}}\right)$$

holds.

In practically relevant situations where $\alpha\beta \leq d_x$,[15] a naive comparison of the welfare loss convergence rate of the plug-in rule presented here with that of EWM (Theorems 2.3 and 2.4) would suggest that in terms of the welfare loss converge rate, the EWM rule would outperform the nonparametric plug-in rule. It is, however, important to notice that the classes of data generating processes over which the uniform rates are ensured differ between the two cases. $\mathcal{P}_{smooth}(M, \kappa, \alpha, \eta, \beta)$ is constrained by smooth regression equations and continuously distributed X, whereas $\mathcal{P}_{FB}(M, \kappa, \alpha, \eta)$

[15] In an analogy to the Proposition 3.4 of Audibert and Tsybakov (2007), when the class of data generating processes is assumed to have $\alpha\beta > d_x$, no data generating process in this class can have the conditional treatment effect $\tau(x) = 0$ in an interior of the support of P_X. In the practice of causal inference, we a priori would not restrict the plausible data generating processes only to these extreme cases; therefore, the class of data generating processes with $\alpha\beta > d_x$ would be less relevant in practice.

considered in Theorems 2.3 and 2.4 allows for discontinuous regression equations and no restriction on the marginal distribution of X's. Assumption 2.2 (FB) on $\mathcal{P}_{FB}(M, \kappa, \alpha, \eta)$ requires that $\{x : \tau(x) \geq 0\}$ belongs to the pre-specified VC-class $\mathcal{G}$, whereas $\mathcal{P}_{smooth}(M, \kappa, \alpha, \eta, \beta)$ is free from such assumption. This non-nested relationship between $\mathcal{P}_{FB}(M, \kappa, \alpha, \eta)$ and $\mathcal{P}_{smooth}(M, \kappa, \alpha, \eta, \beta)$ makes the naive rate comparison between (4.4) and Theorem 2.3 less meaningful because a data generating process in $\mathcal{P}_{smooth}(M, \kappa, \alpha, \eta, \beta)$ that yields the slowest convergence rate for the nonparametric plug-in rule is in fact excluded from $\mathcal{P}_{FB}(M, \kappa, \alpha, \eta)$. Accordingly, unless we can assess which one of $\mathcal{P}_{smooth}(M, \kappa, \alpha, \eta, \beta)$ and $\mathcal{P}_{FB}(M, \kappa, \alpha, \eta)$ is more likely to contain the true data generating process, these rate results offer us limited guidance on the procedure that should be used in a given application.

In practical terms, we consider these two distinct approaches as complementary, and our choice between them should be based on available assumptions and the dimension of covariates in a given application. A practical advantage of the EWM rule is that the welfare loss convergence rate does not directly depend on the dimension of X, so when an available credible assumption on the level set $\{x : \tau(x) \geq 0\}$ implies a certain class of decision sets with a finite VC-dimension, the EWM approach offers a practical solution to get around the curse of dimensionality of X. A potential drawback of using the EWM rule is the risk of misspecification of $\mathcal{G}$, i.e., if Assumption 2.2 (FB) is not valid, the EWM rule only attains the second-best welfare, whereas the nonparametric plug-in rule is guaranteed to yield the first-best welfare in the limit. Another aspect of comparison is that the performance of the EWM rule is stable regardless of whether the underlying data generating processes, including the marginal distribution of X and the regression equations $m_1(X)$ and $m_0(X)$, are smooth. Furthermore, in terms of implementation, the EWM approach does not require the user to specify smoothing parameters once the class of candidate decision sets $\mathcal{G}$ is given. In contrast, the nonparametric plug-in rule requires smoothing parameters. The statistical performance of the nonparametric plug-in rule can be sensitive to the choice of smoothing parameters, and the theoretical results of the convergence rate given in (4.4) assume the user's ability to choose the smoothing parameter properly.

5 Computing EWM Treatment Rules

The Empirical Welfare Maximization estimator $\hat{G}$, as well as hybrid estimators $\hat{G}_{m-hybrid}$, and $\hat{G}_{e-hybrid}$, share the same structure

$$\hat{G} \in \arg\max_{G \in \mathcal{G}} \sum_{1 \leq i \leq n} g_i \cdot 1\{X_i \in G\}, \tag{5.1}$$

where each g_i is a function of the data, i.e., for the EWM rule $\hat{G}_{EWM}$, $g_i = \frac{1}{n}\left(\frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)}\right)$, for the e-hybrid rule $\hat{G}_{e-hybrid}$, $g_i = \hat{\tau}_i^e/n$, and for the m-hybrid rule $\hat{G}_{m-hybrid}$, $g_i = \hat{\tau}^m(X_i)/n$. The objective function in (5.1) is non-convex and discontinuous in G, thus finding $\hat{G}$ could be computationally challenging. In this section, we propose a set of convenient tools that permit solving this optimization problem and performing inference using widely available software[16] for practically important classes of sets $\mathcal{G}$ defined by linear eligibility scores.

5.1 Single Linear Index Rules

We start with the problem of computing optimal treatment rules that assign treatments based on a linear index (linear eligibility score; LES, see Examples 2.1 and 2.2). To reduce notational complexity, we include a constant in the covariate vector X throughout the exposition of this section. An LES rule can be expressed as $1\{X^T\beta \geq 0\}$. This type of treatment rule is commonly used in practice because it offers a simple way to reduce the dimension of observable characteristics. Furthermore, it is easy to enforce monotonicity of treatment assignment in specific covariates by imposing sign restrictions on the components of β.

Let $\mathcal{G}_{LES}$ be a collection of half-spaces of the covariate space $\mathcal{X}$, which are the upper contour sets of linear functions:

$$\begin{aligned} \mathcal{G}_{LES} &= \left\{G_\beta : \beta \in \mathbf{B} \subset \mathbb{R}^{d_x+1}\right\}, \\ G_\beta &= \left\{x : x^T\beta \geq 0\right\}. \end{aligned}$$

Then the optimization problem (5.1) becomes:

$$\max_{\beta \in \mathbf{B}} \sum_{1 \leq i \leq n} g_i \cdot 1\left\{X_i^T\beta \geq 0\right\}. \tag{5.2}$$

This problem is similar to the maximum weighted score problem analyzed in Florios and Skouras (2008). They observe that the maximum score objective function could be rewritten as a Mixed

[16] For the empirical illustration we used IBM ILOG CPLEX Optimization Studio, which is available free for academic use through the IBM Academic Initiative.

Integer Linear Programming problem with additional binary parameters $(z_1, ..., z_n)$ that replace the indicator functions $1\left\{X_i^T\beta \geq 0\right\}$. The equality $z_i = 1\left\{X_i^T\beta \geq 0\right\}$ is imposed by a combination of linear inequality constraints and the restriction that z_i's are binary. The advantage of a MILP representation is that it is a standard optimization problem that could be solved by multiple commercial and open-source solvers. The branch-and-cut algorithms implemented in these solvers are faster than brute force combinatorial optimization.

We propose replacing (5.2) by its equivalent problem:

$$\max_{\substack{\beta\in\mathbf{B},\\ z_1,\ldots,z_n\in\mathbb{R}}} \quad \sum_{1\leq i\leq n} g_i \cdot z_i \tag{5.3}$$

$$\text{s.t.} \quad \frac{X_i^T\beta}{C_i} < z_i \leq 1 + \frac{X_i^T\beta}{C_i} \text{ for } i = 1, \ldots, n, \tag{5.4}$$
$$z_i \in \{0, 1\},$$

where constants C_i should satisfy $C_i > \sup_{\beta\in\mathbf{B}} |X_i^T\beta|$. Then the inequality constraints (5.4) and the restriction that z_i's are binary imply that $z_i = 1$ if and only if $X_i^T\beta \geq 0$. It follows that the maximum value of (5.4) for each value of β is the same as the value of (5.2).

The problem (5.3) is a linear optimization problem with linear inequality constraints and integer constraints on z_i's if the set $\mathbf{B}$ is defined by linear inequalities that could be passed to any MILP solver. Florios and Skouras (2008) impose only one side of the inequality constraint (5.4) for each i. For $g_i > 0$, it is sufficient to impose only the upper bound on z_i and for $g_i < 0$ only the lower bound. The other side of the bound is always satisfied by the solution due to the direction of the objective function.

Our formulation has significant advantages. Despite a larger number of inequalities, it reduces the computation time in our applications by a factor of 10-40. Furthermore, it is not sufficient to impose only one side of the inequalities on z_i's for optimization with a capacity constraint considered further below.

Inference on the welfare gain $V(\hat{G}_{EWM})$ of the empirical welfare maximizing policy requires computing $\bar{\nu}_n^* = \sup_{G\in\mathcal{G}} \sqrt{n}\,(V_n^*(G) - V_n(G))$ in each bootstrap sample. Denoting the bootstrap weights by $\{w_i^*\}$, $\sum_{i=1}^n w_i^* = n$, $\bar{\nu}_n^*$ could be expressed as

$$\bar{\nu}_n^* = \sqrt{n} \sup_{G\in\mathcal{G}} \sum_{1\leq i\leq n} (w_i^* - 1) g_i \cdot 1\left\{X_i^T\beta \geq 0\right\} \tag{5.5}$$

The optimization problem for $\bar{\nu}_n^*$ is analogous to the optimization problem for $\hat{G}_{EWM}$. Furthermore, solving it does not require the knowledge of $\hat{G}_{EWM}$, hence all bootstrap computations could be performed in parallel with the main EWM problem.

5.2 Multiple Linear Index Rules

We extend this method to compute treatment rules based on multiple linear scores. These rules construct J scores that are linear in covariates (or in their functions) and assign an individual to treatment if each score exceeds a specific threshold. An example of a multiple index treatment rule with three indices is when an individual is assigned to a job training program if ($25 \leq$ age ≤ 35) AND (wage at the previous job $<$ \$15). The results are easily extended to treatment rules that apply if any of the indices exceeds its threshold, for example, (age ≥ 40) OR (length of unemployment $\geq$ 2 years).

Let the treatment assignment set G be defined as an intersection of upper contour sets of J linear functions:

$$\begin{aligned} \mathcal{G} &= \left\{ G_{\beta^1,...,\beta^J}, \beta^1, ..., \beta^J \in \mathbf{B} \right\}, \\ G_{\beta^1,...,\beta^J} &= \left\{ x : x^T\beta^1 \geq 0, ..., x^T\beta^J \geq 0 \right\}. \end{aligned}$$

Then the optimization problem (5.1) becomes

$$\max_{\beta^1,...,\beta^J \in \mathbf{B}} \sum_{1 \leq i \leq n} g_i \cdot 1\{X_i^T\beta^1 \geq 0, \ldots, X_i^T\beta^J \geq 0\}. \tag{5.6}$$

We propose its equivalent formulation as a MILP problem with auxiliary binary variables $\left\{(z_i^1, \ldots, z_i^J, z_i^*), i = 1, \ldots, n\right\}$:

$$\max_{\substack{\beta^1,...,\beta^J \in \mathbf{B}, \\ z_i^1,...,z_i^J, z_i^* \in \mathbb{R}}} \quad \sum_{1 \leq i \leq n} g_i \cdot z_i^* \tag{5.7}$$

$$\text{s.t.} \quad \frac{X_i^T\beta^j}{C_i} < z_i^j \leq 1 + \frac{X_i^T\beta^j}{C_i} \text{ for } 1 \leq i \leq n, 1 \leq j \leq J, \tag{5.8}$$

$$1 - J + \sum_{1 \leq j \leq J} z_i^j \leq z_i^* \leq J^{-1} \sum_{1 \leq j \leq J} z_i^j \text{ for } 1 \leq i \leq n, \tag{5.9}$$

$$z_i^1, \ldots, z_i^J, z_i^* \in \{0, 1\} \text{ for } 1 \leq i \leq n.$$

Similarly to the single index problem, the inequalities (5.8) and the constraint that z_i^j's are binary imply together that $z_i^j = 1\{X_i^T\beta^j \geq 0\}$. Linear inequalities (5.9) and the binary constraints imply together that

$$z_i^* = z_i^1 \cdot ... \cdot z_i^J = 1\{X_i^T\beta^1 \geq 0\} \cdot ... \cdot 1\{X_i^T\beta^J \geq 0\}.$$

The problem for a collection of sets defined by the union of linear inequalities

$$G_{\beta^1,...,\beta^J} = \left\{ X : X^T\beta^1 \geq 0 \text{ or } \ldots \text{or } X^T\beta^J \geq 0 \right\}$$

could also be written as a MILP problem with the inequality constraint (5.9) replaced by

$$J^{-1} \sum_{1 \le j \le J} z_i^j \le z_i^* \le \sum_{1 \le j \le J} z_i^j \text{ for } i = 1, \ldots, n. \tag{5.10}$$

5.3 Optimization with a Capacity Constraint

When there is a capacity constraint K on the proportion of population that could be assigned to treatment 1, Empirical Welfare Maximization problem (4.1) on a set $\mathcal{G}$ of half-spaces becomes

$$\max_{\beta \in \mathbf{B}} \left[\min \left\{ 1, \frac{Kn}{\sum_{i=1}^n 1\{X_i^T \beta \ge 0\}} \right\} \sum_{1 \le i \le n} g_i \cdot 1\left\{ X_i^T \beta \ge 0 \right\} \right]. \tag{5.11}$$

This problem cannot be rewritten as a linear optimization problem in the same way as (5.3) because the factor $\min \left\{ 1, \frac{Kn}{\sum_{i=1}^n 1\{X_i^T \beta \ge 0\}} \right\}$ varies with β. This factor could take fewer than n different values and the maximum of (5.11) could be obtained by solving a sequence of optimization problems each of which holds this factor constant.

$$\begin{array}{rl} \text{For} & k = \lfloor Kn \rfloor, \ldots, n \\ \max\limits_{\substack{\beta \in \mathbf{B}, \\ z_1, \ldots, z_n \in \mathbb{R}}} & \min \left\{ 1, \dfrac{Kn}{k} \right\} \sum\limits_{1 \le i \le n} g_i \cdot z_i \\ \text{s.t.} & \dfrac{X_i^T \beta}{C_i} < z_i \le 1 + \dfrac{X_i^T \beta}{C_i} \text{ for } 1 \le i \le n, \\ & z_i \in \{0, 1\}, \\ & \sum\limits_{1 \le i \le n} z_i \le k. \end{array}$$

The capacity constrained problem with multiple indexes could be solved similarly.

6 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 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 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. We use the same sample of 11,204 adults (22 years and older) used in the original evaluation of the program and in the subsequent studies (Bloom et al., 1997, Heckman et al., 1997, Abadie et al., 2002). The probability of being assigned to the treatment was one third in this sample.

We use two simple welfare outcome measures for our illustration. The first is the total individual earnings in the 30-month period following the program assignment. The second outcome measure is the 30-month earnings minus \$1,000 cost for assigning individuals to the treatment, which is close to the average cost (per assignee) of the additional services provided by the JTPA program, as estimated by Bloom et al. (1997). The first outcome measure reflects social preferences that put no weight on the costs of the program incurred by the government. The second measure weighs participants' gains and the government's losses equally.

For all treatment rules, we report the estimated *intention-to-treat* effect of assigning all individuals with covariates $X \in G$ to the treatment. The take-up of different program services varied across individuals. We view the policy maker's problem as a choice of eligibility criteria for the program and not a choice of the take-up rate (which is 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 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.

Pre-treatment variables based on which we consider conditioning the treatment assignment are 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, since this may be easily perceived as discrimination. Education and earnings are generally verifiable characteristics. This is an important feature for implementing the proposed treatment assignment because the empirical welfare estimates are inaccurate for the target population if the individuals could manipulate their characteristics to obtain the desired treatment.

Table 1 reports the estimated welfare gains of alternative treatment rules. All of them are estimated by inverse probability weighting. The average effect of the program on 30-month earnings for the whole study population is \$1,269. If treatment cost of \$1,000 per assignee is taken into account, the net average effect of assigning everyone to treatment is estimated to be \$269.

We consider two candidate classes of treatment rules for EWM. The first is the class of *quadrant*

treatment rules:

$$\mathcal{G}_Q \equiv \left\{ \begin{array}{c} \{x : s_1(\text{education} - t_1) > 0 \ \& \ s_2(\text{prior earnings} - t_2) > 0\}, \\ s_1, s_2 \in \{-1, 0, 1\}, t_1, t_2 \in \mathbb{R} \end{array} \right\}. \tag{6.1}$$

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 have to be above (or below) some specific thresholds.

Figure 1 demonstrates the quadrant treatment rules selected by the EWM criterion. The entire shaded area covers individuals who would be assigned to treatment if it were costless. The dark shaded area shows the EWM treatment rule that takes into account \$1,000 treatment cost. The size of black dots indicates the number of individuals with different covariate values. Both rules set a minimum threshold of ten years of education and a maximum threshold on pre-program earnings (\$12,200 and \$6,500). The estimated proportions of population assigned to treatment under these rules are 83% and 73%.

Second, we consider the class of *linear treatment rules:*

$$\mathcal{G}_{LES} \equiv \{\{x : \beta_0 + \beta_1 \cdot \text{education} + \beta_2 \cdot \text{prior earnings} > 0\}, \beta_0, \beta_1, \beta_2 \in \mathbb{R}\}. \tag{6.2}$$

Figure 2 displays the treatment rules from this class chosen according to the EWM criterion. They are nearly identical for no treatment cost and for a cost of \$1,000, assigning 82% of the population to treatment. At higher treatment costs, EWM selects a much smaller subset of the population.

Linear treatment rules that maximize empirical welfare are markedly different from the plug-in rule derived from linear regressions, which are shown in Figure 3. Without treatment costs, linear regression predicts positive treatment effects for the entire range of feasible covariate values. With a cost of \$1,000, the regression predicts positive net treatment effect for about 82% of individuals. Noticeably, the direction of the treatment assignment differs between regression plug-in and linear EWM rules. The regression puts a positive coefficient on prior earnings, whereas the equation characterizing linear EWM rule puts a negative coefficient on them. If the linear regression is correctly specified, the regression plug-in and EWM rules have identical large sample limits. If the regression is misspecified, however, only linear EWM treatment rules converge with sample size to the welfare-maximizing limit. The welfare yielded by regression plug-in rules converges to a lower limit with sample size.

Figure 4 shows plug-in treatment rules based on Kernel regressions of treatment and control outcomes on the covariates. The bandwidths were chosen by Silverman's rule of thumb. The class of nonparametric plug-in rules is 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.

7 Conclusion

The EWM approach proposed in this paper directly maximizes a sample analog of the welfare criterion of a utilitarian policy maker. This welfare-function-based statistical procedure for 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 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. Second, we ruled out the case in which the data are subject to selection on unobservables. With self-selection into the treatment, the welfare criterion could be only set-identified, and it is not clear how to extend the EWM idea to this case. Third, we restricted our analysis to the utilitarian social welfare criterion, but in some contexts, policy makers have a non-utilitarian social welfare criterion. We leave these issues for future research.

Table 1: Estimated welfare gain of alternative treatment assignment rules that condition on education and pre-program earnings.

