Statistica Sinica **31** (2021), 1353-1374
doi:https://doi.org/10.5705/ss.202018.0451

A SEQUENTIAL SIGNIFICANCE TEST FOR TREATMENT BY COVARIATE INTERACTIONS

Min Qian[1], Bibhas Chakraborty[2,3], Raju Maiti[2] and Ying Kuen Cheung[1]

[1]*Columbia University,* [2]*National University of Singapore and* [3]*Duke University*

Abstract: Biomedical and clinical research is gradually shifting from a traditional "one-size-fits-all" approach to a new paradigm of personalized medicine. An important step in this direction is to identify the treatment-covariate interactions. Our setting may include many covariates of interest. Numerous machine learning methodologies have been proposed to aid in treatment selection in this setting. However, few have adopted formal hypothesis testing procedures. As such, we present a novel testing procedure based on an m-out-of-n bootstrap that can be used to sequentially identify variables that interact with a treatment. We study the theoretical properties of the method, and use simulations to show that it outperforms competing methods in terms of controlling the type-I error rate and achieving satisfactory power. The usefulness of the proposed method is illustrated using real-data examples, from a randomized trial and an observational study.

Key words and phrases: Double robustness, forward stepwise testing, m-out-of-n bootstrap, non-regular asymptotics, personalized medicine.

1. Introduction

Owning to patients' heterogeneous responses to treatment programs, biomedical and clinical research is shifting from the traditional "one-size-fits-all" treatments to paradigms of personalized medicine. An important step in this direction is to identify the treatment-covariate interactions. In the conventional approach, investigators would first identify a set of key covariates. Then they would examine the treatment-covariate interactions either by comparing the treatment and control in subgroups defined by the key covariates, or by testing the regression coefficients of the treatment-covariate interaction terms in a multivariable linear model. However, this approaches are either infeasible or unreliable owning to overfitting in the case of a moderate to large number of covariates.

Several methodologies have been proposed to identify treatment-covariate interaction effects based on a large set of covariates, and thus optimize the treat-

Corresponding author: Min Qian, Department of Biostatistics, Columbia University, NYC, NY 10032, USA. E-mail: mq2158@cumc.columbia.edu.

ment selection. These include the ranking methods (Gunter, Zhu and Murphy (2011); Tian and Tibshirani (2011); Chen et al. (2017)), regression-based methods (Qian and Murphy (2011); Lu, Zhang and Zeng (2013); Tian et al. (2014); Fan, Lu and Song (2016)), weighted classification-type learning methods (Orellana, Rotnitzky and Robins (2010), Zhang et al. (2012); Zhao et al. (2012); Huang and Fong (2014); Liu et al. (2018)), tree-based methods (Su et al. (2008); Laber and Zhao (2015); Tsai et al. (2016)), and functional data approaches (McKeague and Qian (2014); Ciarleglio et al. (2015); Laber and Staicu (2018)), among others. Nevertheless, very few studies investigate formal hypothesis testing procedures that take variable selection into account.

A few works have proposed novel hypothesis testing approaches that identify subgroups of enhanced treatment effects. Shen and He (2015) developed a likelihood-based test for the existence of a subgroup based on linear logistic-normal mixture models. Fan, Song and Lu (2017) proposed a method to test and identify a subgroup using change-point techniques. Wager and Athey (2018) investigated a forest-based method for treatment effect estimations and inferences. Shi, Song and Lu (2019) proposed a nonparametric test to assess the incremental value of a given set of new variables when deciding on an optimal treatment, conditional on an existing set of prescriptive variables. These methods test nonlinear treatment effects, and work well with a relatively small set of covariates. Although Shi, Song and Lu (2019) apply their method in a forward stepwise fashion and study its variable-selection properties, the derived p-value loses its interpretation when used with forward selection. In the presence of a large number of covariates, Shen and Cai (2016) proposed a kernel-based method to identify interactions between a treatment and a group of covariates; however, this method can be used only in a randomized trial setting. Zhao, Small and Ertefaie (2017) considered a Lasso-based inference approach to identify informative covariates under the condition that the propensity score and the main effect of the covariates on the outcome are both well estimated.

In a methodologically analogous setting, gene–environment interactions have been studied extensively in the field of genetics. The most common approach is to test the interaction of each genetic marker and an environmental exposure separately, and then to adjust for multiple comparisons. To improve the power and reduce the burden of such comparisons, several global tests have been proposed to assess the joint interaction effect between a marker set and an environmental variable (e.g., Lin et al. (2013); Marceau et al. (2015)). However, the validity of these methods relies on having an approximately correct model of the main effects of the markers and environmental factors.