Treatment rule:	Share of population assigned to treatment	Empirical welfare gain per population member	Lower 90% CI (bootstrap)
Outcome variable: 30-month post-program earnings. No treatment cost.			
Average treatment effect:	1	\$1,269	\$719
EWM quadrant treatment rule	0.827	\$1,614	\$787
EWM linear treatment rule	0.820	\$1,657	\$765
Linear regression plug-in rule	1	\$1,269	
Nonparametric plug-in rule	0.816	\$1,924	
With \$1000 expected cost per treatment assignment			
Average treatment effect:	1	\$269	-\$281
EWM quadrant treatment rule	0.732	\$797	-\$20
EWM linear treatment rule	0.819	\$837	-\$46
Linear regression plug-in rule	0.824	\$760	
Nonparametric plug-in rule	0.675	\$1,164	

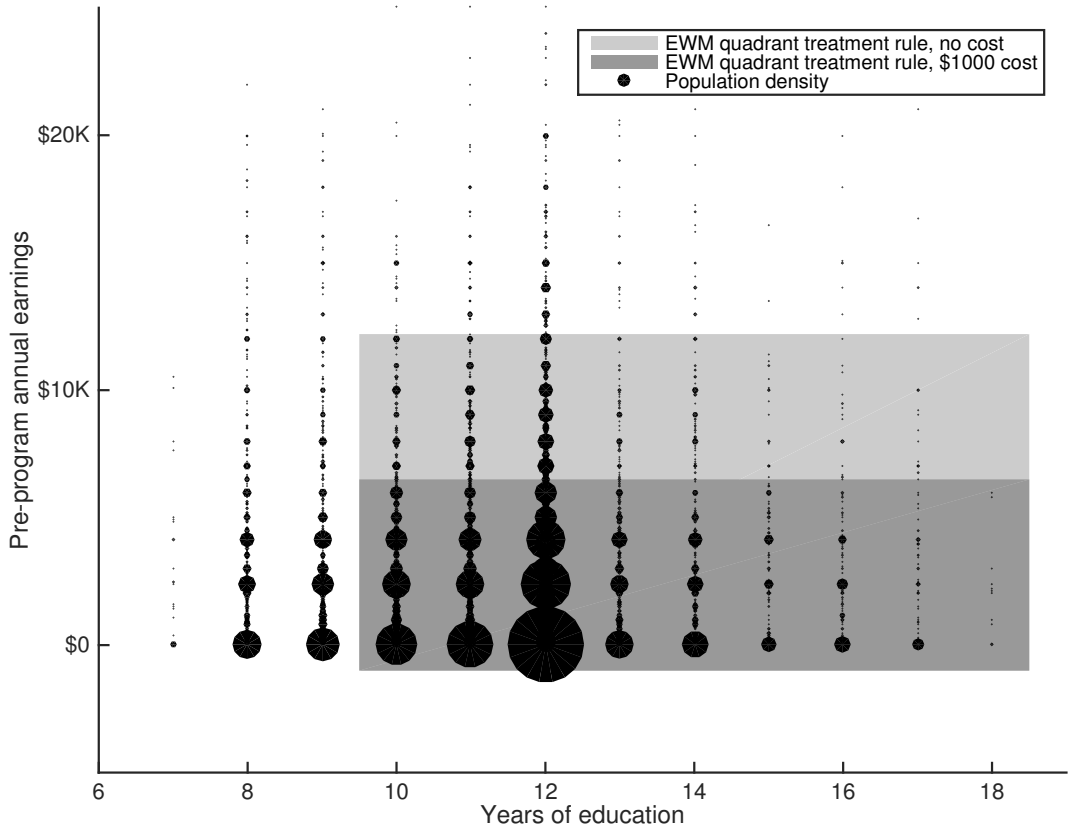

Figure 1: Empirical Welfare-Maximizing treatment rules from the quadrant class conditioning on years of education and pre-program earnings

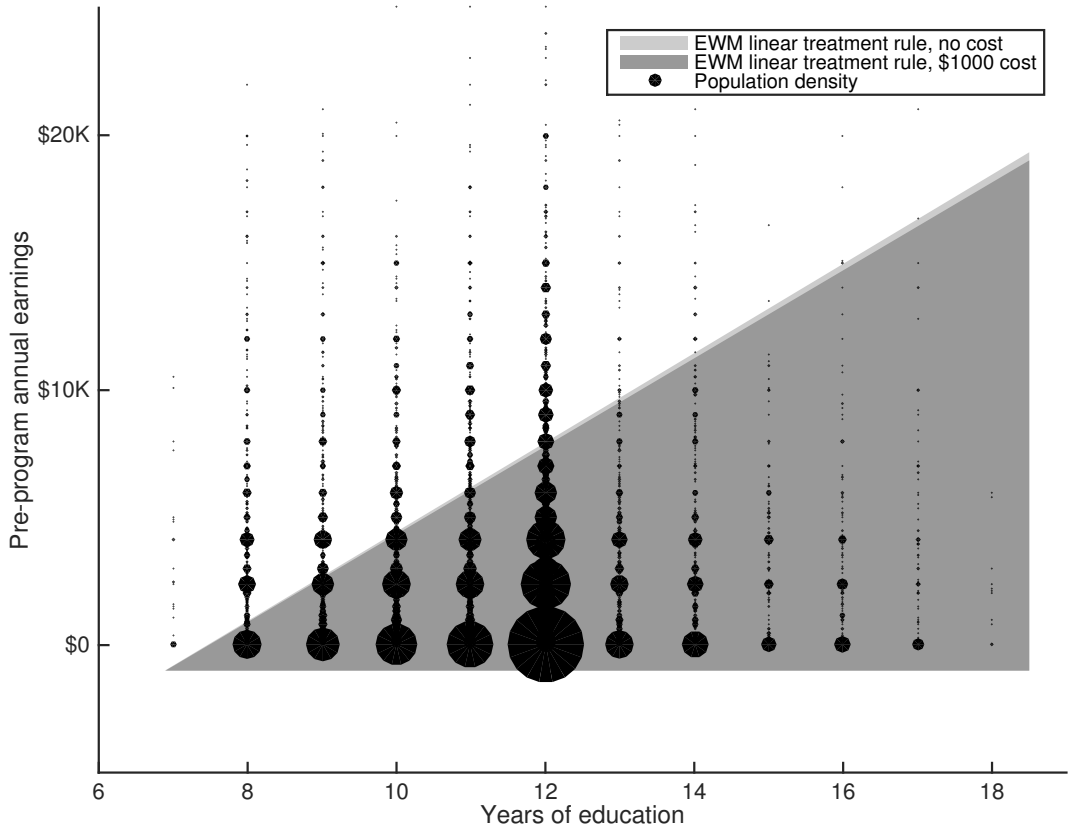

Figure 2: Empirical Welfare-Maximizing treatment rules from the linear class conditioning on years of education and pre-program earnings

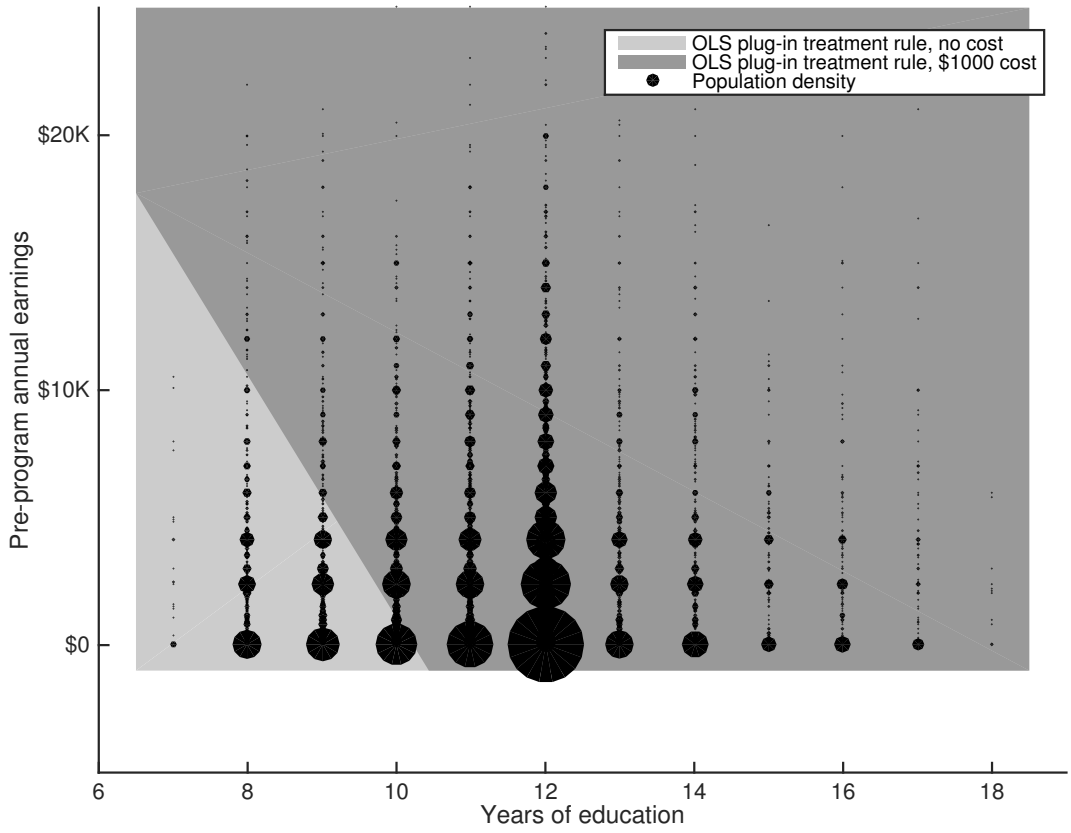

Figure 3: Parametric plug-in treatment rules based on the linear regressions of treatment outcomes on years of education and pre-program earnings

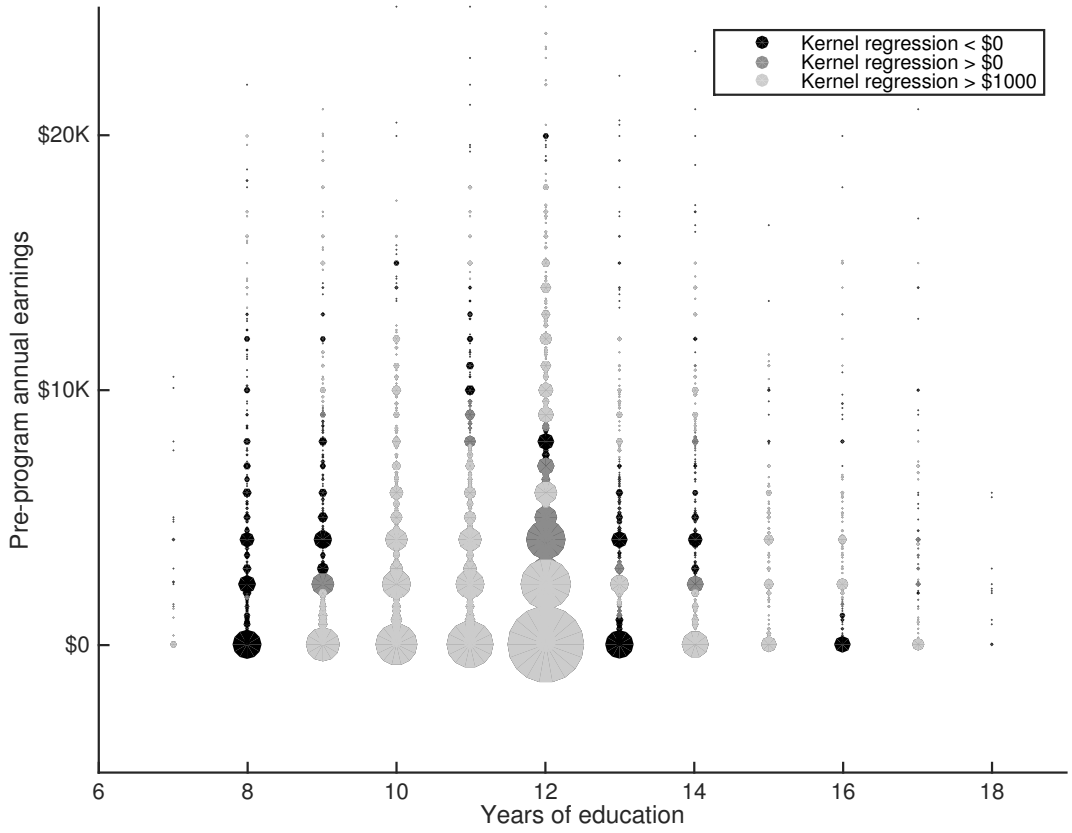

Figure 4: Nonparametric plug-in treatment rules based on the kernel regressions of treatment outcomes on years of education and pre-program earnings

A Appendix: Lemmas and Proofs

A.1 Notations and Basic Lemmas

Let $Z_i = (Y_i, D_i, X_i) \in \mathcal{Z}$. The *subgraph* of a real-valued function $f : \mathcal{Z} \mapsto \mathbb{R}$ is the set

$$SG(f) \equiv \{(z,t) \in \mathcal{Z} \times \mathbb{R} : 0 \leq t \leq f(z) \text{ or } f(z) \leq t \leq 0\}.$$

The following lemma establishes a link between the VC-dimension of a class of subsets in the covariate space $\mathcal{X}$ and the VC-dimension of a class of subgraphs of functions on $\mathcal{Z} = \mathbb{R} \times \{0,1\} \times \mathcal{X}$ (their subgraphs will be in $\mathcal{Z} \times \mathbb{R}$).

Lemma A.1. *Let $\mathcal{G}$ be a VC-class of subsets of $\mathcal{X}$ with VC-dimension $v < \infty$. Let g and h be two given functions from $\mathcal{Z}$ to $\mathbb{R}$. Then the set of functions from $\mathcal{Z}$ to $\mathbb{R}$*

$$\mathcal{F} = \{f_G(z) = g(z) \cdot 1\{x \in G\} + h(z) 1\{x \notin G\} : G \in \mathcal{G}\}$$

is a VC-subgraph class of functions with VC-dimension less than or equal to v.

Proof. Let $z_i = (y_i, d_i, x_i)$. By the assumption, no set of $(v+1)$ points in $\mathcal{X}$ could be shattered by $\mathcal{G}$. Take an arbitrary set of $(v+1)$ points in $\mathcal{Z} \times \mathbb{R}$, $A = \{(z_1, t_1), ..., (z_{v+1}, t_{v+1})\}$. Denote the collection of subgraphs of $\mathcal{F}$ by $SG(\mathcal{F}) \equiv \{\text{SG}(f_G), G \in \mathcal{G}\}$. We want to show that $SG(\mathcal{F})$ doesn't shatter A.

If for some $i \in \{1, \ldots, (v+1)\}$, $(z_i, t_i) \in \text{SG}(g) \cap \text{SG}(h)$ then $SG(\mathcal{F})$ cannot pick out all of the subsets of A because the i-th point is included in any $S \in SG(\mathcal{F})$. Similarly, if for some $i \in \{1, \ldots, (v+1)\}$, $(z_i, t_i) \in \text{SG}(g)^c \cap \text{SG}(h)^c$, then point i cannot be included in any $S \in SG(\mathcal{F})$.

The remaining case is that, for each i, either $(z_i, t_i) \in \text{SG}(g) \cap \text{SG}(h)^c$ or $(z_i, t_i) \in \text{SG}(g)^c \cap \text{SG}(h)$ holds. Indicate the former case by $\delta_i = 0$ and the latter case by $\delta_i = 1$. The points with $\delta_i = 0$ could be picked by $\text{SG}(f_G)$ if and only if $x_i \notin G$. The points with $\delta_i = 1$ could be picked if and only if $x_i \in G$. Given that $\mathcal{G}$ is a VC-class with VC-dimension v, there exists a subset X_0 of $\{x_1, \ldots, x_{v+1}\}$ such that $X_0 \neq (\{x_1, \ldots, x_{v+1}\} \cap G)$ for any $G \in \mathcal{G}$. Then there could be no set $S \in SG(\mathcal{F})$ that picks out the set (possibly empty)

$$\{(z_i, t_i) : (x_i \in X_0 \text{ and } \delta_i = 1) \text{ or } (x_i \notin X_0 \text{ and } \delta_i = 0)\}, \tag{A.1}$$

because this set of points could only be picked out by $\text{SG}(f_G)$ if $(\{x_1, \ldots, x_{v+1}\} \cap G) = X_0$. Hence, $\mathcal{F}$ is a VC subgraph class of functions with VC-dimension less than or equal to v. □

In addition to the notations introduced in the main text, the following notations are used throughout the appendix. The empirical probability distribution based on an iid size n sample of $Z_i = (Y_i, D_i, X_i)$ is denoted by P_n. $L_2(P)$ metric for f is denoted by $\|f\|_{L_2(P)} = \left[\int_{\mathcal{Z}} f^2 dP\right]^{1/2}$, and the sup-metric of f is denoted by $\|f\|_\infty$. Positive constants that only depend on the class of data generating processes, not on the sample size nor the VC-dimension, are denoted by $c_1, c_2, c_3, \ldots$. The universal constants are denoted by the capital letter $C_1, C_2, \ldots$.

In what follows, we present lemmas that will be used in the proofs of Theorems 2.1 and 2.3. Lemmas A.2 and A.3 are classical inequalities whose proofs can be found, for instance, in Lugosi (2002).

Lemma A.2. *Hoeffding's Lemma: let X be a random variable with $EX = 0$, $a \leq X \leq b$. Then, for $s > 0$,*

$$E\left(e^{sX}\right) \leq e^{s^2(b-a)^2/8}.$$

Lemma A.3. *Let $\lambda > 0$, $n \geq 2$, and let $Y_1, \ldots, Y_n$ be real-valued random variables such that for all $s > 0$ and $1 \leq i \leq n$, $E(e^{sY_i}) \leq e^{s^2\lambda^2/2}$ holds. Then,*

$$\begin{aligned} &(i) \quad E\left(\max_{i\leq n} Y_i\right) \leq \lambda\sqrt{2\ln n}, \\ &(ii) \quad E(\max_{i\leq n} |Y_i|) \leq \lambda\sqrt{2\ln(2n)}. \end{aligned}$$

The next two lemmas give maximal inequalities that bound the mean of a supremum of centered empirical processes indexed by a VC-subgraph class of functions. The first maximal inequality (Lemma A.4) is standard in the empirical process literature, and it yields our Theorem 2.1 as a corollary. Though its proof can be found elsewhere (e.g., Dudley (1999), van der Vaart and Wellner (1996)), we present it here for the sake of completeness and for later reference in the proof of Lemma A.5. The second maximal inequality (Lemma A.5) concerns the class of functions whose diameter is constrained by the $L_2(P)$-norm. Lemma A.5 will be used in the proofs of Theorem 2.3. A lemma similar to our Lemma A.5 appears in Massart and Nédélec (2006, Lemma A.3).

Lemma A.4. *Let $\mathcal{F}$ be a class of uniformly bounded functions, i.e., there exists $\bar{F} < \infty$ such that $\|f\|_\infty \leq \bar{F}$ for all $f \in \mathcal{F}$. Assume that $\mathcal{F}$ is a VC-subgraph class with VC-dimension $v < \infty$. Then, there is a universal constant C_1 such that*

$$E_{P^n}\left[\sup_{f\in\mathcal{F}} |E_n(f) - E_P(f)|\right] \leq C_1 \bar{F}\sqrt{\frac{v}{n}}$$

holds for all $n \geq 1$.

Proof. Introduce $(Z'_1, \ldots, Z'_n)$, an independent copy of $(Z_1, \ldots, Z_n) \sim P^n$. We denote the probability law of $(Z'_1, \ldots, Z'_n)$ by $P^{n'}$, its expectation by $E_{P^{n'}}(\cdot)$, and the sample average with respect to $(Z'_1, \ldots, Z'_n)$ by $E'_n(\cdot)$. Define iid Rademacher variables $\sigma_1, \ldots, \sigma_n$ such that $\Pr(\sigma_1 = -1) = \Pr(\sigma_1 = 1) = 1/2$ and they are independent of $Z_1, Z'_1, \ldots, Z_n, Z'_n$. Then,

$$\begin{aligned}
E_{P^n}\left[\sup_{f\in\mathcal{F}}|E_n(f) - E_P(f)|\right] &= E_{P^n}\left[\sup_{f\in\mathcal{F}}\left|E_{P^{n'}}\left[E_n(f) - E'_n(f)|Z_1, \ldots, Z_n\right]\right|\right] \\
&\leq E_{P^n}\left[\sup_{f\in\mathcal{F}}E_{P^{n'}}\left[\left|E_n(f) - E'_n(f)\right||Z_1, \ldots, Z_n\right]\right] \\
&(\ \because \text{ Jensen's inequality}) \\
&\leq E_{P^n, P^{n'}}\left[\sup_{f\in\mathcal{F}}\left|E_n(f) - E'_n(f)\right|\right] \\
&= \frac{1}{n}E_{P^n, P^{n'}}\left\{\sup_{f\in\mathcal{F}}\left|\sum_{i=1}^n\left(f(Z_i) - f(Z'_i)\right)\right|\right\} \\
&= \frac{1}{n}E_{P^n, P^{n'}, \sigma}\left\{\sup_{f\in\mathcal{F}}\left|\sum_{i=1}^n\sigma_i\left(f(Z_i) - f(Z'_i)\right)\right|\right\} \\
&(\ \because \ f(Z_i) - f(Z'_i) \sim \sigma_i\left(f(Z_i) - f(Z'_i)\right) \text{ for all } i\ ) \\
&\leq \frac{1}{n}\ E_{P^n, P^{n'}, \sigma}\left\{\sup_{f\in\mathcal{F}}\left|\sum_{i=1}^n\sigma_i f(Z_i)\right| + \sup_{f\in\mathcal{F}}\left|\sum_{i=1}^n\sigma_i f(Z'_i)\right|\right\} \\
&= \frac{2}{n}E_{P^n, \sigma}\left[\sup_{f\in\mathcal{F}}\left|\sum_{i=1}^n\sigma_i f(Z_i)\right|\right] \\
&= \frac{2}{n}E_{P^n}\left\{E_\sigma\left[\sup_{f\in\mathcal{F}}\left|\sum_{i=1}^n\sigma_i f(Z_i)\right||Z_1, \ldots, Z_n\right]\right\}. \qquad \text{(A.2)}
\end{aligned}$$

Fix $Z_1, \ldots, Z_n$, and define $\mathbf{f} \equiv (f(Z_1), \ldots, f(Z_n)) = (f_1, \ldots, f_n)$, which is a vector of length n corresponding to the value of $f \in \mathcal{F}$ evaluated at each of $(Z_1, \ldots, Z_n)$. Let $\mathbf{F} = \{\mathbf{f} : f \in \mathcal{F}\} \subset \mathbb{R}^n$, which is a bounded set in $\mathbb{R}^n$ with radius $\bar{F} = M/\kappa$, since $\mathcal{F}$ is the set of uniformly bounded functions with $|f(\cdot)| \leq M/\kappa$. Introduce the Euclidean norm to $\mathbf{F}$,

$$\rho(\mathbf{f}, \mathbf{f}') = \left(\frac{1}{n}\sum_{i=1}^n\left(f_i - f'_i\right)^2\right)^{1/2}.$$