In this study, we consider data from either randomized trials or observational studies. We aim to identify those covariates that interact with a treatment, from among a large set of candidate covariates, using a sequential testing procedure. First, a marginal screening test is used to detect whether any covariate interacts significantly with the treatment. If so, we test whether there are additional treatment-covariate interactions in a forward stepwise fashion. The procedure continues until the p-value exceeds a prespecified level of significance. Forward stepwise regressions have been studied extensively (Barron and Cohen (2008); Donoho and Stodden (2006); Wang (2009); Ing and Lai (2011)). However, most of these studies focus on studying the variable selection consistency properties, instead of hypothesis testing. In real applications, to enhance reproducibility, clinicians like to have inferential guarantees for the selected covariates, which is not provided by variable selection consistency theory alone. These types of selective inference problems have drawn much attention from statisticians, and various methods have been proposed in the prediction literature (e.g., Bühlmann (2013); Zhang and Zhang (2014); Lockhart et al. (2014); van de Geer et al. (2014); McKeague and Qian (2015); Ning and Liu (2017); Luedtke and van der Laan (2018)).

We propose calibrating our test statistic either by directly sampling from the null (if it is estimable), or by using an m-out-of-n bootstrap. The m-out-of-n bootstrap is a general tool for conducting valid statistical inferences for nonregular parameters (Shao (1994); Bickel, Gotze and van Zwet (1997)). It is the usual nonparametric bootstrap (Efron (1979)), except that the resample size m is of a smaller order than the original sample size n. With an appropriate choice of m, the m-out-of-n bootstrap acts as a smoothing operation on the empirical distribution of the data. Data-driven methods for choosing m in various contexts have been proposed by Hall, Horowitz and Jing (1995), Lee (1999), Bickel and Sakov (2008) and Chakraborty, Laber and Zhao (2013). In this study, we develop an adaptive choice of m. The proposed test is valid as long as the propensity score is known (or modeled correctly) or the main-effects model of the covariates on the outcome is specified correctly.

The remainder of the paper is organized as follows. In Section 2, we set up the framework and describe the initial marginal screening test used to identify the variable that most strongly interacts with the treatment in a randomized trial setting. In Section 3, we present the sequential test procedure. In Section 4, we extend our method to allow for double robustness, which is particularly useful in observational studies in which the propensity score model is unknown. In Section 5, we conduct simulations comparing the proposed methods with existing

competitors, and illustrate the methods using two data examples. We conclude with a discussion in Section 6. The proofs of the theorems and details of the simulations are presented in the online Supplementary Material.

2. Marginal Screening Test for Randomized Trials

2.1. Marginal regression

Suppose we are given the pre-treatment information, treatment assignments, and outcomes for n patients. Further suppose there are only two competing treatments $A \in \{0, 1\}$. Let $\boldsymbol{X} \in \mathbb{R}^p$ be the vector of pre-treatment variables, and Y be a scalar outcome. Let $q_0(\boldsymbol{x}) := P(A = 1|\boldsymbol{X} = \boldsymbol{x})$ be the propensity of receiving treatment 1 in the observed data as a function of the pre-treatment variables $\boldsymbol{X} = \boldsymbol{x}$. In this section and Section 3, we assume the data come from a randomized trial; thus, $q_0(\boldsymbol{x})$ is known.

We frame the problem in terms of the following model:

$$Y = h_0(\boldsymbol{X}) + (\alpha_0 + \boldsymbol{X}^\mathsf{T}\boldsymbol{\beta}_0)A + \epsilon, \tag{2.1}$$

where $\boldsymbol{\beta}_0 \in \mathbb{R}^p$, $h_0(\boldsymbol{X}) := E(Y|\boldsymbol{X}, A = 0)$, and the error term ϵ has mean zero, finite variance, and is uncorrelated with $A - q_0(\boldsymbol{X})$ and $(A - q_0(\boldsymbol{X}))\boldsymbol{X}$. The term $\alpha_0 + \boldsymbol{X}^\mathsf{T}\boldsymbol{\beta}_0$ models $T(\boldsymbol{X}) := E(Y|\boldsymbol{X}, A = 1) - E(Y|\boldsymbol{X}, A = 0)$, the causal treatment effect for patients with pre-treatment information $\boldsymbol{X}$; thus, $(\alpha_0 + \boldsymbol{X}^\mathsf{T}\boldsymbol{\beta}_0)A$ is the treatment-by-covariate interaction model. We propose a sequential testing procedure that identifies which covariates interact with a treatment.