Let $\mathbf{f}^{(0)} = (0, \ldots, 0)$, and $\mathbf{f}^* = (f_1^*, \ldots, f_n^*)$ be a random element in $\mathbf{F}$ maximizing $|\sum_{i=1}^n \sigma_i f_i|$. Let $B_0 = \left\{\mathbf{f}^{(0)}\right\}$ and construct $\left\{B_k : k = 1, \ldots, \bar{K}\right\}$ a sequence of covers of $\mathbf{F}$, such that $B_k \subset \mathbf{F}$ is a minimal cover with radius $2^{-k}\bar{F}$ and $B_{\bar{K}} = \mathbf{F}$. Note that such $\bar{K} < \infty$ exists at given n

and $(Z_1, \ldots, Z_n)$. Define also $\left\{\mathbf{f}^{(k)} \in B_k : k = 1, \ldots, \bar{K}\right\}$ be a random sequence such that $\mathbf{f}^{(k)} \in \arg\min_{\mathbf{f}\in\mathbf{B}_k} \rho(\mathbf{f}, \mathbf{f}^*)$. Since B_k is a cover with radius $2^{-k}\bar{F}$, $\rho\left(\mathbf{f}^{(k)}, \mathbf{f}^*\right) \leq 2^{-k}\bar{F}$ holds. In addition, we have

$$\rho\left(\mathbf{f}^{(k-1)}, \mathbf{f}^{(k)}\right) \leq \rho\left(\mathbf{f}^{(k)}, \mathbf{f}^*\right) + \rho\left(\mathbf{f}^{(k-1)}, \mathbf{f}^*\right) \leq 3 \cdot 2^{-k}\bar{F}.$$

By a telescope sum,

$$\begin{aligned}\sum_{i=1}^{n} \sigma_i f_i^* &= \sum_{i=1}^{n} \sigma_i f_i^{(0)} + \sum_{k=1}^{\bar{K}} \sum_{i=1}^{n} \sigma_i \left(f_i^{(k)} - f_i^{(k-1)}\right) \\ &= \sum_{k=1}^{\bar{K}} \sum_{i=1}^{n} \sigma_i \left(f_i^{(k)} - f_i^{(k-1)}\right).\end{aligned}$$

We hence obtain

$$\begin{aligned}E_\sigma \left|\sum_{i=1}^{n} \sigma_i f_i^*\right| &\leq \sum_{k=1}^{\bar{K}} E_\sigma \left|\sum_{i=1}^{n} \sigma_i \left(f_i^{(k)} - f_i^{(k-1)}\right)\right| \\ &\leq \sum_{k=1}^{\bar{K}} E_\sigma \max_{\mathbf{f}\in B_k, \mathbf{g}\in B_{k-1}:\rho(\mathbf{f},\mathbf{g})\leq 3\cdot 2^{-k}\bar{F}} \left|\sum_{i=1}^{n} \sigma_i (f_i - g_i)\right|. \qquad \text{(A.3)}\end{aligned}$$

We apply Lemma A.2 to obtain

$$\begin{aligned}E_\sigma \left(e^{s\sum_{i=1}^{n} \sigma_i (f_i - g_i)}\right) &= \prod_{i=1}^{n} E_{\sigma_i}\left[e^{s\sigma_i (f_i - g_i)}\right] \\ &\leq \prod_{i=1}^{n} e^{s^2 (f_i - g_i)^2/2} \\ &= \exp\left(s^2 n \rho^2(\mathbf{f}, \mathbf{g})/2\right) \\ &\leq \exp\left(s^2 n \left(3 \cdot 2^{-k}\bar{F}\right)^2/2\right).\end{aligned}$$

An application of Lemma A.3 (ii) with $\lambda = 3\sqrt{n} \cdot 2^{-k}\bar{F}$ and $n = |B_k|\,|B_{k-1}| \leq |B_k|^2$ then yields

$$\begin{aligned}E_\sigma \max_{\mathbf{f}\in B_k, \mathbf{g}\in B_{k-1}:\rho(\mathbf{f},\mathbf{g})\leq 3\cdot 2^{-k}\bar{F}} \left|\sum_{i=1}^{n} \sigma_i (f_i - g_i)\right| &\leq 3\sqrt{n} \cdot 2^{-k}\bar{F}\sqrt{2\ln 2\,|B_k|^2} \\ &= 3\sqrt{n} \cdot 2^{-k}\bar{F}\sqrt{2\ln 2 N(2^{-k}\bar{F}, \mathbf{F}, \rho)^2} \\ &= 6\sqrt{n} \cdot 2^{-k}\bar{F}\sqrt{\ln 2^{1/2} N(2^{-k}\bar{F}, \mathbf{F}, \rho)},\end{aligned}$$

where $N(r, \mathbf{F}, \rho)$ is the covering number of $\mathbf{F}$ with radius r in terms of norm ρ. Accordingly,

$$\begin{aligned} E_\sigma \left| \sum_{i=1}^n \sigma_i f_i^* \right| &\leq \sum_{k=1}^K 6\sqrt{n} \cdot 2^{-k} \bar{F} \sqrt{\ln 2^{1/2} N(2^{-k}\bar{F}, \mathbf{F}, \rho)} \\ &\leq 12\sqrt{n} \sum_{k=1}^\infty 2^{-(k+1)} \bar{F} \sqrt{\ln 2^{1/2} N(2^{-k}\bar{F}, \mathbf{F}, \rho)} \\ &\leq 12\sqrt{n} \int_0^1 \epsilon \bar{F} \sqrt{\ln 2^{1/2} N(\epsilon \bar{F}, \mathbf{F}, \rho)} d\epsilon, \end{aligned} \tag{A.4}$$

where the last line follows from the fact that $N(\epsilon \bar{F}, \mathbf{F}, \rho)$ is decreasing in ϵ.

To bound (A.4) from above, we apply a uniform entropy bound for the covering number. In Theorem 2.6.7 of van der Vaart and Wellner (1996), by setting $r = 2$ and Q at the empirical probability measure of $(Z_1, \ldots, Z_n)$, we have,

$$N(\epsilon \bar{F}, \mathbf{F}, \rho) \leq K(v+1) (16e)^{(v+1)} \left(\frac{1}{\epsilon} \right)^{2v}, \tag{A.5}$$

where $K > 0$ is a universal constant. Plugging this into (A.4) leads to

$$\begin{aligned} E_\sigma \left| \sum_{i=1}^n \sigma_i f_i^* \right| &\leq 12\bar{F}\sqrt{n} \int_0^1 \epsilon \sqrt{\ln(2^{1/2}K) + \ln(v+1) + (v+1)\ln(16e) - 2v \ln \epsilon} d\epsilon \\ &\leq 12\bar{F}\sqrt{nv} \int_0^1 \epsilon \sqrt{\ln(2^{1/2}K) + \ln 2 + 2\ln(16e) - 2\ln \epsilon} d\epsilon \\ &= C' \bar{F} \sqrt{nv}, \end{aligned} \tag{A.6}$$

where $C' = 12 \int_0^1 \epsilon \sqrt{\ln(2^{1/2}K) + \ln 2 + 2\ln(16e) - 2\ln \epsilon} d\epsilon < \infty$. Combining (A.6) with (A.2) and setting $C_1 = 2C'$ leads to the conclusion. □

Lemma A.5. *Let $\mathcal{F}$ be a class of uniformly bounded functions with $\|f\|_\infty \leq \bar{F} < \infty$ for all $f \in \mathcal{F}$. Assume that $\mathcal{F}$ is a VC-subgraph class with VC-dimension $v < \infty$. Assume further that $\sup_{f \in \mathcal{F}} \|f\|_{L_2(P)} \leq \delta$. Then, there exists a positive universal constant C_2 such that*

$$E_{P^n} \left[\sup_{f \in \mathcal{F}} \left(E_n(f) - E_P(f) \right) \right] \leq C_2 \delta \bar{F} \sqrt{\frac{v}{n}}$$

holds for all $n \geq C_1 \bar{F}^2 v / \delta^2$, where C_1 is the universal constant defined in Lemma A.4.

Proof. By the same symmetrization argument and the same use of Rademacher variables as in the proof of Lemma A.4, we have

$$E_{P^n} \left[\sup_{f \in \mathcal{F}} \left(E_n(f) - E_P(f) \right) \right] \leq \frac{2}{n} E_{P^n} \left\{ E_\sigma \left[\sup_{f \in \mathcal{F}} \sum_{i=1}^n \sigma_i f(Z_i) | Z_1, \ldots, Z_n \right] \right\}. \tag{A.7}$$

Fix the values of $Z_1, \ldots, Z_n$, and define $\mathbf{f}$, $\mathbf{f}^{(0)}$, $\mathbf{F}$, and norm $\rho(\mathbf{f}, \mathbf{f}')$ as in the proof of Lemma A.4. Let $\mathbf{f}^*$ be a maximizer of $\sum_{i=1}^n \sigma_i f(Z_i)$ in $\mathbf{F}$ and let $\delta_n = \sup_{\mathbf{f}\in\mathbf{F}} \rho(\mathbf{f}^{(0)}, \mathbf{f}) \leq \bar{F}$. Let $B_0 = \left\{\mathbf{f}^{(0)}\right\}$ and construct $\left\{B_k : k = 1, \ldots, \bar{K}\right\}$ a sequence of covers of $\mathbf{F}$, such that $B_k \subset \mathbf{F}$ is a minimal cover with radius $2^{-k}\delta_n$ and $B_{\bar{K}} = \mathbf{F}$. We define $\left\{\mathbf{f}^{(k)} \in B_k : k = 1, \ldots, \bar{K}\right\}$ to be a random sequence such that $\mathbf{f}^{(k)} \in \arg\min_{\mathbf{f}\in\mathbf{B}_k} \rho\left(\mathbf{f}, \mathbf{f}^*\right)$. By applying the chaining argument in the proof of Lemma A.4, Lemma A.3 (i), and the uniform bound of the covering number (A.5), we obtain

$$\begin{aligned} E_\sigma \sum_{i=1}^n \sigma_i f_i^* &\leq 12\sqrt{n} \int_0^1 \epsilon\delta_n \sqrt{\log N(\epsilon\delta_n, \mathbf{F}, \rho)} d\epsilon, \\ &\leq C'\delta_n\sqrt{nv}. \end{aligned}$$

for the universal constant C' defined in the proof of Lemma A.4. Hence, from (A.7), we have

$$\begin{aligned} E_{P^n}\left[\sup_{f\in\mathcal{F}} \left(E_n\left(f\right) - E_P(f)\right)\right] &\leq C_1\sqrt{\frac{v}{n}} E_{P^n}\left(\delta_n\right) \\ &= C_1\sqrt{\frac{v}{n}} E_{P^n}\left(\left[\sup_{f\in\mathcal{F}} E_n\left(f^2\right)\right]^{1/2}\right) \\ &\leq C_1\sqrt{\frac{v}{n}} \left[E_{P^n}\left(\sup_{f\in\mathcal{F}} E_n\left(f^2\right)\right)\right]^{1/2}. \end{aligned} \tag{A.8}$$

Note that $E_n\left(f^2\right)$ is bounded by

$$\begin{aligned} E_n\left(f^2\right) &= E_n\left(f^2 - E_P\left(f^2\right)\right) + E_P(f^2) \\ &= E_n\left[\left(f - \|f\|_{L_2(P)}\right)\left(f + \|f\|_{L_2(P)}\right)\right] + \|f\|_{L_2(P)}^2 \\ &\leq 2\bar{F}E_n\left[f - \|f\|_{L_2(P)}\right] + \|f\|_{L_2(P)}^2 \\ &\because \|f\|_{L_2(P)} \geq E_P(f) \text{ by the Cauchy-Schwartz inequality} \\ &\leq 2\bar{F}E_n\left[f - E_P(f)\right] + \|f\|_{L_2(P)}^2. \end{aligned}$$

Combining this inequality with (A.8) yields

$$E_{P^n}\left[\sup_{f\in\mathcal{F}} \left(E_n\left(f\right) - E_P(f)\right)\right] \leq C_1\sqrt{\frac{v}{n}} \sqrt{2\bar{F}E_{P^n}\left[\sup_{f\in\mathcal{F}} \left(E_n\left(f\right) - E_P(f)\right)\right] + \delta^2}.$$

Solving this inequality for $E_{P^n}\left[\sup_{f\in\mathcal{F}} \left(E_n\left(f\right) - E_P(f)\right)\right]$ leads to

$$E_{P^n}\left[\sup_{f\in\mathcal{F}} \left(E_n\left(f\right) - E_P(f)\right)\right] \leq \bar{F}C_1\sqrt{\frac{v}{n}} \left(\sqrt{\frac{v}{n}} + \sqrt{\frac{v}{n} + \frac{\delta^2}{\bar{F}^2 C_1}}\right).$$

For $\frac{v}{n} \leq \frac{\delta^2}{\bar{F}^2 C_1}$, that is, $n \geq \frac{C_1 \bar{F}^2 v}{\delta^2}$, the upper bound can be further bounded by $(1+\sqrt{2})\sqrt{C_1}\bar{F}\delta\sqrt{\frac{v}{n}}$, so the conclusion of the lemma follows with $C_2 = (1+\sqrt{2})\sqrt{C_1}$. □

The next lemma is a version of Bousquet's concentration inequality (Bousquet, 2002).

Lemma A.6. *Let $\mathcal{F}$ be a countable family of measurable functions, such that $\sup_{f\in\mathcal{F}} E_P(f^2) \leq \delta^2$ and $\sup_{f\in\mathcal{F}} \|f\|_\infty \leq \bar{F}$ for some constants δ^2 and $\bar{F}$. Let $S = \sup_{f\in\mathcal{F}} (E_n(f) - E_P(f))$. Then, for every positive t,*

$$P^n\left(S - E_{P^n}(S) \geq \sqrt{\frac{2\left[\delta^2 + 4\bar{F}E_{P^n}(S)\right]t}{n}} + \frac{2\bar{F}t}{3n}\right) \leq \exp(-t).$$

A.2 Proofs of Theorems 2.1 and 2.2

Proof of Theorem 2.1. Define

$$f(Z_i; G) = \left[\frac{Y_i D_i}{e(X_i)} \cdot 1\{X_i \in G\} + \frac{Y_i(1-D_i)}{1-e(X_i)} \cdot 1\{X_i \notin G\}\right],$$

and the class of functions on $\mathcal{Z}$

$$\mathcal{F} = \{f(\cdot; G) : G \in \mathcal{G}\}.$$

With these notations, we can express inequality (2.2) as

$$W_{\mathcal{G}}^* - W(\hat{G}_{EWM}) \leq 2\sup_{f\in\mathcal{F}} |E_n(f) - E_P(f)|. \tag{A.9}$$

Note that Assumption 2.1 (BO) and (SO) imply that $\mathcal{F}$ has uniform envelope $\bar{F} = M/\kappa$. Also, by Assumption 2.1 (VC) and Lemma A.1, $\mathcal{F}$ is a VC-subgraph class of functions with VC-dimension at most v. We apply Lemma A.4 to (A.9) to obtain

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

Since this upper bound does not depend on $P \in \mathcal{P}(M,\kappa)$, the upper bound is uniform over $\mathcal{P}(M,\kappa)$. □

Proof of Theorem 2.2. In obtaining the rate lower bound, we normalize the support of outcomes to $Y_{1i}, Y_{0i} \in \left[-\frac{1}{2}, \frac{1}{2}\right]$. That is, we focus on bounding $\sup_{P \in \mathcal{P}(1,\kappa)} E_{P^n}\left[W^*_{\mathcal{G}} - W(G_n)\right]$. The lower bound of the original welfare loss $\sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n}\left[W^*_{\mathcal{G}} - W(G_n)\right]$ is obtained by multiplying by M the lower bound of $\sup_{P \in \mathcal{P}(1,\kappa)} E_{P^n}\left[W^*_{\mathcal{G}} - W(G_n)\right]$.

We consider a suitable subclass $\mathcal{P}^* \subset \mathcal{P}(1,\kappa)$, for which the worst case welfare loss can be bounded from below by a distribution-free term that converges at rate $n^{-1/2}$. The construction of $\mathcal{P}^*$ proceeds as follows. First, let $x_1, \ldots, x_v \in \mathcal{X}$ be v points that are shattered by $\mathcal{G}$. We constrain the marginal distribution of X to being supported only on $(x_1, \ldots, x_v)$. We put mass p at x_i, $i < v$, and mass $1-(v-1)p$ 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_1, Y_0, X), and D follows the Bernoulli distribution with $\Pr(D = 1) = 1/2$. Let $\mathbf{b} = (b_1, \ldots, b_{v-1}) \in \{0,1\}^{v-1}$ be a bit vector used to index a member of $\mathcal{P}^*$, i.e., $\mathcal{P}^*$ consists of finite number of DGPs. For each $j = 1, \ldots, (v-1)$, and depending on $\mathbf{b}$, construct the following conditional distribution of Y_1 given $X = x_j$; if $b_j = 1$,

$$Y_1 = \begin{cases} \frac{1}{2} & \text{with prob. } \frac{1}{2} + \gamma, \\ -\frac{1}{2} & \text{with prob. } \frac{1}{2} - \gamma, \end{cases} \tag{A.10}$$

and, if $b_j = 0$,

$$Y_1 = \begin{cases} \frac{1}{2} & \text{with prob. } \frac{1}{2} - \gamma, \\ -\frac{1}{2} & \text{with prob. } \frac{1}{2} + \gamma, \end{cases} \tag{A.11}$$

where $\gamma \in \left[0, \frac{1}{2}\right]$ is chosen properly in a later step of the proof. For $j = v$, the conditional distribution of Y_1 given $X = x_v$ is degenerate at $Y_1 = 0$. As for Y_0's conditional distribution, we consider the degenerate distribution at $Y_0 = 0$ at every $X = x_j$, $j = 1, \ldots, v$. That is, when $b_j = 1$, $\tau(x_j) = \gamma$, and when $b_j = 0$, $\tau(x_j) = -\gamma$. Each $\mathbf{b} \in \{0,1\}^{v-1}$ induces a unique joint distribution of $(Y_1, Y_0, D, X) \sim P_{\mathbf{b}}$ and, clearly, $P_{\mathbf{b}} \in \mathcal{P}(1,\kappa)$. We accordingly define $\mathcal{P}^* = \left\{P_{\mathbf{b}} : \mathbf{b} \in \{0,1\}^{v-1}\right\}$.

With knowledge of $P_{\mathbf{b}} \in \mathcal{P}^*$, the optimal treatment assignment rule is

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

which is feasible $G^*_{\mathbf{b}} \in \mathcal{G}$ by the construction of the support points of X. The maximized social welfare is

$$W(G^*_{\mathbf{b}}) = p\gamma \left(\sum_{j=1}^{v-1} b_j\right).$$

Let $\hat{G}$ be an arbitrary treatment choice rule as a function of $(Z_1, \ldots, Z_n)$, and $\hat{\mathbf{b}} \in \{0,1\}^{(v-1)}$ be a binary vector whose j-th element is $\hat{b}_j = 1\left\{x_j \in \hat{G}\right\}$. Consider $\pi(\mathbf{b})$ a prior distribution for $\mathbf{b}$ such that $b_1, \ldots, b_{v-1}$ are iid and $b_1 \sim Ber(1/2)$. The welfare loss satisfies the following inequalities,

$$\begin{aligned}
\sup_{P \in \mathcal{P}(1,\kappa)} E_{P^n}\left[W_{\mathcal{G}}^* - W(\hat{G})\right] &\geq \sup_{P_{\mathbf{b}} \in \mathcal{P}^*} E_{P_{\mathbf{b}}^n}\left[W(G_{\mathbf{b}}^*) - W(\hat{G})\right] \\
&\geq \int_{\mathbf{b}} E_{P_{\mathbf{b}}^n}\left[W(G_{\mathbf{b}}^*) - W(\hat{G})\right] d\pi(\mathbf{b}) \\
&= \gamma \int_{\mathbf{b}} E_{P_{\mathbf{b}}^n}\left[P_X\left(G_{\mathbf{b}}^* \triangle \hat{G}\right)\right] d\pi(\mathbf{b}) \\
&= \gamma \int_{\mathbf{b}} \int_{Z_1,\ldots,Z_n} P_X\left(\left\{b(X) \neq \hat{b}(X)\right\}\right) dP^n(Z_1, \ldots, Z_n|\mathbf{b})\, d\pi(\mathbf{b}) \\
&\geq \inf_{G_n} \gamma \int_{\mathbf{b}} \int_{Z_1,\ldots,Z_n} P_X\left(\left\{b(X) \neq \hat{b}(X)\right\}\right) dP^n(Z_1, \ldots, Z_n|\mathbf{b})\, d\pi(\mathbf{b})
\end{aligned}$$

where each $b(X)$ and $\hat{b}(X)$ is an element of $\mathbf{b}$ and $\hat{\mathbf{b}}$ such that $b(x_j) = b_j$ and $\hat{b}(x_j) = \hat{b}_j$. We define $\mathbf{b}(x_v) = \hat{\mathbf{b}}(x_v) = 0$. 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,