As an initial step, we test whether a treatment-by-covariate interaction exists. That is, we want to test

$$H_0 : \boldsymbol{\beta}_0 = \mathbf{0} \text{ vs. } H_a : \boldsymbol{\beta}_0 \neq \mathbf{0}. \tag{2.2}$$

Our proposed method is based on fitting p working marginal regression models, and then conducting a single test on the marginal regression coefficient of the most informative predictor of the causal treatment effect $T(\boldsymbol{X})$. Specifically, we can write $E(Y|\boldsymbol{X}, A) = E(Y|\boldsymbol{X}) + T(\boldsymbol{X})W$, where $W := A - q_0(\boldsymbol{X})$ (see Robins (1994)). For $k = 1, \ldots, p$, consider the working model $T(\boldsymbol{X}) = \alpha_k + \theta_k X_k$. The kth marginal regression model aims to estimate

$$(\alpha_k, \theta_k) = \underset{(\alpha,\theta)}{\operatorname{argmin}} E\left\{\left[Y - E(Y|\boldsymbol{X}) - (\alpha + \theta X_k)W\right]^2\right\}. \tag{2.3}$$

In addition, the index of the most informative predictor of $T(\boldsymbol{X})$ is

$$k_0 = \underset{k}{\operatorname{argmin}}\, E\left\{\left[Y - E(Y|\boldsymbol{X}) - (\alpha_k + \theta_k X_k)W\right]^2\right\}.$$

By taking the first-order derivative of (2.3) with respect to (α, θ), and noting that $E(W|\boldsymbol{X}) = 0$, we have under model (2.1),

$$\theta_k = \frac{E\big[W(Y - E(Y|\boldsymbol{X}))X_k'\big]}{E(WX_k')^2} = \frac{\operatorname{Cov}(WX_k', W\boldsymbol{X}^T\boldsymbol{\beta}_0)}{E(WX_k')^2}$$
$$\text{and} \quad k_0 = \underset{k}{\operatorname{argmax}}\big|\operatorname{Corr}(WX_k', W\boldsymbol{X}^T\boldsymbol{\beta}_0)\big|,$$

where $X_k' = X_k - E(W^2X_k)/EW^2$. Assume that k_0 is unique when $\boldsymbol{\beta}_0 \neq \mathbf{0}$. We can verify that $\boldsymbol{\beta}_0 = \mathbf{0}$ if and only if $\theta_k = 0$, for all $k = 1, \ldots, p$. Thus, hypothesis (2.2) is equivalent to

$$H_0 : \theta_0 = 0 \text{ vs. } H_a : \theta_0 \neq 0, \tag{2.4}$$

where $\theta_0 = \theta_{k_0}$. Using the randomized trial data, we can estimate k_0 and θ_0 using

$$\hat{k}_n = \underset{k\in\{1,\ldots,p\}}{\operatorname{argmin}}\, \mathbb{P}_n\left\{\left[Y - \hat{\phi}_n(\boldsymbol{X}) - (\hat{\alpha}_k + \hat{\theta}_k X_k)W\right]^2\right\}$$
$$\text{and} \quad \hat{\theta}_n := \hat{\theta}_{\hat{k}_n} = \left\{\mathbb{P}_n\left[\left(W\hat{X}_{\hat{k}_n}'\right)^2\right]\right\}^{-1}\mathbb{P}_n\left[W(Y - \hat{\phi}_n(\boldsymbol{X}))\hat{X}_{\hat{k}_n}'\right],$$

respectively, where $\mathbb{P}_n$ denotes the sample average, $\hat{\phi}_n(\boldsymbol{X})$ is an estimate of $E(Y|\boldsymbol{X})$, $(\hat{\alpha}_k, \hat{\theta}_k) = \operatorname{argmin}_{(\alpha,\theta)}\{\mathbb{P}_n\big[Y - \hat{\phi}_n(\boldsymbol{X}) - (\alpha + \theta X_k)W\big]^2\}$, and $\hat{X}_k' = X_k - \mathbb{P}_n(W^2X_k)/\mathbb{P}_nW^2$.

Our first result gives the asymptotic distribution of $\hat{\theta}_n$.

Theorem 1. *Assume i) $EX_k^4 < \infty$, for $k = 1, \ldots, p$; ii) the error term ϵ in model* (2.1) *has mean zero, finite variance, and is uncorrelated with W and $W\boldsymbol{X}$; and iii) $\hat{\phi}_n(\boldsymbol{X})$ is estimated from a P-Donsker class of measurable functions, and there exists some fourth-moment integrable function $\tilde{\phi}(\boldsymbol{X})$, such that $E[\hat{\phi}_n(\boldsymbol{X}) - \tilde{\phi}(\boldsymbol{X})]^4 \xrightarrow{P} 0$, as $n \to \infty$. Suppose k_0 is unique when $\boldsymbol{\beta}_0 \neq \mathbf{0}$. Then, under model* (2.1),

$$n^{1/2}(\hat{\theta}_n - \theta_0) \xrightarrow{d} \frac{Z_{k_0}}{E[WX_{k_0}']^2}1_{\boldsymbol{\beta}_0\neq\mathbf{0}} + \frac{Z_K}{E[WX_K']^2}1_{\boldsymbol{\beta}_0=\mathbf{0}},$$

where $X_k' = X_k - E(W^2X_k)/EW^2$, $K = \operatorname{argmax}_{k=1,\ldots,p}Z_k^2/E(WX_k')^2$, and $(Z_1, \ldots, Z_p)^{\mathsf{T}}$ is a mean-zero normal random vector with covariance matrix Σ given by that of the random vector with components

$$WX_k'\left\{Y-\tilde{\phi}(\boldsymbol{X})-\frac{E[W(Y-\tilde{\phi}(\boldsymbol{X}))]}{EW^2}W-\frac{E[WX_k'(Y-\tilde{\phi}(\boldsymbol{X}))]}{E(WX_k')^2}WX_k'\right\},$$

for $k=1,\ldots,p$, *where* Σ *is assumed to exist.*

Remark 1. The uniqueness of k_0 is assumed to ensure that the parameter θ_0 is well defined under H_a. Note that for hypothesis testing purposes, the test is always calibrated under the null distribution, which does not rely on this condition. Nonetheless, this condition can be removed with a slight modification of hypothesis (2.4) and the test statistic. A modified version of Theorem 1 without the uniqueness condition is presented in Section S3 of the Supplementary Material.

Remark 2. Note that $\hat{\phi}_n(\boldsymbol{X})$ is an estimate of $E(Y|\boldsymbol{X})$, which can be obtained from the sample mean of Y, the (regularized) regression from a linear model, or other nonparametric methods based on a Donsker class of functions that guarantee the convergence of $\hat{\phi}_n(\boldsymbol{X})$ to some function $\tilde{\phi}(\boldsymbol{X})$. The result does not require $\hat{\phi}_n(\boldsymbol{X})$ to be a consistent estimate of $E(Y|\boldsymbol{X})$. However, a good estimate of $E(Y|\boldsymbol{X})$ may help reduce the variance of each Z_k, and thus $\hat{\theta}_n$. For example, assume $q_0(\boldsymbol{X})=1/2$. We can verify that

$$\begin{aligned}\mathrm{Var}(Z_k)=&\mathrm{Var}\left(WX_k'\left[E(Y|X)-\tilde{\phi}(\boldsymbol{X})\right]\right)+\mathrm{Var}\left(WX_k'\epsilon\right)\\&+\mathrm{Var}\left(W^2X_k'\left[(\boldsymbol{X}-E\boldsymbol{X})^{\mathsf{T}}\boldsymbol{\beta}_0-\theta_kX_k'\right]\right).\end{aligned}$$

Therefore, the variance of each Z_k is minimized when $\tilde{\phi}(\boldsymbol{X})=E(Y|\boldsymbol{X})$.

One way to test $\theta_0=0$ is to estimate the null distribution of $n^{1/2}\hat{\theta}_n$ by setting $\boldsymbol{\beta}_0=\mathbf{0}$ in Theorem 1, and then replacing $\tilde{\phi}(\boldsymbol{X})$ with $\hat{\phi}_n(\boldsymbol{X})$ and the expectation with the sample average. The p-value can be calculated by comparing the observed test statistic $n^{1/2}\hat{\theta}_n$ with the estimated null distribution. As an alternative, we introduce an m-out-of-n bootstrap method to estimate the asymptotic distribution of $\hat{\theta}_n$. This method can be extended easily to cases where the null distribution is difficult to estimate (see Section 4).