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

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

$$\begin{aligned}
&\gamma \int_{Z_1,\ldots,Z_n} E_X\left[\min\left\{\pi(b(X) = 1|Z_1, \ldots, Z_n), 1 - \pi(b(X) = 1|Z_1, \ldots, Z_n)\right\}\right] d\tilde{P}^n \\
= \;& \gamma \int_{Z_1,\ldots,Z_n} \sum_{j=1}^{v-1} p\left[\min\left\{\pi(b_j = 1|Z_1, \ldots, Z_n), 1 - \pi(b_j = 1|Z_1, \ldots, Z_n)\right\}\right] d\tilde{P}^n, \qquad \text{(A.12)}
\end{aligned}$$

where $\tilde{P}^n$ is the marginal likelihood of $\{(Y_{1,i}, Y_{0,i}, D_i, X_i) : i = 1, \ldots, n\}$ corresponding to prior $\pi(\mathbf{b})$. For each $j = 1, \ldots, (v-1)$, let

$$\begin{aligned}
K_j &= \#\{i : X_i = x_j\}, \\
k_j^+ &= \#\left\{i : X_i = x_j,\, Y_i D_i = \frac{1}{2}\right\}, \\
k_j^- &= \#\left\{i : X_i = x_j,\, Y_i D_i = -\frac{1}{2}\right\}.
\end{aligned}$$

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

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

Hence,

$$\begin{aligned} & \min\left\{\pi\left(b_j = 1|Z_1, \ldots, Z_n\right), 1 - \pi(b_j = 1|Z_1, \ldots, Z_n)\right\} \\ = & \frac{\min\left\{\left(\frac{1}{2}+\gamma\right)^{k_j^+}\left(\frac{1}{2}-\gamma\right)^{k_j^-}, \left(\frac{1}{2}+\gamma\right)^{k_j^-}\left(\frac{1}{2}-\gamma\right)^{k_j^+}\right\}}{\left(\frac{1}{2}+\gamma\right)^{k_j^+}\left(\frac{1}{2}-\gamma\right)^{k_j^-}+\left(\frac{1}{2}+\gamma\right)^{k_j^-}\left(\frac{1}{2}-\gamma\right)^{k_j^+}} \\ = & \frac{\min\left\{1, \left(\frac{\frac{1}{2}+\gamma}{\frac{1}{2}-\gamma}\right)^{k_j^+-k_j^-}\right\}}{1+\left(\frac{\frac{1}{2}+\gamma}{\frac{1}{2}-\gamma}\right)^{k_j^+-k_j^-}} \\ = & \frac{1}{1+a^{\left|k_j^+-k_j^-\right|}}, \quad \text{where } a = \frac{1+2\gamma}{1-2\gamma} > 1. \end{aligned} \tag{A.13}$$

Since $k_j^+ - k_j^- = \sum_{i:X_i=x_j} 2Y_iD_i$, plugging (A.13) into (A.12) yields

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

where $E_{\tilde{P}^n}(\cdot)$ is the expectation with respect to the marginal likelihood of $\{(Y_{1,i}, Y_{0,i}, D_i, X_i), i = 1, \ldots, n\}$. 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 $Y_{1,i}$ is $\Pr(Y_{1,i} = 1/2) = \Pr(Y_{1,i} = -1/2) = 1/2$, so

$$\begin{aligned} E_{\tilde{P}^n}\left|\sum_{i:X_i=x_j} 2Y_iD_i\right| &= E_{\tilde{P}_n}\left|\sum_{i=1:X_i=x_j, D_i=1} 2Y_{1,i}\right| \\ &= \sum_{k=0}^{n}\binom{n}{k}\left(\frac{p}{2}\right)^k\left(1-\frac{p}{2}\right)^{n-k} E\left|B(k, \frac{1}{2}) - \frac{k}{2}\right| \end{aligned}$$

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

$$\begin{aligned} E_{\tilde{P}^n}\left|\sum_{i:X_i=x_j} 2Y_iD_i\right| &\leq \sum_{k=0}^{n}\binom{n}{k}\left(\frac{p}{2}\right)^k\left(1-\frac{p}{2}\right)^{n-k}\sqrt{\frac{k}{4}} \\ &= E\sqrt{\frac{B\left(n,\frac{p}{2}\right)}{4}} \\ &\leq \sqrt{\frac{np}{8}}. \quad (\because \text{ Jensen's inequality}). \end{aligned}$$

Hence, the Bayes risk is bounded from below by

$$\begin{aligned} & \frac{\gamma}{2}p(v-1)a^{-\sqrt{\frac{np}{8}}} \\ \geq\ & \frac{\gamma}{2}p(v-1)e^{-(a-1)\sqrt{\frac{np}{8}}} \quad (\because 1+x\leq e^x\ \forall x) \\ =\ & \frac{p\gamma}{2}(v-1)e^{-\frac{4\gamma}{1-2\gamma}\sqrt{\frac{np}{8}}}. \end{aligned} \tag{A.14}$$

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 $\gamma = \sqrt{\frac{v-1}{n}}$. Since $(v-1)^{-1} \geq p \geq v^{-1}$, we have

$$\begin{aligned} \frac{p\gamma}{2}(v-1)e^{-\frac{4\gamma}{1-2\gamma}\sqrt{\frac{np}{8}}} &\geq \frac{1}{2}\sqrt{\frac{v-1}{n}}\left(1-\frac{1}{v}\right)\exp\left\{-\frac{\sqrt{2}}{1-2\gamma}\right\} \\ &\geq \frac{1}{4}\sqrt{\frac{v-1}{n}}\exp\left\{-2\sqrt{2}\right\}, \quad \text{if } 1-2\gamma\geq\frac{1}{2}. \end{aligned}$$

The condition $1-2\gamma \geq \frac{1}{2}$ is equivalent to $n \geq 16(v-1)$. This completes the proof. □

A.3 Proofs of Theorems 2.3 and 2.4

In proving Theorem 2.3, it is convenient to work with the *normalized welfare difference*,

$$d(G,G') = \frac{\kappa}{M}\left[W(G)-W(G')\right],$$

and its sample analogue

$$d_n(G, G') = \frac{\kappa}{M}\left[W_n(G) - W_n(G')\right]. \tag{A.15}$$

By Assumption 2.1 (BO) and (SO), both $d(G, G')$ and $d_n(G, G')$ are bounded in $[-1, 1]$, and the normalized welfare difference relates to the original welfare loss of decision set G as

$$d(G^*_{FB}, G) = \frac{\kappa}{M}\left[W(G^*_{FB}) - W(G)\right] \in [0, 1]. \tag{A.16}$$

Hence, the welfare loss upper bound of $\hat{G}_{EWM}$ can be obtained by multiplying M/κ by the upper bound of $d(G^*_{FB}, \hat{G}_{EWM})$.

Note that $d(G^*_{FB}, G)$ can be bounded from above by $P_X(G^*_{FB} \triangle G)$, since

$$\begin{aligned} d(G^*_{FB}, G) &= \frac{\kappa}{M}\int_{G^*_{FB}\triangle G} |\tau(X)|\, dP_X \\ &\leq \kappa P_X(G^*_{FB}\triangle G) \\ &\leq P_X(G^*_{FB}\triangle G). \end{aligned} \tag{A.17}$$

On the other hand, with Assumption 2.2 (MA) imposed, $P_X(G^*_{FB}\triangle G)$ can be bounded from above by a function of $d(G^*_{FB}, G)$, as the next lemma shows. We borrow this lemma from Tsybakov (2004).

Lemma A.7. *Suppose Assumption 2.2 (MA) holds with margin coefficient $\alpha \in (0, \infty)$. Then*

$$P_X(G^*_{FB}\triangle G) \leq c_1(M, \kappa, \eta, \alpha) d(G^*_{FB}, G)^{\frac{\alpha}{1+\alpha}}$$

holds for all $G \in \mathcal{G}$, where $c_1(M, \kappa, \eta, \alpha) = \left(\frac{M}{\kappa\eta\alpha}\right)^{\frac{\alpha}{1+\alpha}}(1+\alpha)$.

Proof. Let $A = \{x : |\tau(x)| > t\}$ and consider the following inequalities,

$$\begin{aligned} W(G^*_{FB}) - W(G) &= \int_{G^*_{FB}\triangle G} |\tau(x)|\, dP_X \\ &\geq \int_{(G^*_{FB}\triangle G)} |\tau(X)|\, 1\{x \in A\}\, dP_X \\ &\geq tP_X\left((G^*_{FB}\triangle G) \cap A\right) \\ &\geq t\left[P_X\left(G^*_{FB}\triangle G\right) - P_X(A^c)\right] \\ &\geq t\left[P_X\left(G^*_{FB}\triangle G\right) - \left(\frac{t}{\eta}\right)^{\alpha}\right], \end{aligned}$$

where the final line uses the margin condition. The right-hand side is maximized at $t = \eta(1+\alpha)^{-\frac{1}{\alpha}} \left[P_X\left(G_{FB}^* \triangle G\right)\right]^{\frac{1}{\alpha}} \leq \eta$, so it holds

$$W(G_{FB}^*) - W(G) \geq \eta\alpha \left(\frac{1}{1+\alpha}\right)^{\frac{1+\alpha}{\alpha}} \left[P_X\left(G_{FB}^* \triangle G\right)\right]^{\frac{1+\alpha}{\alpha}}.$$

This, in turn, implies

$$P_X\left(G_{FB}^* \triangle G\right) \leq \left(\frac{M}{\kappa\eta\alpha}\right)^{\frac{\alpha}{1+\alpha}} (1+\alpha) d(G_{FB}^*, G)^{\frac{\alpha}{1+\alpha}}.$$

□

Proof of Theorem 2.3. Let $a = \sqrt{kt}\epsilon_n$ with $k \geq 1$, $t \geq 1$, and $\epsilon_n > 0$, where $t \geq 1$ is arbitrary, k is a constant that we choose later, and ϵ_n is a sequence indexed by sample size n whose proper choice will be discussed in a later step. The normalized welfare loss can be bounded by

$$d(G_{FB}^*, \hat{G}_{EWM}) \leq d(G_{FB}^*, \hat{G}_{EWM}) - d_n\left(G_{FB}^*, \hat{G}_{EWM}\right),$$

as $d_n\left(G_{FB}^*, \hat{G}_{EWM}\right) \leq 0$ by Assumption 2.2 (FB). Define a class of functions induced by $G \in \mathcal{G}$.

$$\begin{aligned} \mathcal{H} &\equiv \{h(Z_i; G) : G \in \mathcal{G}\}, \\ h(Z_i; G) &\equiv \frac{\kappa}{M}\left(\frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)}\right)\left[1\{X_i \in G\} - 1\{X_i \in G_{FB}^*\}\right]. \end{aligned}$$

By Assumption 2.1 (VC) and Lemma A.1, $\mathcal{H}$ is a VC-subgraph-class with VC-dimension at most $v < \infty$ with envelope $\bar{H} = 1$. Using $h(Z_i; G)$, we can write $d(G_{FB}^*, G) = -E_P\left(h(Z_i; G)\right)$. Since $d(G_{FB}^*, G) \geq 0$ for all $G \in \mathcal{G}$, it holds $-E_P(h) \geq 0$ for all $h \in \mathcal{H}$.

Since we have

$$d(G_{FB}^*, \hat{G}_{EWM}) - d_n\left(G_{FB}^*, \hat{G}_{EWM}\right) = E_n\left(h(Z_i; \hat{G}_{EWM})\right) - E_P\left(h(Z_i; \hat{G}_{EWM})\right)$$

and $d_n\left(G_{FB}^*, \hat{G}_{EWM}\right) \leq 0$, the normalized welfare loss can be bounded by

$$\begin{aligned} d(G_{FB}^*, \hat{G}_{EWM}) &\leq E_n\left(h(Z_i; \hat{G}_{EWM})\right) - E_P\left(h(Z_i; \hat{G}_{EWM})\right) \\ &\leq V_a\left[d(G_{FB}^*, \hat{G}_{EWM}) + a^2\right], \end{aligned}$$

where

$$\begin{aligned} V_a &= \sup_{h \in \mathcal{H}}\left\{\frac{E_n(h) - E_P(h)}{-E_P(h) + a^2}\right\} \\ &= \sup_{h \in \mathcal{H}}\left\{E_n\left(\frac{h}{-E_P(h) + a^2}\right) - E_P\left(\frac{h}{-E_P(h) + a^2}\right)\right\} \end{aligned}$$

On event $V_a < \frac{1}{2}$, $d(G^*_{FB}, \hat{G}_{EWM}) \leq a^2$ holds, so this implies

$$P^n\left(d(G^*_{FB}, \hat{G}_{EWM}) \geq a^2\right) \leq P^n\left(V_a \geq \frac{1}{2}\right). \tag{A.18}$$

In what follows, our aim is to construct an exponential inequality for $P^n\left(V_a \geq \frac{1}{2}\right)$ involving only t, and we make use of such exponential tail bound to bound $E_{P^n}\left(d(G^*_{FB}, \hat{G}_{EWM})\right)$.

To apply the Bousquet's inequality (Lemma A.6) to V_a, note first that,

$$\begin{aligned}
E_P\left(\left(\frac{h}{-E_P(h)+a^2}\right)^2\right) &\leq \frac{P_X(G^*_{FB} \triangle G)}{(E_P(h)+a^2)^2} \\
&\leq c_1 \frac{[-E_P(h)]^{\frac{\alpha}{1+\alpha}}}{(-E_P(h)+a^2)^2} \\
&\because \text{ by Lemma A.7 and } d(G^*_{FB}, G) = -E_P\left(h(Z_i;G)\right)) \\
&\leq c_1 \sup_{\epsilon \geq 0} \frac{\epsilon^{\frac{2\alpha}{1+\alpha}}}{\left(\epsilon^2+a^2\right)^2} \\
&\leq c_1 \frac{1}{a^2} \sup_{\epsilon \geq 0} \frac{\epsilon^{\frac{2\alpha}{1+\alpha}}}{\epsilon^2+a^2} \\
&\leq c_1 \frac{1}{a^2} \sup_{\epsilon \geq 0} \left(\frac{\epsilon^{\frac{\alpha}{1+\alpha}}}{\epsilon \vee a}\right)^2 \\
&\leq c_1 \frac{1}{a^4} a^{\frac{2\alpha}{1+\alpha}},
\end{aligned}$$

where c_1 is a constant that depends only on $(M, \kappa, \eta, \alpha)$ as defined in Lemma A.7. We, on the other hand, have

$$\sup_{h \in \mathcal{H}} \left| \sup_Z \frac{h}{-E_P(h)+a^2} \right| \leq \frac{1}{a^2}.$$

Hence, Lemma A.6 gives, with probability larger than $1 - \exp(-t)$,

$$V_a \leq E_{P^n}(V_a) + \sqrt{\frac{\left[c_1 a^{\frac{2\alpha}{1+\alpha}-2} + 4E_{P^n}(V_a)\right] t}{na^2}} + \frac{2t}{3na^2}. \tag{A.19}$$

Next, we derive an upper bound of $E_{P^n}(V_a)$ by applying the maximal inequality of Lemma A.5. Let $r > 1$ be arbitrary and consider partitioning $\mathcal{H}$ by $\mathcal{H}_0, \mathcal{H}_1, \ldots$, where $\mathcal{H}_0 = \left\{h \in \mathcal{H} : -E_P(h) \leq a^2\right\}$

and $\mathcal{H}_j = \left\{h \in \mathcal{H} : r^{2(j-1)}a^2 < -E_P(h) \leq r^{2j}a^2\right\}$, $j = 1, 2, \ldots$. Then,

$$\begin{aligned} V_a &\leq \sup_{h\in\mathcal{H}_0}\left\{\frac{E_n(h) - E_P(h)}{-E_P(h) + a^2}\right\} + \sum_{j\geq 1}\sup_{h\in\mathcal{H}_j}\left\{\frac{E_n(h) - E_P(h)}{-E_P(h) + a^2}\right\} \\ &\leq \frac{1}{a^2}\left[\sup_{h\in\mathcal{H}_0}\left(E_n(h) - E_P(h)\right) + \sum_{j\geq 1}(1 + r^{2(j-1)})^{-1}\sup_{h\in\mathcal{H}_j}\left(E_n(h) - E_P(h)\right)\right] \\ &\leq \frac{1}{a^2}\left[\begin{array}{c}\sup_{-E_P(h)\leq a^2}\left(E_n(h) - E_P(h)\right) \\ +\sum_{j\geq 1}(1 + r^{2(j-1)})^{-1}\sup_{-E_P(h)\leq r^{2j}a^2}\left(E_n(h) - E_P(h)\right)\end{array}\right]. \end{aligned} \tag{A.20}$$

Since it holds $\|h\|^2_{L_2(P)} \leq P_X(G^*_{FB}\triangle G) \leq c_1(M,\kappa,\eta,\alpha)\left[-E_P(h)\right]^{\frac{\alpha}{1+\alpha}}$, where the latter inequality follows from Lemma A.7, $-E_P(h) \leq r^{2j}a^2$ implies $\|h\|_{L_2(P)} \leq c_1^{1/2} r^{\frac{\alpha}{1+\alpha}j} a^{\frac{\alpha}{1+\alpha}}$. Hence, (A.20) can be further bounded by

$$V_a \leq \frac{1}{a^2}\left[\begin{array}{c}\sup_{\|h\|_{L_2(P)}\leq c_1^{1/2}a^{\frac{\alpha}{1+\alpha}}}\left(E_n(h) - E_P(h)\right) \\ +\sum_{j\geq 1}(1 + r^{2(j-1)})^{-1}\sup_{\|h\|_{L_2(P)}\leq c_1^{1/2}r^{\frac{\alpha}{1+\alpha}j}a^{\frac{\alpha}{1+\alpha}}}\left(E_n(h) - E_P(h)\right)\end{array}\right].$$

We apply Lemma A.5 to each supremum term, and obtain

$$\begin{aligned} E_{P^n}(V_a) &\leq C_2\frac{c_1^{\frac{1}{2}}}{a^2}\sqrt{\frac{v}{n}}a^{\frac{\alpha}{1+\alpha}}\sum_{j\geq 0}\frac{r^{\frac{\alpha}{1+\alpha}j}}{1 + r^{2(j-1)}} \\ &\leq C_2 c_1^{\frac{1}{2}}\sqrt{\frac{v}{n}}a^{\frac{\alpha}{1+\alpha}-2}\left(\frac{r^2}{1 + r^{-\frac{2+\alpha}{1+\alpha}}}\right) \\ &\leq c_2\sqrt{\frac{v}{n}}a^{\frac{\alpha}{1+\alpha}-2} \end{aligned}$$

for

$$n \geq \frac{C_1 v}{c_1 a^{\frac{2\alpha}{1+\alpha}}} \quad \Longleftrightarrow \quad a \geq \left(\frac{C_1}{c_1}\right)^{\frac{1+\alpha}{2\alpha}}\left(\frac{v}{n}\right)^{\frac{1+\alpha}{2\alpha}} \tag{A.21}$$

where C_1 and C_2 are universal constants defined in Lemmas A.4 and A.5, and $c_2 = C_2 c_1^{\frac{1}{2}}\left(\frac{r^2}{1+r^{-\frac{2+\alpha}{1+\alpha}}}\right) \vee 1$ is a constant greater than or equal to one and depends only on (M,κ,η,α), as $r > 1$ is fixed. We plug in this upper bound into (A.19) to obtain

$$V_a \leq c_2\sqrt{\frac{v}{n}}a^{\frac{\alpha}{1+\alpha}-2} + \sqrt{\frac{\left[c_1 a^{\frac{2\alpha}{1+\alpha}-2} + 4c_2\sqrt{\frac{v}{n}}a^{\frac{\alpha}{1+\alpha}-2}\right]t}{na^2}} + \frac{2t}{3na^2}. \tag{A.22}$$

Choose ϵ_n as the root of $c_2\sqrt{\frac{v}{n}}a^{\frac{\alpha}{1+\alpha}-2}=1$, i.e.,

$$\epsilon_n = \left(c_2\sqrt{\frac{v}{n}}\right)^{\frac{1+\alpha}{2+\alpha}}. \tag{A.23}$$

Note that the right hand side of (A.22) is decreasing in a, and $a \geq \epsilon_n$ by the construction. Hence, if ϵ_n satisfies inequality (A.21), i.e.,

$$n \geq c_2^{-\alpha}\left(\frac{C_1}{c_1}\right)^{1+\frac{\alpha}{2}} v,$$

which can be reduced to an innocuous restriction $n \geq 1$ by inflating, if necessary, c_1 large enough, we can substitute ϵ_n for a to bound the right hand side of (A.22). In particular, by noting

$$\begin{aligned} c_2\sqrt{\frac{v}{n}}a^{\frac{\alpha}{1+\alpha}-2} &\leq \frac{\epsilon_n}{a} = \frac{1}{\sqrt{kt}} \leq \frac{1}{\sqrt{k}} \quad \text{and} \\ a^{\frac{2\alpha}{1+\alpha}-2} &= a^{2\left(\frac{\alpha}{1+\alpha}-2\right)}a^2 \leq \left[\epsilon_n^{\frac{\alpha}{1+\alpha}-2}\right]^2 \epsilon_n^2 = c_2^{-2}v^{-1}n\epsilon_n^2, \end{aligned}$$

the right-hand side of (A.22) can be bounded by

$$\begin{aligned} V_a &\leq \frac{1}{\sqrt{k}} + \sqrt{\frac{c_1c_2^{-2}v^{-1}n\epsilon_n^2+4}{nk\epsilon_n^2}} + \frac{2}{3nk\epsilon_n^2} \\ &= \frac{1}{\sqrt{k}} + \sqrt{\frac{c_1c_2^{-2}v^{-1}}{k} + \frac{4}{nk\epsilon_n^2}} + \frac{2}{3nk\epsilon_n^2} \\ &\leq \frac{1}{\sqrt{k}} + \sqrt{\frac{c_1c_2^{-2}v^{-1}}{k} + \frac{4}{k}} + \frac{2}{3k} \quad \text{for } n\epsilon_n^2 \geq 1. \end{aligned} \tag{A.24}$$

Note that condition $n\epsilon_n^2 \geq 1$ used to derive the last line is valid for all n, since it is equivalent to $n \geq c_2^{-2(1+\alpha)}v^{-(1+\alpha)}$, which holds for all $n \geq 1$ since $c_2 \geq 1$ and $v \geq 1$. By choosing k large enough so that the right-hand side of (A.24) is less than $\frac{1}{2}$, we can conclude

$$\Pr(V_a < \frac{1}{2}) \geq 1 - \exp(-t). \tag{A.25}$$

Hence, (A.18) yields

$$P^n\left(d(G^*_{FB}, \hat{G}_{EWM}) \geq kt\epsilon_n^2\right) \leq \exp(-t)$$

for all $t \geq 1$. From this exponential bound, we obtain

$$\begin{aligned}
& E_{P^n}\left(d(G^*_{FB}, \hat{G}_{EWM})\right) \\
= & \int_0^\infty P^n\left(d(G^*_{FB}, \hat{G}_{EWM}) > t'\right) dt' \\
\leq & \int_0^{k\epsilon_n^2} P^n\left(d(G^*_{FB}, \hat{G}_{EWM}) \geq t'\right) dt' + \int_{k\epsilon_n^2}^\infty P^n\left(d(G^*_{FB}, \hat{G}_{EWM}) \geq t'\right) dt' \\
\leq & k\epsilon_n^2 + k\epsilon_n^2 e^{-1} \\
= & (1+e^{-1}) k c_2^{\frac{2(1+\alpha)}{2+\alpha}} \left(\frac{v}{n}\right)^{\frac{1+\alpha}{2+\alpha}}.
\end{aligned}$$

So, setting $c = \frac{M}{\kappa}(1+e^{-1}) k c_2^{\frac{2(1+\alpha)}{2+\alpha}}$ leads to the conclusion. □

Proof of Theorem 2.4. As in the proof of Theorem 2.2, we work with the normalized outcome support, $Y_{1i}, Y_{0i} \in \left[-\frac{1}{2}, \frac{1}{2}\right]$. With the normalized outcome, constant η of the margin assumption satisfies $\eta \leq 1$.