2.2. The m-out-of-n bootstrap

The m-out-of-n bootstrap is a general tool that remedies the bootstrap inconsistency due to nonsmoothness. When the resample size m is of a smaller order than n, the empirical distribution converges to the true distribution at a faster rate than the m-out-of-n bootstrap sample empirical distribution converges to the empirical distribution. Intuitively, this implies that the empirical distribution converges to the true distribution first and, thus, the bootstrap resamples

behave as if they were drawn from the true distribution.

For a selected resample size m, let $\hat{\theta}_m^*$ be the analog of $\hat{\theta}_n$ based on the bootstrap sample of size m. Theorem 2 shows the bootstrap consistency results. The p-value can be calculated by comparing $n^{1/2}\hat{\theta}_n$ with the distribution of $m^{1/2}(\hat{\theta}_m^* - \hat{\theta}_n)$.

Theorem 2. *Assume all conditions in Theorem* 1 *hold. Suppose* $m/n = O(1)1_{\boldsymbol{\beta}\neq\mathbf{0}} + o(1)1_{\boldsymbol{\beta}=\mathbf{0}}$*, and* $m \to \infty$ *as* $n \to \infty$*. Then,* $m^{1/2}(\hat{\theta}_m^* - \hat{\theta}_n)$ *converges to the same limiting distribution as* $n^{1/2}(\hat{\theta}_n - \theta_0)$ *conditionally (on the data), in probability.*

A key challenge when using an m-out-of-n bootstrap is the choice of m. Bickel and Sakov (2008) proposed using an adaptive choice of m to construct confidence intervals for the extrema, and proved that the chosen m satisfies the conditions in Theorem 2. In our simulation studies, we found that their approach, when applied to our setting, did not achieve sufficient power, even though it successfully controlled the type-I error rate at the nominal level. As an alternative, in the context of using Q-learning to estimate optimal dynamic treatment regimes, Chakraborty, Laber and Zhao (2013) developed a scheme for selecting m for inferences of stage-one regression parameters that adapts to the degree of nonregularity. However, the tuning parameter in the procedure is chosen using a double bootstrap, which is very time consuming in our setting.

We extend Bickel and Sakov's method for choosing m by adding a crude pre-testing step that improves the power without inflating the type-I error rate. In particular, we define

$$\hat{r} = 1\left\{|\sqrt{n}T_n| < \max\left(\sqrt{c\log n},\ \text{upper } \frac{\alpha}{2p}\text{-quantile of } N(0,1)\right)\right\}, \tag{2.5}$$

where $T_n = \hat{\theta}_n/\hat{\sigma}_n$ is the conventional t-statistic based on the selected covariate $X_{\hat{k}_n}$, α is the level of significance, and $c > 0$ is a tuning parameter. Because $\sqrt{n}T_n = O_p(1)$ under H_0 and $\sqrt{n}T_n = O_p(\sqrt{n})$ under H_a, $\sqrt{c\log n}$ on the right-hand side of (2.5) guarantees that $\hat{r} \xrightarrow{P} 1_{\theta_0=0}$. The second component within max on the right-hand side of (2.5) controls the type-I error rate for small samples (see McKeague and Qian (2015)). Note that $\sqrt{c\log n}$ is an n-term and the second term in (2.5) is a p-term. Intuitively, when p is large, we expect that it to play a more important role than n; thus, the p-term should dominate the n-term, and vice-versa. The tuning parameter c controls the balance of these two terms. We recommend using $c = 2$ because we find it works well in various simulation settings. If $\hat{r} = 0$, we consider this a crude rejection of H_0, and propose using the

regular n-out-of-n bootstrap to conduct a refined test. On the other hand, $\hat{r} = 1$ indicates that there might be some nonregularity; in this case we propose using Bickel and Sakov's method to choose m. The complete algorithm is given below.

1. Calculate $\hat{r}$ defined in (2.5). If $\hat{r} = 0$, then choose $\hat{m} = n$. Otherwise, continue with steps 2–4 to obtain Bickel and Sakov's estimate $\hat{m}^{BS}$.

2. Consider a sequence of m's of the form $m_j = \lceil d^j n \rceil$, for $j = 0, 1, 2, \ldots$ and $0 < d < 1$, where $\lceil x \rceil$ denotes the smallest integer $\geq x$, and d is a tuning parameter.