Let $\alpha \in (0, \infty)$ and $\eta \in (0, 1]$ be given. Similarly to the proof of Theorem 2.2, we consider constructing a suitable subclass $\mathcal{P}^* \subset \mathcal{P}(1, \kappa, \eta, \alpha)$. Let $x_1, \ldots, x_v \in \mathcal{X}$ be v points that are shattered by $\mathcal{G}$, and let γ be a positive number satisfying $\gamma \leq \min\left\{\eta, \frac{1}{2}\right\}$, whose proper choice will be given later. We fix the marginal distribution of X at the one supported only on $(x_1, \ldots, x_v)$ and having the probability mass function,

$$\begin{aligned}
P_X(X_i &= x_j) = \frac{1}{v-1}\left(\frac{\gamma}{\eta}\right)^\alpha, \quad \text{for } j = 1, \ldots, (v-1), \quad \text{and} \\
P_X(X_i &= x_v) = 1 - \frac{1}{v-1}\left(\frac{\gamma}{\eta}\right)^\alpha.
\end{aligned}$$

Thus-constructed marginal distribution of X is common in $\mathcal{P}^*$. As in the proof of Theorem 2.2, we specify D to be independent of (Y_1, Y_0, X) and follow the Bernoulli distribution with $\Pr(D = 1) = 1/2$. Let $\mathbf{b} = (b_1, \ldots, b_{v-1}) \in \{0,1\}^{v-1}$ be a binary vector that uniquely indexes a member of $\mathcal{P}^*$, and, accordingly, write $\mathcal{P}^* = \left\{P_{\mathbf{b}} : \mathbf{b} \in \{0,1\}^{v-1}\right\}$. For each $j = 1, \ldots, (v-1)$, we specify the conditional distribution of Y_1 given $X = x_j$ to be (A.10) if $b_j = 1$ and (A.11) if $b_j = 0$. For $j = v$, the conditional distribution of Y_1 given $X = x_v$ is degenerate at $Y_1 = \frac{1}{2}$. As for the conditional distribution of Y_0 given $X = x_j$, we consider the degenerate distribution at $Y_0 = 0$ for $j = 1, \ldots, (v-1)$, and the degenerate distribution at $Y_0 = -\frac{1}{2}$ for $X = x_v$. In this specification of

$\mathcal{P}^*$, it holds

$$P_X(|\tau(x)| \leq t) = \begin{cases} 0 & \text{for } t \in [0, \gamma), \\ \left(\frac{\gamma}{\eta}\right)^\alpha & \text{for } t \in [\gamma, \eta), \\ 1 & \text{for } t \in [\eta, 1]. \end{cases},$$

for every $P_{\mathbf{b}} \in \mathcal{P}^*$. Furthermore, by the construction of the support points, for every $P_{\mathbf{b}} \in \mathcal{P}^*$, the first-best decision rule is contained in $\mathcal{G}$. Hence, it holds $\mathcal{P}^* \subset \mathcal{P}_{FB}(1, \kappa, \eta, \alpha)$.

Let $\pi(\mathbf{b})$ be a prior distribution for $\mathbf{b}$ such that $b_1, \ldots, b_{v-1}$ are iid and $b_1 \sim Ber(1/2)$. By following the same line of reasonings as used in obtaining (A.12), for arbitrary estimated treatment choice rule $\hat{G}$, we obtain

$$\begin{aligned} & \sup_{P \in \mathcal{P}(1,\kappa,\eta,\alpha)} E_{P^n}\left[W(G^*) - W(\hat{G})\right] \\ \geq \; & \frac{\gamma}{v-1}\left(\frac{\gamma}{\eta}\right)^\alpha \int_{Z_1,\ldots,Z_n} \sum_{j=1}^{v-1} \left[\min\left\{\pi\left(b_j = 1 | Z_1, \ldots, Z_n\right), 1 - \pi(b_j = 1 | Z_1, \ldots, Z_n)\right\}\right] d\tilde{P}^n. \end{aligned}$$

Furthermore, by repeating the same bounding arguments as in the proof of Theorem 2.2, this Bayes risk can be bounded from below by (A.14) with $p = \frac{1}{v-1}\left(\frac{\gamma}{\eta}\right)^\alpha$,

$$\sup_{P \in \mathcal{P}(1,\kappa,\eta,\alpha)} E_{P^n}\left[W(G^*) - W(\hat{G})\right] \geq \frac{\gamma}{2}\left(\frac{\gamma}{\eta}\right)^\alpha \exp\left\{-\frac{4\gamma}{1-2\gamma}\sqrt{\frac{n}{8(v-1)}\left(\frac{\gamma}{\eta}\right)^\alpha}\right\}.$$

The slowest convergence rate of this lower bound can be obtained by tuning γ to be converging at the rate of $n^{-\frac{1}{2+\alpha}}$. In particular, by choosing $\gamma = \eta^{\frac{\alpha}{2+\alpha}}\left(\frac{v-1}{n}\right)^{\frac{1}{2+\alpha}}$ assuming $\gamma \leq \frac{1}{4}$, the exponential term can be bounded from below by $\exp\left\{-2\sqrt{2}\right\}$, so we obtain the following lower bound,

$$\frac{1}{2}\eta^{\frac{\alpha}{2+\alpha}}(v-1)^{\frac{\alpha}{2+\alpha}} n^{-\frac{1+\alpha}{2+\alpha}} \exp\left\{-2\sqrt{2}\right\}. \tag{A.26}$$

Recall that γ is constrained to $\gamma \leq \min\left\{\eta, \frac{1}{4}\right\}$. This implies that the obtained bound is valid for

$$n \geq \left(\max\left\{\eta^{-1}, 4\right\}\right)^{2+\alpha} \eta^\alpha (v-1),$$

whose stricter but simpler form is given by

$$n \geq \max\left\{\eta^{-2}, 4^{2+\alpha}\right\}(v-1). \tag{A.27}$$

The lower bound presented in this theorem follows by denormalizing the outcomes, i.e., multiply M to (A.26) and substitute η/M for η appearing in (A.26) and (A.27). □

A.4 Proof of Theorem 2.6

The next lemma gives a linearized solution of a certain polynomial inequality. We owe this lemma to Shin Kanaya (2014, personal communication). The technique of applying the mean value expansion to an implicit function defined as the root of a polynomial equation has been used in the context of bandwidth choice in Kanaya and Kristensen (2014).

Lemma A.8. *Let $A \geq 0$, $B \geq 0$, and $X \geq 0$. For any $\alpha \geq 0$, $X \leq AX^{\frac{\alpha}{1+\alpha}} + B$ implies*

$$X \leq A^{1+\alpha} + (1+\alpha)B.$$

Proof. When $A = B = 0$, the conclusion trivially holds. When $B > 0$, $X = AX^{\frac{\alpha}{1+\alpha}} + B$ has a unique root, and we denote it by $X^* = g(A, B)$. When $A > 0$ and $B = 0$, we mean by $g(A, 0)$ the nonzero root of $X = AX^{\frac{\alpha}{1+\alpha}}$. Let $f(X, A, B) = X - AX^{\frac{\alpha}{1+\alpha}} - B$. By the form of the inequality, the original inequality can be equivalently written as $X \leq X^* = g(A, B)$, so we aim to verify that X^* is bounded from above by $A^{1+\alpha} + (1+\alpha)B$. Consider the mean value expansion of $g(A, B)$ in B at $B = 0$,

$$X^* = g(A, 0) + \frac{\partial g}{\partial B}\left(A, \tilde{B}\right) \times B \quad \text{for some } 0 \leq \tilde{B} \leq B.$$

Note $g(A, 0) = A^{1+\alpha}$. In addition, by the implicit function theorem, we have, with $\tilde{X} = g(A, \tilde{B})$,

$$\begin{aligned}
\frac{\partial g}{\partial B}\left(A, \tilde{B}\right) &= -\frac{\frac{\partial f}{\partial B}(\tilde{X}, A, \tilde{B})}{\frac{\partial f}{\partial X}(\tilde{X}, A, \tilde{B})} \\
&= \frac{1}{1 - \frac{\alpha}{1+\alpha} A \tilde{X}^{-\frac{1}{1+\alpha}}} \\
&= \frac{\tilde{X}}{\frac{\tilde{X}}{1+\alpha} + \frac{\alpha}{1+\alpha}\left(\tilde{X} - A\tilde{X}^{\frac{\alpha}{1+\alpha}}\right)} \\
&= \frac{\tilde{X}}{\frac{\tilde{X}}{1+\alpha} + \frac{\alpha}{1+\alpha}\tilde{B}} \\
&\leq 1 + \alpha.
\end{aligned}$$

Hence, $X^* \leq A^{1+\alpha} + (1+\alpha)B$ holds. □

The next lemma provides an exponential tail probability bound of the supremum of the centered empirical processes. This lemma follows from Theorem 2.14.9 in van der Vaart and Wellner (1996) combined with their Theorem 2.6.4.

Lemma A.9. *Assume $\mathcal{G}$ is a VC-class of subsets in $\mathcal{X}$ with VC-dimension $v < \infty$. Let $P_{X,n}(\cdot)$ be the empirical probability distribution on $\mathcal{X}$ constructed upon $(X_1, \ldots, X_n)$ generated iid from $P_X(\cdot)$. Then,*

$$P^n\left(\sup_{G\in\mathcal{G}} |P_{X,n}(G) - P_X(G)| > t\right) \leq \left(\frac{C_4 t}{\sqrt{2v}}\right)^{2v} n^v \exp\left(-nt^2\right)$$

holds for every $t > 0$, where C_4 is a universal constant.

Proof of Theorem 2.6. We first consider the m-hybrid case. Set $\tilde{G} = G^*_{FB}$ in (2.4) and rewrite (2.4) in terms of the normalized welfare loss for $\hat{G}_{m-hybrid}$,

$$\begin{aligned} d(G^*_{FB}, \hat{G}_{m-hybrid}) &\leq \frac{\kappa}{M}\left[W^\tau_n(G^*_{FB}) - \hat{W}^\tau_n(G^*_{FB}) - W^\tau_n(\hat{G}_{m-hybrid}) + \hat{W}^\tau_n\left(\hat{G}_{m-hybrid}\right)\right] \\ &\quad + d(G^*_{FB}, \hat{G}_{m-hybrid}) - d^\tau_n\left(G^*_{FB}, \hat{G}_{m-hybrid}\right) \\ &\leq \frac{1}{n}\sum_{i=1}^n \frac{\kappa}{M}\left[\tau(X_i) - \hat{\tau}^m(X_i)\right]\left[1\{X_i \in G^*_{FB}\} - 1\left\{X_i \in \hat{G}_{m-hybrid}\right\}\right] \\ &\quad + d(G^*_{FB}, \hat{G}_{m-hybrid}) - d^\tau_n\left(G^*_{FB}, \hat{G}_{m-hybrid}\right) \\ &\leq \rho_n + d(G^*_{FB}, \hat{G}_{m-hybrid}) - d^\tau_n\left(G^*_{FB}, \hat{G}_{m-hybrid}\right) \qquad \text{(A.28)} \end{aligned}$$

where $d(G^*_{FB}, \hat{G}_{m-hybrid})$ is as defined in equation (A.16) in Appendix A.3, $d^\tau_n\left(G^*_{FB}, \hat{G}_{m-hybrid}\right) = W^\tau_n(G^*_{FB}) - W^\tau_n(\hat{G}_{m-hybrid})$,

$$\rho_n \equiv \frac{\kappa}{M}\max_{1\leq i\leq n} |\hat{\tau}^m(X_i) - \tau(X_i)|\, P_{X,n}\left(G^*_{FB} \triangle \hat{G}_{m-hybrid}\right),$$

and $P_{X,n}$ is the empirical distribution on $\mathcal{X}$ constructed upon $(X_1, \ldots, X_n)$. Similarly to the proof of Theorem 3.2, define a class of functions generated by $G \in \mathcal{G}$

$$\begin{aligned} \mathcal{H}^\tau &\equiv \{h(Z_i; G) : G \in \mathcal{G}\}, \\ h(Z_i; G) &\equiv \frac{\kappa}{M}\tau(X_i)\cdot\left[1\{X_i \in G\} - 1\{X_i \in G^*_{FB}\}\right], \end{aligned}$$

which is a VC-subgraph class with the VC-dimension at most v with envelope $\bar{H} = 1$ by Lemma A.1. Let $a = \sqrt{kt}\epsilon_n$ be as defined in the proof of Theorem 2.3 and $V^\tau_a \equiv \sup_{h\in\mathcal{H}^\tau}\left\{\frac{E_n(h) - E_P(h)}{-E_P(h) + a^2}\right\}$. By noting

$$d(G^*_{FB}, \hat{G}_{m-hybrid}) - d^\tau_n\left(G^*_{FB}, \hat{G}_{m-hybrid}\right) \leq V^\tau_a(d(G^*_{FB}, \hat{G}_{m-hybrid}) + a^2),$$

inequality (A.28) implies

$$d(G^*_{FB}, \hat{G}_{m-hybrid}) \leq \rho_n + V_a^\tau (d(G^*_{FB}, \hat{G}_{m-hybrid}) + a^2). \tag{A.29}$$

Denote event $\left\{V_a^\tau < \frac{1}{2}\right\}$ by Ω_t, which is equivalent to event $\left\{d(G^*_{FB}, \hat{G}_{m-hybrid}) \leq 2\rho_n + k\epsilon_n^2 t\right\}$. Along the same line of argument that leads to (A.25) in the proof of Theorem 2.3, we obtain, for $t \geq 1$,

$$P^n(\Omega_t) = P^n\left(d(G^*_{FB}, \hat{G}_{m-hybrid}) \leq 2\rho_n + k\epsilon_n^2 t\right) \geq 1 - \exp(-t), \tag{A.30}$$

where ϵ_n is given in (A.23). We bound ρ_n from above by

$$\rho_n \leq \frac{\kappa}{M}\left[\max_{1\leq i\leq n} |\hat{\tau}^m(X_i) - \tau(X_i)| P_X\left(G^*_{FB} \triangle \hat{G}_{m-hybrid}\right) + \mathcal{V}_{0,n} \max_{1\leq i\leq n} |\hat{\tau}^m(X_i) - \tau(X_i)|\right],$$

where

$$\mathcal{V}_{0,n} = \sup_{G\in\mathcal{G}:} |P_{X,n}(G^*_{FB} \triangle G) - P_X(G^*_{FB} \triangle G)|,$$

Let $\lambda > 0$, that will be chosen properly later. Define events

$$\begin{aligned}
\Lambda_1 &= \left\{\mathcal{V}_{0,n} \leq n^{-\lambda}\right\}, \\
\Lambda_2 &= \left\{P_X\left(G^*_{FB} \triangle \hat{G}_{m-hybrid}\right) \geq n^{-\lambda}\right\}.
\end{aligned}$$

Then, on $\Lambda_1 \cap \Lambda_2$, it holds $\mathcal{V}_{0,n} \leq P_X\left(G^*_{FB} \triangle \hat{G}_{m-hybrid}\right)$. Therefore, on $\Lambda_1 \cap \Lambda_2 \cap \Omega_t$, $d(G^*_{FB}, \hat{G}_{m-hybrid})$ can be bounded by

$$\begin{aligned}
d(G^*_{FB}, \hat{G}_{m-hybrid}) &\leq 4\frac{\kappa}{M} \max_{1\leq i\leq n} |\hat{\tau}^m(X_i) - \tau(X_i)| P_X\left(G^*_{FB} \triangle \hat{G}_{m-hybrid}\right) + k\epsilon_n^2 t \\
&\leq 4c_1 \frac{\kappa}{M} \max_{1\leq i\leq n} |\hat{\tau}^m(X_i) - \tau(X_i)| d(G^*_{FB}, \hat{G}_{m-hybrid})^{\frac{\alpha}{1+\alpha}} + k\epsilon_n^2 t,
\end{aligned}$$

where the second line follows from Lemma A.7 with the same definition of c_1 given there. By Lemma A.8 and substituting (A.23) to ϵ_n, we obtain, on event $\Lambda_1 \cap \Lambda_2 \cap \Omega_t$,

$$d(G^*_{FB}, \hat{G}_{m-hybrid}) \leq c_6\left[\max_{1\leq i\leq n} |\hat{\tau}^m(X_i) - \tau(X_i)|\right]^{1+\alpha} + c_7\left(\frac{v}{n}\right)^{\frac{1+\alpha}{2+\alpha}} t, \tag{A.31}$$

where constants c_6 and c_7 depend only on $(M, \kappa, \eta, \alpha)$.

Using the upper bound derived in (A.31), we obtain, for $t \geq 1$,

$$\begin{aligned}
& E_{P^n}\left(d(G^*_{FB}, \hat{G}_{m-hybrid})\right) \\
= \; & E_{P^n}\left(d(G^*_{FB}, \hat{G}_{m-hybrid}) 1\left\{\Lambda_1 \cap \Lambda_2 \cap \Omega_t\right\}\right) + E_{P^n}\left(d(G^*_{FB}, \hat{G}_{m-hybrid}) 1\left\{\Lambda_1^c \cup \Lambda_2^c \cup \Omega_t^c\right\}\right) \\
\leq \; & c_6 E_{P^n}\left(\left[\max_{1\leq i\leq n} \left|\hat{\tau}^m(X_i) - \tau(X_i)\right|\right]^{1+\alpha}\right) + c_7\left(\frac{v}{n}\right)^{\frac{1+\alpha}{2+\alpha}} t + P^n\left(\Lambda_1^c\right) \\
& + E_{P^n}\left(d(G^*_{FB}, \hat{G}_{m-hybrid}) 1\{\Lambda_2^c\}\right) + P^n\left(\Omega_t^c\right) \\
\leq \; & \underbrace{c_6 \psi_n^{-(1+\alpha)} E_{P^n}\left(\left[\psi_n \max_{1\leq i\leq n} \left|\hat{\tau}^m(X_i) - \tau(X_i)\right|\right]^{1+\alpha}\right)}_{A_{1,n}} + \underbrace{c_7\left(\frac{v}{n}\right)^{\frac{1+\alpha}{2+\alpha}} t}_{A_{2,n}} \\
& + \underbrace{\left(\frac{C_4}{2v}\right)^{2v} n^{-2v\left(\lambda-\frac{1}{2}\right)} \exp\left(-n^{-2\left(\lambda-\frac{1}{2}\right)}\right)}_{A_{3,n}} + \underbrace{n^{-\lambda}}_{A_{4,n}} + \underbrace{\exp(-t)}_{A_{5,n}}
\end{aligned}$$

where ψ_n is a sequence as specified in Condition 2.7. In these inequalities, the third line uses (A.31) and $d(G^*_{FB}, \hat{G}_{m-hybrid}) \leq 1$. In the fourth line, $A_{3,n}$ follows from Lemma A.9, $A_{4,n}$ follows from $d(G^*_{FB}, \hat{G}_{m-hybrid}) \leq P_X\left(G^*_{FB} \triangle \hat{G}_{m-hybrid}\right)$ and $P_X\left(G^*_{FB} \triangle \hat{G}_{m-hybrid}\right) < n^{-\lambda}$ on Λ_2^c, and $A_{5,n}$ follows from (A.30).