3. For a given data set (with estimate $\hat{\theta}_n$), and for all j, define the bootstrap empirical distribution function $R^B_{m_j}(t, \hat{\theta}_n) = \sum_{b=1}^{B} 1_{m_j^{1/2}(\hat{\theta}^{*,b}_{m_j} - \hat{\theta}_n) \leq t}/B$, where $\hat{\theta}^{*,b}_{m_j}$ is the m_j-out-of-n bootstrap version of the estimate $\hat{\theta}_n$ from the bth bootstrap sample, for $b = 1, \ldots, B$.

4. Following Bickel and Sakov (2008), set m as the minimizer of the sup-norm of the successive differences between the bootstrap empirical distribution functions: $\hat{m}^{BS} = \text{argmin}_{m_j} \sup_t |R^B_{m_j}(t, \hat{\theta}_n) - R^B_{m_{j+1}}(t, \hat{\theta}_n)|$.

5. Output $\hat{m} = (1 - \hat{r})n + \hat{r}\hat{m}^{BS}$.

The tuning parameter d in Step 2 can be viewed as a step size because $m_{j+1}/m_j \approx d$. Bickel and Sakov (2008) used $d = 0.75$ in their simulation study, and reported robustness to other values. In our setting, we find that our method is robust to a choice of $d \in [0.7, 0.9]$ (see Section S5.4 in the Supplementary Material). In the simulation, we use $d = 0.8$ for our analysis.

3. Conditional Sequential Test for Randomized Trials

In the previous section, we proposed a marginal screening test to detect whether any covariate interacts with the treatment. If the null in (2.4) is rejected, we select $\hat{k}_n$ as the most informative predictor of the causal treatment effect. In this section, we extend our test to detect additional treatment-by-covariate interactions.

The procedure is carried out in a forward stepwise fashion. At each step, let $J \subset \{1, \ldots, p\}$ denote the index set such that $\{X_j : j \in J\}$ have been identified in previous steps as having a significant interaction with the treatment. We aim to test whether there is any $X_k \in \{X_j : j \in J^C\}$ that interacts with the treatment. Specifically, we rewrite model (2.1) as

$$Y = h_0(\boldsymbol{X}) + (\alpha_0 + \boldsymbol{X}_J^{\mathsf{T}}\boldsymbol{\beta}_{0,J} + \boldsymbol{X}_{J^C}^{\mathsf{T}}\boldsymbol{\beta}_{0,J^C})A + \epsilon, \tag{3.1}$$

where $\boldsymbol{X}_J = \{X_j : j \in J\}$ and $\boldsymbol{X}_{J^C} = \{X_j : j \in J^C\}$. The goal here is to test $\boldsymbol{\beta}_{0,J^C} = \mathbf{0}$.

Note that the index set J includes previously selected covariates. Thus, the null hypothesis is actually a function of the observed data. This makes any sequential test beyond the initial step a conditional hypothesis test. Nonetheless, under the alternative hypothesis, the selected covariate is the most informative predictor at each step, with probability tending to one (see the proofs of Theorems 1 and 3). Thus, the index set J converges to a fixed set if the alternative hypotheses in the previous steps are true.

For each $k \in J^C$, let $U_k = X_k - \widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\gamma}_k$, where $\widetilde{\boldsymbol{X}}_J = (1, \boldsymbol{X}_J^{\mathsf{T}})^{\mathsf{T}}$ and $\boldsymbol{\gamma}_k = \operatorname{argmin}_{\boldsymbol{\gamma}}\{E[W(X_k - \widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\gamma})]^2\}$. That is, $\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\gamma}_k$ is the weighted projection of X_k on the space spanned by $\widetilde{\boldsymbol{X}}_J$. After algebraic simplification, we can reformulate model (3.1) as

$$Y = h_0'(\boldsymbol{X}) + (\alpha_0' + \boldsymbol{U}^{\mathsf{T}}\boldsymbol{\beta}_{0,J^C})A + \epsilon', \tag{3.2}$$

where $\boldsymbol{U} = \{U_k : k \in J^C\}$; at the same time,

$$h_0'(\boldsymbol{X}) = h_0(\boldsymbol{X}) + q_0(\boldsymbol{X})\left[\widetilde{\boldsymbol{X}}_J - \frac{E(W^2\widetilde{\boldsymbol{X}}_J)}{EW^2}\right]^{\mathsf{T}}\boldsymbol{\lambda},$$

$$\alpha_0' = \frac{E(W^2\widetilde{\boldsymbol{X}}_J^{\mathsf{T}})}{EW^2}\boldsymbol{\lambda}, \text{ and } \epsilon' = \epsilon + W\left[\widetilde{\boldsymbol{X}}_J - \frac{E(W^2\widetilde{\boldsymbol{X}}_J)}{EW^2}\right]^{\mathsf{T}}\boldsymbol{\lambda},$$

where $\boldsymbol{\lambda} = (\alpha_0, \boldsymbol{\beta}_{0,J}^{\mathsf{T}})^{\mathsf{T}} + \boldsymbol{\Gamma}\boldsymbol{\beta}_{0,J^C}$, with $\boldsymbol{\Gamma}$ being a parameter matrix with columns consisting of $\{\boldsymbol{\gamma}_k : k \in J^C\}$.

Note that (3.2) is similar in form to model (2.1), where $\boldsymbol{X}$ is replaced with $\boldsymbol{U}$. Thus, testing $\boldsymbol{\beta}_{0,J^C} = 0$ is equivalent to testing $H_0 : \theta_0' = 0$ vs. $H_a : \theta_0' \neq 0$, where

$$\theta_0' = \frac{E[W\big(U_{k_0'} - E(W^2U_{k_0'})/EW^2\big)(Y - E(Y|\boldsymbol{X}))]}{E\Big\{\Big[W\Big(U_{k_0'} - E(W^2U_{k_0'})/EW^2\Big)\Big]^2\Big\}} = \frac{\operatorname{Cov}(WU_{k_0'}, W\boldsymbol{U}^{\mathsf{T}}\boldsymbol{\beta}_{0,J^C})}{E\Big(W^2U_{k_0'}^2\Big)},$$

and $k_0' = \operatorname{argmax}_{k:k\in J^C}\big|\operatorname{Corr}(WU_k, W\boldsymbol{U}^{\mathsf{T}})\boldsymbol{\beta}_{0,J^C}\big|$.

Similarly to the approach described in Section 2, we can estimate θ_0' by $\hat{\theta}_n'$ using empirical quantities. The theorem below gives the asymptotic distribution of $\hat{\theta}_n'$ and establishes the bootstrap consistency.

Theorem 3. *Assume conditions i)–iii) in Theorem* 1 *hold. Suppose* $k_0' = \operatorname{argmax}_{k:k\in J^C}\big|\operatorname{Corr}(WU_k, W\boldsymbol{U}^{\mathsf{T}}\boldsymbol{\beta}_{0,J^C})\big|$ *is unique when* $\boldsymbol{\beta}_{0,J^C} \neq \mathbf{0}$. *Then, under*

model (3.1),

$$n^{1/2}(\hat{\theta}'_n - \theta'_0) \xrightarrow{d} \frac{Z_{k_0}}{E(WU_{k_0})^2} 1_{\boldsymbol{\beta}_{0,J^C}\neq \mathbf{0}} + \frac{Z_K}{E(WU_K)^2} 1_{\boldsymbol{\beta}_{0,J^C}=\mathbf{0}},$$

where $K = \text{argmax}_{k in J^C} Z_k^2/E(WU_K)^2$, *and* $\{Z_k : k \in J^C\}$ *is a mean-zero normal random vector with a covariance matrix* Σ *given by that of the random vector with components*

$$WU_k \Bigg\{ Y - \tilde{\phi}(\boldsymbol{X}) - \frac{E[WU_k(Y - \tilde{\phi}(\boldsymbol{X}))]}{E(WU_k)^2} WU_k$$
$$-W\widetilde{\boldsymbol{X}}_J^{\mathsf{T}} \left(EW^2 \widetilde{\boldsymbol{X}}_J \widetilde{\boldsymbol{X}}_J^{\mathsf{T}} \right)^{-1} E\left[W\widetilde{\boldsymbol{X}}_J \left(Y - \tilde{\phi}(\boldsymbol{X}) \right) \right] \Bigg\},$$

for $k \in J^C$, *where* Σ *is assumed to exist.*

Furthermore, Let $\hat{\theta}'^*_m$ *be the* m*-out-of-*n *bootstrap analog of* $\hat{\theta}'_n$. *Assume* $m/n = O(1)1_{\boldsymbol{\beta}_{0,J^C}\neq \mathbf{0}} + o(1)1_{\boldsymbol{\beta}_{0,J^C}=\mathbf{0}}$, *and* $m \to \infty$ *as* $n \to \infty$. *Then,* $m^{1/2}(\hat{\theta}'^*_m - \hat{\theta}'_n)$ *converges to the same limiting distribution as that of* $n^{1/2}(\hat{\theta}'_n - \theta'_0)$ *conditionally (on the data), in probability.*

The test of $\theta'_0 = 0$ can be conducted using either the sampling from the null procedure or the m-out-of-n bootstrap procedure, as described in Section 2.