We now discuss convergence rates of $A_{j,n}$, $j = 1, \ldots, 5$, individually with suitable choices of t and λ. Condition (2.9) assumed in this theorem implies

$$\begin{aligned}
& \sup_{P\in\mathcal{P}_m} E_{P^n}\left(\left(\psi_n \max_{1\leq i\leq n} \left|\hat{\tau}^m(X_i) - \tau(X_i)\right|\right)^{1+\alpha}\right) \\
\leq \; & \sup_{P\in\mathcal{P}_m} E_{P^n}\left(\left[\left(\psi_n \max\nolimits_{1\leq i\leq n} \left|\hat{\tau}^m(X_i) - \tau(X_i)\right|\right)^2\right]^{\frac{1+\alpha}{2}}\right) \\
\leq \; & \left(\left[\sup_{P\in\mathcal{P}_m} E_{P^n}\left(\psi_n \max\nolimits_{1\leq i\leq n} \left|\hat{\tau}^m(X_i) - \tau(X_i)\right|\right)^2\right]^{\frac{1+\alpha}{2}}\right) \\
= \; & O(1),
\end{aligned}$$

where the third line follows from Jensen's inequality. Hence, $A_{1,n}$ satisfies $\sup_{P\in\mathcal{P}_m} A_{1,n} = O\left(\psi_n^{-(1+\alpha)}\right)$. By setting $t = (1+\alpha)\log\psi_n$, we can make the convergence rate of $A_{5,n}$ equal to that of $A_{1,n}$. At the same time, by choosing $\lambda > \frac{1+\alpha}{2+\alpha} \geq \frac{1}{2}$, we can make $A_{3,n}$ and $A_{4,n}$ converge faster than $A_{2,n}$. Hence, the uniform convergence rate of $E_{P^n}\left(d(G^*_{FB}, \hat{G}_{m-hybrid})\right)$ over $P \in \mathcal{P}_m \cap \mathcal{P}_{FB}(M, \kappa, \eta, \alpha)$

is bounded by the convergence rates of the $A_{1,n}$ and $A_{2,n}$,

$$\left(\sup_{P \in \mathcal{P}_m} A_{1,n} \vee \sup_{P \in \mathcal{P}_{FB}(M,\kappa,\eta,\alpha)} A_{2,n} \right) = \left(\psi_n^{-(1+\alpha)} \vee n^{-\frac{1+\alpha}{2+\alpha}} \log \psi_n \right).$$

This completes the proof for the m-hybrid case.

A proof for the e-hybrid case follows almost identically to the proof of the m-hybrid case. The differences are that ρ_n in inequality (A.28) is given by

$$\rho_n = \frac{\kappa}{M} \max_{1 \le i \le n} |\hat{\tau}_i^e - \tau_i| \, P_{X,n} \left(G_{FB}^* \triangle \hat{G}_{e-hybrid} \right).$$

and that inequality (A.32) is replaced by

$$d(G_{FB}^*, \hat{G}_{e-hybrid}) \le \rho_n + V_a(d(G_{FB}^*, \hat{G}_{e-hybrid}) + a^2), \tag{A.32}$$

where V_a is as defined in the proof of Theorem 2.3. The rest of the proof goes similarly to the proof of the first claim except that the rate ϕ_n given in Condition 2.1 replaces ψ_n in the first claim. □

A.5 Proofs of Theorems 4.1 and 4.2

Proof of Theorem 4.1. Since $W_K(G) - W_K(G') = V_K(G) - V_K(G')$ for all G, G',

$$\sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n} \left[\sup_{G \in \mathcal{G}} W_K(G) - W_K(\hat{G}_K) \right] = \sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n} \left[\sup_{G \in \mathcal{G}} V_K(G) - V_K(\hat{G}_K) \right], \tag{A.33}$$

and we focus on bounding the latter expression.

Since $\hat{G}_K$ maximizes $V_{K,n}(G)$, $V_{K,n}(\tilde{G}) \le V_{K,n}(\hat{G}_K)$ for any $\tilde{G} \in \mathcal{G}$ and

$$\begin{aligned} V_K(\tilde{G}) &\le V_{K,n}(\tilde{G}) + \sup_{G \in \mathcal{G}} |V_{K,n}(G) - V_K(G)| \\ &\le V_{K,n}(\hat{G}_K) + \sup_{G \in \mathcal{G}} |V_{K,n}(G) - V_K(G)| \\ &\le V_K(\hat{G}_K) + 2 \sup_{G \in \mathcal{G}} |V_{K,n}(G) - V_K(G)|. \end{aligned}$$

Applying the inequality for all $\tilde{G} \in \mathcal{G}$, we obtain

$$\sup_{G \in \mathcal{G}} V_K(G) - V_K(\hat{G}_K) \le 2 \sup_{G \in \mathcal{G}} |V_{K,n}(G) - V_K(G)|,$$

which is also true in expectation over P^n.

The welfare gain estimation error for any treatment rule G could be bounded from above by:

$$
\begin{aligned}
&|V_{K,n}(G) - V_K(G)| = \left| \frac{K}{\max\{K, P_{X,n}(G)\}} \cdot V_n(G) - \frac{K}{\max\{K, P_X(G)\}} \cdot V(G) \right| \\
&\leq \frac{K}{\max\{K, P_{X,n}(G)\}} \cdot |V_n(G) - V(G)| + V(G) \cdot \left| \frac{K}{\max\{K, P_{X,n}(G)\}} - \frac{K}{\max\{K, P_X(G)\}} \right| \\
&\leq |V_n(G) - V(G)| + \frac{M}{K} \cdot |P_{X,n}(G) - P_X(G)|,
\end{aligned}
$$

The second line comes from subtracting and adding $\frac{K}{\max\{K,P_{X,n}(G)\}}V(G)$ and then applying the triangle inequality. The third line uses inequalities $\frac{K}{\max\{K,P_{X,n}(G)\}} \leq 1$ and $V(G) \leq M$ (from Assumption 2.1 (BO)), and the observation that for any $a, b \in \mathbb{R}$ and $c > 0$,

$$
\left| \frac{c}{\max\{c,a\}} - \frac{c}{\max\{c,b\}} \right| = \left| \frac{c(\max\{c,b\} - \max\{c,a\})}{\max\{c,a\} \cdot \max\{c,b\}} \right| \leq \frac{|\max\{c,b\} - \max\{c,a\}|}{c} \leq \frac{|b-a|}{c}.
$$

Then

$$
\begin{aligned}
&\sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n} \left[\sup_{G \in \mathcal{G}} V_K(G) - V_K(\hat{G}_K) \right] \leq 2 \sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n} \left[\sup_{G \in \mathcal{G}} |V_{K,n}(G) - V_K(G)| \right] \\
&\leq 2 \sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n} \left[\sup_{G \in \mathcal{G}} |V_n(G) - V(G)| \right] + 2\frac{M}{K} \sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n} \left[\sup_{G \in \mathcal{G}} |P_{X,n}(G) - P_X(G)| \right]
\end{aligned}
$$

Note that since the class $\mathcal{G}$ has VC-dimension $v < \infty$, the classes of functions

$$
\begin{aligned}
f_G(Y, D, X) &\equiv \left(\frac{YD}{e(X)} - \frac{Y(1-D)}{1-e(X)} \right) \cdot 1\{X \in G\}, \\
h_G(Y, D, X) &\equiv 1\{X \in G\} - 1/2,
\end{aligned}
$$

are VC-subgraph classes with VC-dimension no greater than v by Lemma A.1. These classes of functions are uniformly bounded by $M/(2\kappa)$ and $1/2$. Since $V_n(G) = E_n(f_G)$, $V(G) = E(f_G)$, $P_n(X_i \in G) = E_n(h_G) + 1/2$ and $P(X \in G) = E(h_G) + 1/2$, we could apply Lemma A.4 and obtain

$$
\sup_{P \in \mathcal{P}(M,\kappa)} E_{P^n} \left[\sup_{G \in \mathcal{G}} V_K(G) - V_K(\hat{G}_K) \right] \leq C_1 \frac{M}{\kappa} \sqrt{\frac{v}{n}} + C_1 \frac{M}{K} \sqrt{\frac{v}{n}}.
$$

The theorem's result follows from (A.33). □

Proof of Theorem 4.2. Define the class of functions $\mathcal{F} = \{f(\cdot; G) : G \in \mathcal{G}\}$ on $\mathcal{Z}$:

$$
f(Z_i; G) \equiv \left[\left(\frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)} \right) \cdot \rho(X_i) \cdot 1\{X_i \in G\} \right],
$$

Then $V^T(G) = E_P(f(\cdot, G))$ by (4.3) and $\hat{G}^T_{EWM} = \arg\max_{G in \mathcal{G}} E_n(f(\cdot, G))$.

By Assumptions 2.1 (BO), (SO) and 4.2 (BDR), these functions are uniformly bounded by $\bar{F} = \frac{M\bar{\rho}}{2\kappa}$. By Assumption 2.1 (VC) and Lemma A.1, $\mathcal{F}$ is a VC-subgraph class of functions with VC-dimension at most v. Applying the argument in inequality (2.2) we obtain

$$\sup_{G\in\mathcal{G}} W^T(G) - W^T\left(\hat{G}^T_{EWM}\right) = \sup_{G\in\mathcal{G}} V^T(G) - V^T\left(\hat{G}^T_{EWM}\right) \leq 2 \sup_{f\in\mathcal{F}} |E_n(f) - E_P(f)| .$$

Now we take the expectation with respect to the sampling distribution P^n and apply Lemma A.4:

$$E_{P^n}\left[\sup_{G\in\mathcal{G}} W^T(G) - W^T(\hat{G}^T_{EWM})\right] \leq 2E_{P^n}\left[\sup_{f\in\mathcal{F}} |E_n(f) - E_P(f)|\right] \leq C_1 \frac{M\bar{\rho}}{\kappa}\sqrt{\frac{v}{n}}.$$

□

B Validating Condition 2.1 for Local Polynomial Estimators

This Appendix verifies that Condition 2.1 (m) and (e) hold for local polynomial estimators if the class of data generating processes $\mathcal{P}_m$ or $\mathcal{P}_e$ is constrained by Assumptions 2.3 and 2.4, respectively.

B.1 Notations and Basic Lemmas

In addition to the notations introduced in Section 2.4 in the main text, the following notations are used. Let $\mu : \mathbb{R}^{d_x} \to \mathbb{R}$ be a generic notation for a regression equation onto a vector of covariates $X \in \mathbb{R}^{d_x}$. In case of m-hybrid EWM, $\mu(\cdot)$ corresponds to either of $m_1(\cdot)$ or $m_0(\cdot)$. In case of e-hybrid EWM, $\mu(\cdot)$ corresponds to propensity score $e(\cdot)$. We use n to denote the size of the entire sample indexed by $i = 1, \ldots, n$, and denote by $J_i \subset \{1, \ldots, n\}$ a subsample with which $\mu(X_i)$ is estimated nonparametrically. Since we consider throughout the leave-one-out regression fits of $\mu(X_i)$, J_i does not include i-th observation. In case of m-hybrid EWM, J_i is either the leave-one-out treated sample $\{j \in \{1, \ldots, n\} : D_j = 1, j \neq i\}$ or the leave-one-out control sample $\{j \in \{1, \ldots, n\} : D_j = 0, j \neq i\}$ depending on $\mu(\cdot)$ corresponds to $m_1(\cdot)$ or $m_0(\cdot)$. Note that, in the m-hybrid case, J_i is random as it depends on a realization of $(D_1, \ldots, D_n)$. When the e-hybrid EWM is considered, J_i is non-stochastic and it is given by $J_i = \{1, \ldots, n\} \setminus \{i\}$. The size of J_i is denoted by n_{J_i}, which is equal to $n_1 - 1$ or $n_0 - 1$ in the m-hybrid case, and is equal to $n - 1$ in case of e-hybrid case. With abuse of notations, we use Y_i, $i = 1, \ldots, n$, to denote observations of the regressors and use ξ_i to denote a regression residual, i.e., $Y_i = \mu(X_i) + \xi_i$, $E(\xi_i | X_i) = 0$, holds for all $i = 1, \ldots, n$. For e-hybrid rule, Y_i should be read as the treatment status indicator $D_i \in \{1, 0\}$.

We assume that $\mu(\cdot)$ belongs to a Hölder class of functions with degree $\beta \geq 1$ and constant $0 < L < \infty$. Define the leave-one-out local polynomial regression with degree $l = (\beta - 1)$ by

$$\begin{aligned} \hat{\mu}_{-i}(X_i) &= U^T(0)\hat{\theta}(X_i) \cdot 1\{\lambda(X_i) \geq t_n\}, \\ \hat{\theta}_{-i}(X_i) &= \arg\min_{\theta} \sum_{j \in J_i} \left[Y_j - \theta^T U\left(\frac{X_j - X_i}{h}\right) \right]^2 K\left(\frac{X_j - X_i}{h}\right), \end{aligned} \tag{B.1}$$

where $U\left(\frac{X_j - X_i}{h}\right)$ is a regressor vector, $U\left(\frac{X_j - X_i}{h}\right) \equiv \left(\left(\frac{X_j - X_i}{h}\right)^s\right)_{|s| \leq l}$, $\lambda(X_i)$ is a smallest eigenvalue of $B_{-i}(X_i) \equiv \left(nh^{d_x}\right)^{-1} \sum_{j \in J_i} U\left(\frac{X_j - X_i}{h}\right) U^T\left(\frac{X_j - X_i}{h}\right) K\left(\frac{X_i - X_j}{h}\right)$, and t_n is a sequence of trimming constant converging to zero, whose choice will be discussed later. The standard least squares calculus shows

$$\hat{\theta}_{-i}(X_i) = B_{-i}(X_i)^{-1} \left(\frac{1}{nh^{d_x}} \sum_{j \in J_i} U\left(\frac{X_j - X_i}{h}\right) K\left(\frac{X_j - X_i}{h}\right) \right),$$

so that $\hat{\mu}(X_i)$ can be written as

$$\begin{aligned} \hat{\mu}_{-i}(X_i) &= \left[\sum_{j \in J_i} Y_j \omega_j(X_i) \right] \cdot 1\{\lambda(X_i) \geq t_n\}, \\ \text{where } \omega_j(X_i) &= \frac{1}{nh^{d_x}} U^T(0) \left[B_{-i}(X_i)\right]^{-1} U\left(\frac{X_j - X_i}{h}\right) K\left(\frac{X_j - X_i}{h}\right). \end{aligned} \tag{B.2}$$

We first present lemmas that will be used for proving Lemma B.4 below.

Lemma B.1. *Suppose Assumptions 2.3 (PX) and (Ker).*

(i) Conditional on $(X_1, \ldots, X_n)$ such that $\lambda(X_i) > 0$,

$$\begin{aligned} \max_{j \neq i} |\omega_j(X_i)| &\leq c_5 \frac{1}{nh^{d_x} \lambda(X_i)}, \\ \sum_{j \in J_i} |\omega_j(X_i)| &\leq \frac{c_5}{nh^{d_x} \lambda(X_i)} \sum_{j \in J_i} 1\left\{(X_j - X_i) \in [-h, h]^{d_x}\right\}, \end{aligned}$$

where c_5 is a constant that depends only on β, d_x and $K_{\max}$.

(ii) For any multi-index s such that $|s| \leq (\beta - 1)$, $\sum_{j \in J_i} \left(\frac{X_j - X_i}{h}\right)^s \omega_j(X_i) = 0$.

(iii) Let $\lambda(x)$ be a smallest eigenvalue of $B(x) \equiv \left(nh^{d_x}\right)^{-1} \sum_{j=1}^n U\left(\frac{X_j - x}{h}\right) U^T\left(\frac{X_j - x}{h}\right) K\left(\frac{X_j - x}{h}\right)$ there exist positive constants c_6 and c_7 that depend only on $\underline{c}$, r_0, $\underline{p}_X$, and $K(\cdot)$ such that

$$P^n\left(\{\lambda(x) \leq c_6\}\right) \leq 2\left[\dim U\right]^2 \exp\left(-c_7 nh^{d_x}\right)$$

holds for all x, P_X-almost surely, at every $n \geq 1$.

Proof. (i) Since $\|U(0)\| = 1$, it holds

$$\begin{aligned}
|\omega_j(X_i)| &\leq \frac{1}{nh^{d_x}} \left\| [B_{-i}(X_i)]^{-1} U\left(\frac{X_j - X_i}{h}\right) K\left(\frac{X_j - X_i}{h}\right) \right\| \\
&\leq \frac{K_{\max}}{nh^{d_x}\lambda(X_i)} \left\| U\left(\frac{X_j - X_i}{h}\right) 1\left\{(X_j - X_i) \in [-h,h]^{d_x}\right\} \right\| \\
&\leq \frac{K_{\max} \dim(U)^{1/2}}{nh^{d_x}\lambda(X_i)} \\
&\equiv \frac{c_5}{nh^{d_x}\lambda(X_i)},
\end{aligned}$$

for every $1 \leq j \leq n$. Similarly,

$$\begin{aligned}
\sum_{j \in J_i} |\omega_j(X_i)| &\leq \frac{K_{\max}}{nh^{d_x}\lambda(X_i)} \sum_{j \in J_i} \left\| U\left(\frac{X_j - X_i}{h}\right) \right\| 1\left\{(X_j - X_i) \in [-h,h]^{d_x}\right\} \\
&= \frac{c_5}{nh^{d_x}\lambda(X_i)} \sum_{j \in J_i} 1\left\{(X_j - X_i) \in [-h,h]^{d_x}\right\}.
\end{aligned}$$

(ii) This claim follows from the first order condition for θ in the least square minimization problem in (B.1).

(iii) This lemma is from Equation (6.3, pp. 626) in the proof of Theorem 3.2 in Audibert and Tsybakov (2007), where a suitable choice of constant c_6 is given in Equation (6.2, pp.625) in Audibert and Tsybakov (2007). □

The next lemma provides an exponential tail bound for the local polynomial estimators. The first statement is borrowed from Theorem 3.2 in Audibert and Tsybakov (2007), and the second statement is its immediate extension.

Lemma B.2. *(i) Suppose Assumption 2.3 (PX) and (Ker) hold, and $\mu(\cdot)$ belongs to a Hölder class of functions with degree $\beta \geq 1$ and constant $0 < L < \infty$. Assume J_i is non-stochastic with $n_{J_i} = n - 1$ (e-hybrid case). Then, there exist positive constants c_8, c_9, and c_{10} that depend only on β, d_x, L, $\underline{c}$, r_0, $\underline{p}_X$, and $\bar{p}_X$, such that, for any $0 < h < r_0/\underline{c}$, any $c_8 h^\beta < \delta$, and any $n \geq 2$,*

$$P^{n-1}\left(\left|\hat{\mu}_{-n}(x) - \mu(x)\right| > \delta\right) \leq c_9 \exp\left(-c_{10} n h^{d_x} \delta^2\right),$$

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

(ii) Suppose Assumptions 2.1 (SO), 2.3 (PX), and (Ker) hold, and $\mu(\cdot)$ belongs to a Hölder class of functions with degree $\beta \geq 1$ and constant $0 < L < \infty$. Assume J_i is stochastic (m-hybrid case) with $J_i = \{j \neq i : D_j = d\}$, $d \in \{1, 0\}$. There exist positive constants c_{11}, c_{12}, and c_{13} that

depend only on κ, β, d_x, L, $\underline{c}$, r_0, $\underline{p}_X$, *and* $\bar{p}_X$, *such that for any* $0 < h < r_0/\underline{c}$, *any* $c_{11}h^\beta < \delta$, *and any* $n_{J_n} \geq 1$,

$$P^{n-1}\left(\left|\hat{\mu}_{-n}(x) - \mu(x)\right| > \delta | n_{J_n}\right) \leq c_{12}\exp\left(-c_{13}n_{J_n}h^{d_x}\delta^2\right)$$

holds for almost all x *with respect to* P_X, *where* $P^{n-1}(\cdot|n_{J_n})$ *is the conditional distribution of* $\left\{(Y_i, X_i)_{i=1}^{n-1}\right\}$ *given* $\sum_{j=1}^{n-1} 1\{D_j = d\}$.

Proof. (i) See Theorem 3.2 in Audibert and Tsybakov (2007).

(ii) Under Assumption 2.1 (SO), the conditional distribution of covariates X given $D = d$, $d \in \{1, 0\}$, has the support $\mathcal{X}$ same as the unconditional distribution P_X, and has bounded density on $\mathcal{X}$, since

$$\frac{\kappa}{1-\kappa}\frac{dP_X}{dx} < \frac{dP_{X|D=d}}{dx} < \frac{1-\kappa}{\kappa}\frac{dP_X}{dx}$$

holds for all $x \in \mathcal{X}$. Therefore, when P_X satisfies Assumption 2.3 (PX), the conditional distributions $P_{X|D=d}$, $d \in \{1, 0\}$ also satisfy the support and density conditions analogous to Assumption 2.3 (PX). This implies that, even when we condition on $n_{J_n} = \sum_{j=1}^{n-1} 1\{D_j = d\} \geq 1$, the exponential inequality of (i) in the current lemma is applicable with different constant terms. □

The next lemma concerns an upper bound of the variance of the supremum of centered empirical processes indexed by a class of sets.

Lemma B.3. *Let* $\mathcal{B}$ *be a countable class of sets in* $\mathcal{X}$, *and let* $\{P_{X,n}(B) : B \in \mathcal{B}\}$ *be the empirical distribution based on iid observations,* $(X_1, \ldots, X_n)$, $X_i \sim P_X$.

$$Var\left(\sup_{B\in\mathcal{B}}\{P_{X,n}(B) - P_X(B)\}\right) \leq \frac{2}{n}E\left[\sup_{B\in\mathcal{B}}\{P_{X,n}(B) - P_X(B)\}\right] + \frac{1}{4n}.$$

Proof. In Theorem 11.10 of Boucheron et al. (2013), setting $X_{i,s}$ at the centered indicator function $1\{X_i \in B\} - P_X(B)$, and dividing the inequality of Theorem 11.10 of Boucheron et al. (2013) by n^2 lead to

$$\begin{aligned} Var\left(\sup_{B\in\mathcal{B}}\{P_{X,n}(B) - P_X(B)\}\right) &\leq \frac{2}{n}E\left[\sup_{B\in\mathcal{B}}\{P_{X,n}(B) - P_X(B)\}\right] \\ &\quad + \frac{1}{n}\sup_{B\in\mathcal{B}}\{P_X(B)[1 - P_X(B)]\} \\ &\leq \frac{2}{n}E\left[\sup_{B\in\mathcal{B}}\{P_{X,n}(B) - P_X(B)\}\right] + \frac{1}{4n}. \end{aligned}$$

□

B.2 Main Lemmas and Proofs of Corollaries 2.1 and 2.2

The next lemma yields Corollaries 2.1 and 2.2.

Lemma B.4. *Let $\mathcal{P}_\mu$ be a class of joint distributions of (Y, X) such that $\mu(\cdot)$ belongs to a Hölder class of functions with degree $\beta \geq 1$ and constant $0 < L < \infty$, and Assumption 2.3 (PX) holds. Let $\hat{\mu}_{-i}(\cdot)$ be the leave-one-out local polynomial fit for $\mu(X_i)$ defined in (B.1), whose kernel function satisfies Assumption 2.3 (Ker).*

(i) Then,

$$\sup_{P \in \mathcal{P}_\mu} E_{P^n}\left[\frac{1}{n}\sum_{i=1}^{n}\left|\hat{\mu}_{-i}(X_i) - \mu(X_i)\right|\right] \leq O(h^\beta) + O\left(\frac{1}{\sqrt{nh^{d_x}}}\right) \tag{B.3}$$

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

$$\sup_{P \in \mathcal{P}_\mu} E_{P^n}\left[\frac{1}{n}\sum_{i=1}^{n}\left|\hat{\mu}_{-i}(X_i) - \mu(X_i)\right|\right] \leq O\left(n^{-\frac{1}{2+d_x/\beta}}\right).$$

(ii) Let $t_n = (\log n)^{-1}$. Then,

$$\sup_{P \in \mathcal{P}_\mu} E_{P^n}\left[\left(\max_{1\leq i\leq n}\left|\hat{\mu}_{-i}(X_i) - \mu(X_i)\right|\right)^2\right] \leq O\left(\frac{h^{2\beta}}{t_n^2}\right) + O\left(\frac{\log n}{nh^{d_x}t_n^2}\right) \tag{B.4}$$

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

$$\sup_{P \in \mathcal{P}_\mu} E_{P^n}\left[\left(\max_{1\leq i\leq n}\left|\hat{\mu}_{-i}(X_i) - \mu(X_i)\right|\right)^2\right] \leq O\left((t_n)^{-2}\left(\frac{\log n}{n}\right)^{\frac{2}{2+d_x/\beta}}\right).$$

Proof. (i) First, consider the non-stochastic J_i case with $n_{J_i} = (n-1)$ (e-hybrid case). Since observations are iid (hence exchangeable) and the probability law of $\hat{\mu}_{-i}(\cdot)$ does not depend on X_i, it holds

$$\begin{aligned}
E_{P^n}\left[\frac{1}{n}\sum_{i=1}^{n}\left|\hat{\mu}_{-i}(X_i) - \mu(X_i)\right|\right] &= E_{P^n}\left|\hat{\mu}_{-i}(X_i) - \mu(X_i)\right| \qquad \text{(B.5)}\\
&= E_{P_X}\left[E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(X_n) - \mu(X_n)\right| | X_n\right]\right]\\
&= \int_{\mathcal{X}} E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(x) - \mu(x)\right|\right] dP_X(x)\\
&= \int_{\mathcal{X}}\left[\int_0^\infty P^{n-1}\left(\left|\hat{\mu}_{-n}(x) - \mu(x)\right| > \delta\right) d\delta\right] dP_X(x),
\end{aligned}$$

where $E_{P^{n-1}}[\cdot]$ is the expectation with respect to the first $(n-1)$-observations of (Y_i, X_i). By Lemma B.2 (i), there exist positive constants c_8, c_9, and c_{10} that depend only on β, d_x, L, $\underline{c}$, r_0, $\underline{p}_X$, and $\bar{p}_X$ such that, for any $0 < h < r_0/\underline{c}$, any $c_8 h^\beta < \delta$, and any $n \geq 2$,

$$P^{n-1}\left(\left|\hat{\mu}_{-n}(x) - \mu(x)\right| > \delta\right) \leq c_9 \exp\left(-c_{10} n h^{d_x} \delta^2\right) \tag{B.6}$$

holds for almost all x with respect to P_X. Hence,

$$\begin{aligned}
\int_{\mathcal{X}} \left[\int_0^\infty P^{n-1}\left(\left|\hat{\mu}_{-n}(x) - \mu(x)\right| > \delta\right) d\delta\right] dP_X(x) &\leq c_8 h^\beta + c_9 \int_0^\infty \exp\left(-c_{10} n h^{d_x} \delta^2\right) d\delta \\
&= c_8 h^\beta + \frac{c_{14}}{\sqrt{n h^{d_x}}} \\
&= O(h^\beta) + O\left(\frac{1}{\sqrt{n h^{d_x}}}\right)
\end{aligned} \tag{B.7}$$

where $c_{14} = c_9 (2c_{10})^{-1/2} \int_0^\infty \left(\delta'\right)^{-1/2} \exp\left(-c_{10}\delta'\right) d\delta' < \infty$. Since the upper bound (B.7) does not depend upon $P \in \mathcal{P}_\mu$, this upper bound is uniform over $P \in \mathcal{P}_\mu$, so the conclusion holds.

Next, consider the stochastic J_i case with $n_{J_i} = \sum_{j \neq i} 1\{D_j = d\}$, where $d \in \{1, 0\}$. we can interpret n_{J_i} as a binomial random variable with parameters $(n-1)$ and π, where $\pi = P(D_i = 1)$ when $\mu(\cdot)$ corresponds to $m_1(\cdot)$ and $\pi = P(D_i = 0)$ when $\mu(\cdot)$ corresponds to $m_0(\cdot)$. In either case, $\kappa < \pi < 1 - \kappa$ by Assumption 2.1 (SO). Let $n \geq 1 + \frac{2}{\pi}$ and $\Omega_{\pi,n} \equiv \left\{\left|\frac{n_{J_n}}{n-1} - \pi\right| \leq \frac{1}{2}\pi\right\} = \left\{\frac{(n-1)\pi}{2} \leq n_{J_n} \leq \frac{3(n-1)\pi}{2}\right\}$. Consider

$$\begin{aligned}
E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(x) - \mu(x)\right| \cdot 1\{\Omega_{\pi,n}\}\right] &= \sum_{n_{J_n} \in \Omega_{\pi,n}} E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(x) - \mu(x)\right| | n_{J_n}\right] P^{n-1}(n_{J_n}) \\
&\leq \max_{n_{J_n} \in \Omega_{\pi,n}} \left\{E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(x) - \mu(x)\right| | n_{J_n}\right]\right\} P^{n-1}(\Omega_{\pi,n}) \\
&\leq \max_{n_{J_n} \in \Omega_{\pi,n}} \left\{E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(x) - \mu(x)\right| | n_{J_n}\right]\right\}.
\end{aligned}$$

Since $n_{J_n} \geq \frac{(n-1)\pi}{2} \geq 1$ on $\Omega_{\pi,n}$, Lemma B.2 (ii) implies

$$\begin{aligned}
E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(x) - \mu(x)\right| | n_{J_n}\right] &\leq \int_{\mathcal{X}} \left[\int_0^\infty P^{n-1}\left(\left|\hat{\mu}_{-n}(x) - \mu(x)\right| > \delta | n_{J_n}\right) d\delta\right] dP_X(x) \\
&\leq c_{11} h^\beta + \frac{c_{15}}{\sqrt{n_{J_n} h^{d_x}}},
\end{aligned}$$

where c_{11} and c_{15} are positive constants that depend only on κ, β, d_x, L, $\underline{c}$, r_0, $\underline{p}_X$, and $\bar{p}_X$. Since $n_{J_n} \geq \frac{(n-1)\pi}{2} \geq \frac{n\pi}{4}$ on $\Omega_{\pi,n}$ for $n \geq 2$, it holds

$$\max_{n_{J_n} \in \Omega_{\pi,n}} \left\{E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(x) - \mu(x)\right| | n_{J_n}\right]\right\} \leq c_{11} h^\beta + \frac{2c_{15}}{\sqrt{\pi n h^{d_x}}}.$$

Accordingly, combined with the Hoeffding's inequality $P^{n-1}\left(\Omega_{\pi,n}^{c}\right) \leq 2\exp\left(-\frac{\pi^2}{4}n\right)$, we obtain

$$\begin{aligned} E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(x)-\mu\left(x\right)\right|\right] &\leq E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(x)-\mu\left(x\right)\right|\cdot 1\left\{\Omega_{\pi,n}\right\}\right]+MP^{n-1}\left(\Omega_{\pi,n}^{c}\right) \\ &\leq c_{11}h^{\beta}+\frac{2c_{15}}{\sqrt{\pi nh^{d_x}}}+2M\exp\left(-\frac{\pi^2}{4}n\right). \end{aligned}$$

The third term in the right hand side converges faster than the second term, so we have shown

$$\begin{aligned} E_{P^n}\left[\frac{1}{n}\sum_{i=1}^{n}\left|\hat{\mu}_{-i}(X_i)-\mu\left(X_i\right)\right|\right] &= \int_{\mathcal{X}} E_{P^{n-1}}\left[\left|\hat{\mu}_{-n}(x)-\mu\left(x\right)\right|\right]dP_X(x) \\ &\leq O(h^{\beta})+O\left(\frac{1}{\sqrt{nh^{d_x}}}\right) \end{aligned}$$

holds for the stochastic J_i case as well.

(ii) Let $\Omega_{\lambda,n}$ be an event defined by $\{\lambda\left(X_i\right)\geq t_n,\ \forall i=1,\ldots,n\}$. On $\Omega_{\lambda,n}$, (B.2) implies

$$\begin{aligned} \left|\hat{\mu}_{-i}(X_i)-\mu\left(X_i\right)\right|^2 &\leq \left|\sum_{j\in J_i} Y_j\omega_j\left(X_i\right)-\mu\left(X_i\right)\right|^2 \\ &= \left|\sum_{j\in J_i}\left(\mu\left(X_j\right)-\mu\left(X_i\right)\right)\omega_j\left(X_i\right)+\sum_{j\in J_i}\xi_j\omega_j\left(X_i\right)\right|^2 \\ &\leq 2\left|\sum_{j\in J_i}\left(\mu\left(X_j\right)-\mu\left(X_i\right)\right)\omega_j\left(X_i\right)\right|^2+2\left|\sum_{j\in J_i}\xi_j\omega_j\left(X_i\right)\right|^2, \end{aligned} \tag{B.8}$$

where the second line follows from $Y_j=\mu\left(X_j\right)+\xi_j$ and $\sum_{j\neq i}\omega_j\left(X_i\right)=0$ as implied by Lemma B.1 (ii). Since $\mu\left(\cdot\right)$ is assumed to belong to the Hölder class, Lemma B.1 (ii) and Assumption 2.3 (Ker) imply

$$\begin{aligned} &\left|\sum\nolimits_{j\in J_i}\left(\mu\left(X_j\right)-\mu\left(X_i\right)\right)\omega_j\left(X_i\right)\right|^2=\left|\sum\nolimits_{j\in J_i}\left\|X_j-X_i\right\|^{\beta}\omega_j\left(X_i\right)\right|^2 \\ &=\left|\sum\nolimits_{j\in J_i}\left\|X_j-X_i\right\|^{\beta}\omega_j\left(X_i\right)\cdot 1\left\{\left(X_j-X_i\right)\in[-h,h]^{d_x}\right\}\right|^2 \\ &\leq d_x^{2\beta}h^{2\beta}\left|\sum\nolimits_{j\in J_i}\left|\omega_j\left(X_i\right)\right|\right|^2 \\ &\leq d_x^{2\beta}h^{2\beta}\left(\frac{c_5}{\lambda(X_i)}\right)^2\left(\frac{1}{nh^{d_x}}\sum\nolimits_{j\in J_i}1\left\{\left(X_j-X_i\right)\in[-h,h]^{d_x}\right\}\right)^2 \\ &\leq c_{16}\frac{h^{2\beta}}{t_n^2}\left(\frac{1}{nh^{d_x}}\sum\nolimits_{j\in J_i}1\left\{\left(X_j-X_i\right)\in[-h,h]^{d_x}\right\}\right)^2, \end{aligned}$$

where $c_{16} = d_x^{2\beta} c_5^2$. Under Assumption 2.3 (PX) and conditional on $\Omega_{\lambda,n}$,

$$\begin{aligned}
&\max_{1\le i\le n}\left|\sum\nolimits_{j\in J_i}\left(\mu\left(X_j\right)-\mu\left(X_i\right)\right)\omega_j\left(X_i\right)\right|^2 \le c_{16}\frac{h^{2\beta}}{t_n^2}\left[\frac{1}{h^{d_x}}\sup_{B\in\mathcal{B}_h}P_{X,n}(B)\right]^2\\
&\le c_{16}\frac{h^{2\beta}}{t_n^2}\left[\frac{1}{h^{d_x}}\left(\sup_{B\in\mathcal{B}_h}\left(P_{X,n}(B)-P_X\left(B\right)\right)+\sup_{B\in\mathcal{B}_h}P_X\left(B\right)\right)\right]^2\\
&\le c_{16}\frac{h^{2\beta}}{t_n^2}\left[\frac{1}{h^{d_x}}\sup_{B\in\mathcal{B}_h}\left(P_{X,n}(B)-P_X\left(B\right)\right)+2^{d_x}\cdot\bar{p}_X\right]^2\\
&\le c_{16}\frac{h^{2\beta}}{t_n^2}\left\{\frac{2}{h^{2d_x}}\left[\sup_{B\in\mathcal{B}_h}\left(P_{X,n}(B)-P_X\left(B\right)\right)\right]^2+2^{2d_x+1}\cdot\bar{p}_X^2\right\},
\end{aligned}$$

where $\mathcal{B}_h$ is the class of hypercubes in $\mathbb{R}^{d_x}$, $\mathcal{B}_h \equiv \left\{\prod_{k=1}^{d_x}[x_k-h,x_k+h] : (x_1,\ldots,x_{d_x})\in\mathcal{X}\right\}$. Accordingly,

$$\begin{aligned}
&E_{P^n}\left[\max_{1\le i\le n}\left|\sum\nolimits_{j\neq i}\left(\mu\left(X_j\right)-\mu\left(X_i\right)\right)\omega_j\left(X_i\right)\right|^2\cdot 1\left\{\Omega_{\lambda,n}\right\}\right]\\
\le\ & c_{17}\frac{h^{2\beta}}{t_n^2}+2c_{16}\frac{h^{2\beta}}{t_n^2}\frac{1}{h^{2d_x}}E_{P^n}\left\{\left[\sup\nolimits_{B\in\mathcal{B}_h}\left(P_{X,n}(B)-P_X\left(B\right)\right)\right]^2\right\}\\
\le\ & c_{17}\frac{h^{2\beta}}{t_n^2}+4c_{16}\frac{h^{2\beta}}{t_n^2}\frac{1}{h^{2d_x}}\left\{\begin{array}{c} Var\left(\sup\nolimits_{B\in\mathcal{B}_h}\left(P_{X,n}(B)-P_X\left(B\right)\right)\right)\\ +\left[E_{P^n}\left(\sup\nolimits_{B\in\mathcal{B}_h}\left(P_{X,n}(B)-P_X\left(B\right)\right)\right)\right]^2\end{array}\right\},
\end{aligned}$$

where $c_{17} = 2^{2d_x+1}c_{16}\bar{p}_X^2$. In order to bound the variance and the squared mean terms in the curly brackets, we apply Lemma B.3 and Lemma A.5 with $\bar{F}=1$ and $\delta = \bar{p}_X(2h)^{d_x/2}$. For all n satisfying $nh^{d_x} \ge \frac{C_1 v_{\mathcal{B}_h}}{2^{d_x}\bar{p}_X^2}$, it holds

$$\begin{aligned}
Var\left(\sup_{B\in\mathcal{B}_h}\left(P_{X,n}(B)-P_X\left(B\right)\right)\right) &\le \frac{2}{n}E_{P^n}\left(\sup_{B\in\mathcal{B}_h}\left(P_{X,n}(B)-P_X\left(B\right)\right)\right)+\frac{1}{4n}\\
&\le 2^{\frac{d_x}{2}+1}C_2\bar{p}_X\frac{\sqrt{v_{\mathcal{B}_h}h^{d_x}}}{n^{3/2}}+\frac{1}{4n},\\
\left[E_{P^n}\left(\sup_{B\in\mathcal{B}_h}\left(P_{X,n}(B)-P_X\left(B\right)\right)\right)\right]^2 &\le 2^{d_x}C_2^2\bar{p}_X^2\frac{v_{\mathcal{B}_h}h^{d_x}}{n},
\end{aligned}$$

where $v_{\mathcal{B}_h}<\infty$ is the VC-dimension of $\mathcal{B}_h$ that depends only on d_x. As a result, there exist positive constants c_{18}, and c_{19} that depend only on β, d_x, and $\bar{p}_X$, such that

$$E_{P^n}\left[\max_{1\le i\le n}\left|\sum_{j\neq i}\left(\mu\left(X_j\right)-\mu\left(X_i\right)\right)\omega_j\left(X_i\right)\right|^2\cdot 1\left\{\Omega_{\lambda,n}\right\}\right] \le c_{17}\frac{h^{2\beta}}{t_n^2}+c_{18}\frac{h^{2\beta}}{t_n^2\left(nh^{d_x}\right)}+c_{19}\frac{h^{2\beta}}{t_n^2\left(nh^{d_x}\right)^{3/2}}$$

holds for all n satisfying $nh^{d_x} \geq \frac{C_1 v_{\mathcal{B}_h}}{2^{d_x} \bar{p}_X^2}$. Since $nh^{d_x} \to \infty$ by the assumption, focusing on the leading terms yields

$$\limsup_{n\to\infty} \sup_{P\in P_\mu} E_{P^n}\left[2\max_{1\leq i\leq n}\left|\sum_{j\in J_i}\left(\mu\left(X_j\right)-\mu\left(X_i\right)\right)\omega_j\left(X_i\right)\right|^2 \cdot 1\left\{\Omega_{\lambda,n}\right\}\right] \leq O\left(\frac{h^{2\beta}}{t_n^2}\right). \tag{B.9}$$

In order to bound the second term in the right hand side of (B.8), note first that

$$\begin{aligned}\left|\sum_{j\in J_i}\xi_j\omega_j\left(X_i\right)\right|^2 &\leq \frac{1}{nh^d\lambda^2(X_i)}\left\|\frac{1}{\sqrt{nh^{d_x}}}\sum_{j\in J_i}\xi_j U\left(\frac{X_j-X_i}{h}\right)K\left(\frac{X_j-X_i}{h}\right)\right\|^2 \\ &\leq \frac{K_{\max}^2}{nh^d t_n^2}\max_{1\leq k\leq \dim(U)}\eta_{ik}^2\end{aligned}$$

holds conditional on $\Omega_{\lambda,n}$, where η_{ik}, $1\leq k\leq \dim(U)$, is the k-th entry of vector

$$\frac{1}{\sqrt{nh^{d_x}}}\sum_{j\in J_i}\xi_j U\left(\frac{X_j-X_i}{h}\right)1\left\{\left(X_j-X_i\right)\in[-h,h]^{d_x}\right\}.$$

Therefore,

$$E_{P^n}\left[\max_{1\leq i\leq n}\left|\sum_{j\in J_i}\xi_j\omega_j\left(X_i\right)\right|^2 \cdot 1\left\{\Omega_{\lambda,n}\right\}\right] \leq \frac{K_{\max}^2}{nh^d t_n^2} E_{P^n}\left[\max_{1\leq i\leq n,1\leq k\leq \dim(U)}\eta_{ik}^2\right]. \tag{B.10}$$

Conditional on $(X_1,\ldots,X_n)$, η_{ik} has mean zero and every summand in η_{ik} lies in the interval, $\left[-\frac{M}{\sqrt{nh^{d_x}}}1\left\{\left(X_j-X_i\right)\in[-h,h]^{d_x}\right\}, \frac{M}{\sqrt{nh^{d_x}}}1\left\{\left(X_j-X_i\right)\in[-h,h]^{d_x}\right\}\right]$. The Hoeffding's inequality then implies that, for every $1\leq i\leq n$ and $1\leq k\leq \dim(U)$, it holds

$$\begin{aligned}&P^n\left(|\eta_{ik}|\geq t|X_1,\ldots,X_n\right) \\ \leq\ & 2\exp\left(-\frac{t^2}{\frac{2M^2}{nh^{d_x}}\sum_{j\in J_i}1\left\{\left(X_j-X_i\right)\in[-h,h]^{d_x}\right\}}\right) \\ \leq\ & 2\exp\left(-\frac{t^2}{\frac{2M^2}{nh^{d_x}}\max_{1\leq i\leq n}\sum_{j\in J_i}1\left\{\left(X_j-X_i\right)\in[-h,h]^{d_x}\right\}}\right),\ \forall t>0.\end{aligned}$$

Therefore,

$$
\begin{aligned}
& E_{P^n}\left[\exp\left(\frac{\eta_{ik}^2}{\frac{2M^2}{nh^{d_x}}\max_{1\leq i\leq n}\sum_{j\in J_i}1\left\{(X_j-X_i)\in[-h,h]^{d_x}\right\}}\right)|X_1,\ldots,X_n\right] \\
= \; & 1+\int_1^\infty P^n\left(\exp\left(\frac{\eta_{ik}^2}{\frac{2M^2}{nh^{d_x}}\max\limits_{1\leq i\leq n}\sum\limits_{j\in J_i}1\left\{(X_j-X_i)\in[-h,h]^{d_x}\right\}}\right)\geq t'|X_1,\ldots,X_n\right)dt' \\
= \; & 1+\int_1^\infty P^n\left(|\eta_{ik}|\geq\sqrt{\frac{2M^2}{nh^{d_x}}\max_{1\leq i\leq n}\sum_{j\in J_i}1\left\{(X_j-X_i)\in[-h,h]^{d_x}\right\}\log t'}|X_1,\ldots,X_n\right)dt' \\
\leq \; & 1+2\int_1^\infty \exp\left(-2\log t'\right)dt' \\
= \; & 1+2\int_1^\infty \left(t'\right)^{-2}dt' \\
= \; & 3
\end{aligned}
$$

for all $1\leq i\leq n$ and $1\leq k\leq \dim(U)$. We can therefore apply Lemma 1.6 of Tsybakov (2009) to bound $E_{P^n}\left[\max_{i,k}\eta_{ik}^2|X_1,\ldots,X_n\right]$,

$$
\begin{aligned}
& E_{P^n}\left[\max_{1\leq i\leq n,1\leq k\leq\dim(U)}\eta_{ik}^2|X_1,\ldots,X_n\right] \\
\leq \; & 2M^2\max_{1\leq i\leq n}\left[\frac{1}{nh^{d_x}}\sum_{j\in J_i}1\left\{(X_j-X_i)\in[-h,h]^{d_x}\right\}\right]\log\left(3\dim(U)\,n\right) \\
\leq \; & 2M^2\left[\frac{1}{h^{d_x}}\sup_{B\in\mathcal{B}_h}\left(P_{X,n}(B)-P_X(B)\right)+2^{d_x}\bar{p}_X\right]\log\left(3\dim(U)\,n\right).
\end{aligned}
$$

By applying Lemma A.5 with $\bar{F}=1$ and $\delta=\bar{p}_X(2h)^{d_x/2}$, the unconditional expectation of $\max_{i,k}\eta_{ik}^2$ can be bounded as

$$
E_{P^n}\left[\max_{1\leq i\leq n,1\leq k\leq\dim(U)}\eta_{ik}^2\right]\leq 2M^2\left[C_2 2^{d_x/2}\bar{p}_X\sqrt{\frac{v_{\mathcal{B}_h}}{nh^{d_x}}}+2^{d_x}\bar{p}_X\right]\log\left(3\dim(U)\,n\right) \tag{B.11}
$$

for all n such that $nh^{d_x}\geq\frac{C_1 v_{\mathcal{B}_h}}{2^{d_x}\bar{p}_X^2}$. Plugging (B.11) back into (B.10) and focusing on the leading term give

$$
\limsup_{n\to\infty}\sup_{P\in\mathcal{P}_\mu}E_{P^n}\left[\max_{0\leq i\leq n}\left|\sum\nolimits_{j\neq i}\xi_j\omega_j(X_i)\right|^2\cdot 1\{\Omega_{\lambda,n}\}\right]\leq O\left(\frac{\log n}{nh^{d_x}t_n^2}\right). \tag{B.12}
$$

Combining (B.8), (B.9), and (B.12), we obtain

$$\begin{aligned}
& E_{P^n}\left[\max_{1\leq i\leq n}\left|\hat{\mu}_{-i}(X_i)-\mu\left(X_i\right)\right|^2\right] \\
\leq\ & E_{P^n}\left[\max_{1\leq i\leq n}\left|\hat{\mu}_{-i}(X_i)-\mu\left(X_i\right)\right|^2\cdot 1\left\{\Omega_{\lambda,n}\right\}\right]+M^2P^n\left(\Omega_{\lambda,n}^c\right) \\
\leq\ & 2E_{P^n}\left[\max_{1\leq i\leq n}\left|\sum\nolimits_{j\neq i}\left(\mu\left(X_j\right)-\mu\left(X_i\right)\right)\omega_j\left(X_i\right)\right|^2\cdot 1\left\{\Omega_{\lambda,n}\right\}\right] \\
& +2E_{P^n}\left[\max_{1\leq i\leq n}\left|\sum\nolimits_{j\neq i}\xi_j\omega_j\left(X_i\right)\right|^2\cdot 1\left\{\Omega_{\lambda,n}\right\}\right]+M^2P^n\left(\Omega_{\lambda,n}^c\right), \\
=\ & O\left(\frac{h^{2\beta}}{t_n^2}\right)+O\left(\frac{\log n}{nh^{d_x}t_n^2}\right)+M^2P^n\left(\Omega_{\lambda,n}^c\right),
\end{aligned}$$

so the desired conclusion is proven if $P^n\left(\Omega_{\lambda,n}^c\right)$ is shown to converge faster than the $O\left(\frac{\log n}{nh^{d_x}t_n^2}\right)$ term.

To find the convergence rate of $P^n\left(\Omega_{\lambda,n}^c\right)$, consider first the case of non-stochastic J_i. By applying Lemma B.1 (iii) with the sample size set at $(n-1)$, we have

$$\begin{aligned}
P^n\left(\left\{\lambda\left(X_i\right)\leq c_6,\ \text{for some } 1\leq i\leq n\right\}\right) &= nP^n\left(\left\{\lambda\left(X_n\right)\leq c_6\right\}\right) \\
&= n\int P^n\left(\lambda(X_n)\leq c_6|X_n\right)dP_X \\
&= n\int P^{n-1}\left(\lambda(x)\leq c_6\right)dP_X(x) \qquad \text{(B.13)} \\
&\leq 2n\left[\dim U\right]^2\exp\left(-\frac{c_7}{2}nh^{d_x}\right).
\end{aligned}$$

For the case of stochastic J_i, by viewing n_{J_i} as a binomial random variable with parameters $(n-1)$ and π with $\kappa<\pi<1-\kappa$, and recalling that, when P_X satisfies Assumption 2.3 (PX), the conditional distributions $P_{X|D=d}$, $d\in\{1,0\}$ also satisfy the support and density conditions stated in Assumption 2.3 (PX), we can apply the exponential inequality shown in Lemma B.1 (iii) to bound $P^{n-1}\left(\lambda(x)\leq c_6|n_{J_n}\right)$. Hence, with $\Omega_{\pi,n}\equiv\left\{\left|\frac{n_{J_n}}{n-1}-\pi\right|\leq\frac{1}{2}\pi\right\}=\left\{\frac{(n-1)\pi}{2}\leq n_{J_n}\leq\frac{3(n-1)\pi}{2}\right\}$ used above, we have

$$\begin{aligned}
P^{n-1}\left(\lambda(x)\leq c_6\right) &\leq P^{n-1}\left(\left\{\lambda(x)\leq c_6\right\}\cap\Omega_{\pi,n}\right)+P^{n-1}\left(\Omega_{\pi,n}^c\right) \\
&\leq \max_{n_{J_n}\in\Omega_{\pi,n}}P^{n-1}\left(\lambda(x)\leq c_6|n_{J_n}\right)+P^{n-1}\left(\Omega_{\pi,n}^c\right). \\
&\leq 2\left[\dim U\right]^2\exp\left(-\frac{c_7\pi}{4}nh^{d_x}\right)+2\exp\left(-\frac{\pi^2}{4}n\right),
\end{aligned}$$

Plugging this upper bound into (B.13) and focusing on the leading term leads to

$$P^n\left(\left\{\lambda\left(X_i\right)\leq c_6,\ \text{for some } 1\leq i\leq n\right\}\right)\leq O\left(n\exp\left(-c_7\frac{\pi}{4}nh^{d_x}\right)\right).$$

Hence, in either of the non-stochastic or the stochastic J_i case, since $t_n \leq c_6$ holds for all large n and the obtained upper bounds are uniform over $P \in \mathcal{P}_\mu$, we conclude

$$\limsup_{n\to\infty} \sup_{P\in\mathcal{P}_\mu} E_{P^n}\left[\max_{1\leq i\leq n} \left|\hat{\mu}_{-i}(X_i) - \mu\left(X_i\right)\right|^2\right] \leq O\left(\frac{h^{2\beta}}{t_n^2}\right) + O\left(\frac{\log n}{nh^{d_x}t_n^2}\right) + O\left(n\exp(-nh^{d_x})\right).$$

Since $t_n = (\log n)^{-1}$ by assumption, $O(n\exp\left(-nh^{d_x}\right))$ converges faster than $O\left(\frac{\log n}{nh^{d_x}t_n^2}\right)$, the leading terms are given by the first two terms, $O\left(\frac{h^{2\beta}}{t_n^2}\right) + O\left(\frac{\log n}{nh^{d_x}t_n^2}\right)$. □

Proof of Corollary 2.1. By noting the following inequalities,

$$\begin{aligned}
E_{P^n}\left[\frac{1}{n}\sum\nolimits_{i=1}^n \left|\hat{\tau}^m(X_i) - \tau\left(X_i\right)\right|\right] &\leq E_{P^n}\left[\frac{1}{n}\sum\nolimits_{i=1}^n \left|\hat{m}_1(X_i) - m_1\left(X_i\right)\right|\right] \\
&\quad + E_{P^n}\left[\frac{1}{n}\sum\nolimits_{i=1}^n \left|\hat{m}_0(X_i) - m_0\left(X_i\right)\right|\right] \\
E_{P^n}\left[\max_{1\leq i\leq n}\left(\hat{\tau}^m(X_i) - \tau\left(X_i\right)\right)^2\right] &\leq 2E_{P^n}\left[\max_{1\leq i\leq n}\left(\hat{m}_1(X_i) - m_1\left(X_i\right)\right)^2\right] \\
&\quad + 2E_{P^n}\left[\max_{1\leq i\leq n}\left(\hat{m}_0(X_i) - m_0\left(X_i\right)\right)^2\right],
\end{aligned}$$

we obtain the current corollary by applying Lemma B.4. The resulting uniform convergence rate is given by $\psi_n = n^{\frac{1}{2+d_x/\beta_m}}$. When the assumption (2.9) in Theorem 2.6 is concerned, the corresponding rate is given by $\tilde{\psi}_n = \left[\left(\frac{\log n}{n}\right)^{\frac{1}{2+d_x/\beta_m}}(\log n)^2\right]^{-1}$. □

Proof of Corollary 2.2. (i) Assume that n is large enough so that $\varepsilon_n \leq \kappa/2$ holds. Given $\hat{e}\left(X_i\right) \in [\varepsilon_n, 1-\varepsilon_n]$, $\hat{\tau}_i^e - \tau_i$ can be expressed as

$$\hat{\tau}_i^e - \tau_i = \frac{Y_i D_i}{e(X_i)}\left[\frac{e(X_i) - \hat{e}(X_i)}{\hat{e}(X_i)}\right] + \frac{Y_i\left(1-D_i\right)}{1-e(X_i)}\left[\frac{e(X_i) - \hat{e}(X_i)}{1-\hat{e}\left(X_i\right)}\right],$$

so

$$|\hat{\tau}_i^e - \tau_i| \leq \frac{M}{\kappa}\cdot\frac{1}{\hat{e}\left(X_i\right)\left(1-\hat{e}\left(X_i\right)\right)}\cdot|\hat{e}(X_i) - e(X_i)|$$

holds. On the other hand, when $\hat{e}\left(X_i\right) \notin [\varepsilon_n, 1-\varepsilon_n]$, $\hat{\tau}_i^e = 0$ and $|\tau_i| \leq \frac{M}{\kappa}$ imply $|\hat{\tau}_i^e - \tau_i| \leq \frac{M}{\kappa}$. Hence, the following bounds are valid,

$$|\hat{\tau}_i^e - \tau_i| \leq \begin{cases} \frac{M}{\kappa}\cdot\frac{2}{\kappa(2-\kappa)}\cdot|\hat{e}(X_i) - e(X_i)| & \text{if } \hat{e}\left(X_i\right) \in \left[\frac{\kappa}{2}, 1-\frac{\kappa}{2}\right], \\ \frac{M}{\kappa}\cdot\frac{1}{\varepsilon_n(1-\varepsilon_n)} & \text{if } \hat{e}\left(X_i\right) \notin \left[\frac{\kappa}{2}, 1-\frac{\kappa}{2}\right]. \end{cases} \tag{B.14}$$

Hence,

$$\begin{aligned} E_{P^n}\left[\frac{1}{n}\sum\nolimits_{i=1}^{n}|\hat{\tau}_i^e-\tau_i|\right] &= E_{P^n}\left[|\hat{\tau}_n^e-\tau_n|\right] \\ &\leq \frac{M}{\kappa}\cdot\frac{2}{\kappa\left(2-\kappa\right)}\cdot E_{P^n}\left[|\hat{e}(X_n)-e(X_n)|\right] \\ &\quad +\frac{M}{\kappa}\cdot\frac{1}{\varepsilon_n\left(1-\varepsilon_n\right)}\cdot P^n\left(\hat{e}\left(X_n\right)\notin\left[\frac{\kappa}{2},1-\frac{\kappa}{2}\right]\right). \end{aligned}$$

By Lemma B.4 (i), $\sup_{P\in\mathcal{P}_e}E_{P^n}\left[|\hat{e}(X_n)-e(X_n)|\right]\leq O(n^{-\frac{1}{2+d_x/\beta_e}})$, so the conclusion follows if $P^n\left(\hat{e}\left(X_n\right)\notin\left[\frac{\kappa}{2},1-\frac{\kappa}{2}\right]\right)$ is shown to converge faster than $O(n^{-\frac{1}{2+d_x/\beta_e}})$. To see this claim is true, note that

$$\begin{aligned} P^n\left(\hat{e}\left(X_n\right)\notin\left[\frac{\kappa}{2},1-\frac{\kappa}{2}\right]\right) &= \int_{\mathcal{X}}P^{n-1}\left(\hat{e}\left(x\right)\notin\left[\frac{\kappa}{2},1-\frac{\kappa}{2}\right]\right)dP_X\left(x\right) \\ &\leq \int_{\mathcal{X}}P^{n-1}\left(\left|\hat{e}\left(x\right)-e(x)\right|\geq\frac{\kappa}{2}\right)dP_X\left(x\right) \\ &\leq c_9\exp\left(-\frac{c_{10}\kappa^2}{4}nh^{d_x}\right) \end{aligned}$$

holds for all n satisfying $c_8h^\beta<\kappa/2$, where the c_8, c_9, and c_{10} are the constants defined in Lemma B.2 (i). Since ε_n is assumed to converge at a polynomial rate, $\frac{1}{\varepsilon_n(1-\varepsilon_n)}P^n\left(\hat{e}\left(X_n\right)\notin\left[\frac{\kappa}{2},1-\frac{\kappa}{2}\right]\right)$ converges faster than $O(n^{-\frac{1}{2+d_x/\beta_e}})$.

(ii) By (B.14), we have

$$\begin{aligned} E_{P^n}\left[\max_{1\leq i\leq n}|\hat{\tau}_i^e-\tau_i|^2\right] &\leq \left(\frac{2M}{\kappa^2\left(2-\kappa\right)}\right)^2E_{P^n}\left[\max_{1\leq i\leq n}|\hat{e}(X_i)-e(X_i)|^2\right] \\ &\quad +\left(\frac{M}{\kappa\varepsilon_n\left(1-\varepsilon_n\right)}\right)^2P^n\left(\hat{e}\left(X_i\right)\notin\left[\frac{\kappa}{2},1-\frac{\kappa}{2}\right]\text{ for some }1\leq i\leq n\right). \end{aligned} \tag{B.15}$$

By Lemma B.4 (ii), the first term in (B.15) converges at rate $O\left(n^{-\frac{2}{2+d_x/\beta}}(\log n)^{\frac{2}{2+d_x/\beta}+2}\right)$. To find the convergence rate of the second term in (B.15), consider

$$\begin{aligned} & P^n\left(\hat{e}\left(X_i\right)\notin\left[\frac{\kappa}{2},1-\frac{\kappa}{2}\right]\text{ for some }1\leq i\leq n\right) \\ \leq\ & nP^n\left(\hat{e}\left(X_n\right)\notin\left[\frac{\kappa}{2},1-\frac{\kappa}{2}\right]\right) \\ \leq\ & c_9n\exp\left(-\frac{c_{10}\kappa^2}{4}nh^{d_x}\right), \end{aligned}$$

where the last line follows from Lemma B.2 (i). Since ε_n converges at polynomial rate, we conclude the second term in (B.15) converges faster than the first term. □

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, 91–117.

Abadie, A., M. M. Chingos, and M. R. West (2014): "Endogenous Stratification in Randomized Experiments," *unpublished manuscript*.

Armstrong, T. B. and S. Shen (2014): "Inference on Optimal Treatment Assignments," *unpublished manuscript*.

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.

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.

Boucheron, S., G. Lugosi, and P. Massart (2013): *Concentration Inequalities, A Nonasymptotic Theory of Independence*, Oxford University Press.

Bousquet, O. (2002): "A Bennet Concentration Inequality and its Application to Suprema of Empirical Processes," *Comptes Rendus de l'Académie des Sciences - Series I*, 334, 495–500.

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.

Dehejia (2005): "Program Evaluation as a Decision Problem," *Journal of Econometrics*, 125, 141–173.

Devroye, L., L. Györfi, and G. Lugosi (1996): *A Probabilistic Theory of Pattern Recognition*, Springer.

Devroye, L. and G. Lugosi (1995): "Lower Bounds in Pattern Recognition and Learning," *Pattern Recognition*, 28, 1011–1018.

Dudley, R. (1999): *Uniform Central Limit Theorems*, Cambridge University Press.

Elliott, G. and R. P. Lieli (2013): "Predicting Binary Outcomes," *Journal of Econometrics*, 174, 15 – 26.

Florios, K. and S. Skouras (2008): "Exact Computation of Max Weighted Score Estimators," *Journal of Econometrics*, 146, 86 – 91.

Heckman, J. J., H. Ichimura, and P. E. Todd (1997): "Matching As An Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme," *The Review of Economic Studies*, 64, 605–654.

Hirano, K. and J. R. Porter (2009): "Asymptotics for Statistical Treatment Rules," *Econometrica*, 77, 1683–1701.

Kanaya, S. and D. Kristensen (2014): "Optimal Sampling and Bandwidth Selection for Kernel Estimators of Diffusion Processes," *unpublished manuscript*.

Kasy, M. (2014): "Using Data to Inform Policy," *unpublished manuscript*.

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.

Lieli, R. P. and H. White (2010): "The Construction of Empirical Credit Scoring Rules Based on Maximization Principles," *Journal of Econometrics*, 157, 110 – 119.

Lugosi, G. (2002): "Pattern Classification and Learning Theory," in *Principles of Nonparametric Learning*, ed. by L. Györfi, Vienna: Springer, 1–56.

Mammen, E. and A. B. Tsybakov (1999): "Smooth Discrimination Analysis," *Annals of Statistics*, 27, 1808–1829.

Manski, C. F. (1975): "Maximum Score Estimation of the Stochastic Utility Model of Choice," *Journal of Econometrics*, 3, 205 – 228.

——— (2004): "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 72, 1221–1246.

Manski, C. F. and T. Thompson (1989): "Estimation of Best Predictors of Binary Response," *Journal of Econometrics*, 40, 97 – 123.

Massart, P. and E. Nédélec (2006): "Risk Bounds for Statistical Learning," *The Annals of Statistics*, 34, 2326–2366.

Qian, M. and S. A. Murphy (2011): "Performance Guarantees for Individualized Treatment Rules," *The Annals of Statistics*, 39, 1180–1210.

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.

Tsybakov, A. B. (2004): "Optimal Aggregation of Classifiers in Statistical Learning," *The Annals of Statistics*, 32, 135–166.

——— (2009): *Introduction to Nonparametric Estimation*, Springer.

van der Vaart, A. W. and J. A. Wellner (1996): *Weak Convergence and Empirical Processes*, Springer.

Vapnik, V. N. (1998): *Statistical Learning Theory*, John Wiley & Sons.

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, 1106–1118.