4. Extension to Allow for Double Robustness

The methods presented in the previous sections are designed for scenarios in which the propensity score $q_0(\boldsymbol{X})$ is known. In this section, we extend the procedure to allow for *double robustness*, in the sense that as long as $q_0(\boldsymbol{X})$ or $h_0(\boldsymbol{X})$ is consistently estimated, a valid inferential procedure can be established. This is particularly useful in observational studies, where the propensity score $q_0(\boldsymbol{X})$ is usually unknown.

We start with model (3.1), where the goal is to test $\boldsymbol{\beta}_{0,J^C} = \mathbf{0}$ after $\{X_j : j \in J\}$ have been detected in previous steps. Note that the initial test of $\boldsymbol{\beta}_0 = \mathbf{0}$ is a special case, with $J = \emptyset$. Let $\hat{q}_n(\boldsymbol{X})$ and $\hat{h}_n(\boldsymbol{X})$ be the estimates of $q_0(\boldsymbol{X})$ and $h_0(\boldsymbol{X})$, respectively, based on the data, and let $\tilde{q}(\boldsymbol{X})$ and $\tilde{h}(\boldsymbol{X})$ be the limits of $\hat{q}_n(\boldsymbol{X})$ and $\hat{h}_n(\boldsymbol{X})$, respectively (see Appendix A for the conditions on $\tilde{q}(\boldsymbol{X})$ and $\tilde{h}(\boldsymbol{X})$). Denote $\widetilde{W} := A - \tilde{q}(\boldsymbol{X})$.

To ensure double robustness, parameter estimates are often obtained by solving estimating equations, a method known as the G-estimation for structural mean models (Robins (1989, 1994)). Past research in this area has focused on

fitting a full causal treatment effect model restricted to a small set of variables, and then studying the efficiency of the estimate. In this section, we apply the G-estimation marginally on each covariate, and conduct the test based on the selected most informative covariate. Specifically, for each $k \in J^C$, let $(\boldsymbol{\delta}_k, \psi_k)$ be the solution to

$$E\Big\{(\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}, X_k)^{\mathsf{T}}\widetilde{W}\big[Y - \tilde{h}(\boldsymbol{X}) - (\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\delta} + X_k\psi)A\big]\Big\} = 0. \tag{4.1}$$

This yields $\psi_k = [E(A\widetilde{W}L_k^2)]^{-1}E[\widetilde{W}(Y - \tilde{h}(\boldsymbol{X}))L_k]$, where $L_k = X_k - \widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\eta}_k$ with $\boldsymbol{\eta}_k = \operatorname{argmin}_{\boldsymbol{\eta}} E\big[A\widetilde{W}(X_k - \widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\eta})^2\big]$.

To identify the most informative predictor, we need to identify the optimization objective function corresponding to equation (4.1). Note that under model (3.1), when $\tilde{q}(\boldsymbol{X}) = q_0(\boldsymbol{X})$ or $\tilde{h}(\boldsymbol{X}) = h_0(\boldsymbol{X})$ a.s., the left-hand side of (4.1) is equivalent to $E\big\{(\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}, X_k)^{\mathsf{T}}A\widetilde{W}\big[\alpha_0 + \boldsymbol{X}^{\mathsf{T}}\boldsymbol{\beta}_0 - (\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\delta} + X_k\psi)\big]\big\}$, where the quantity inside the expectation can be viewed as a quasi-likelihood score function. Thus, the solution to (4.1) satisfies

$$(\boldsymbol{\delta}_k, \psi_k) = \operatorname*{argmin}_{(\boldsymbol{\delta},\psi)} E\big[A\widetilde{W}(\alpha_0 + \boldsymbol{X}^{\mathsf{T}}\boldsymbol{\beta}_0 - \widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\delta} - X_k\psi)^2\big].$$

Intuitively, $\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\delta}_k + X_k\psi_k$ can be viewed as the best weighted linear approximation of the causal treatment effect $T(\boldsymbol{X}) = \alpha_0 + \boldsymbol{X}^{\mathsf{T}}\boldsymbol{\beta}_0$ based on $(\widetilde{\boldsymbol{X}}_J, X_k)$. Thus, it is natural to define the most informative predictor in $\{X_j : j \in J^C\}$ as

$$k_0^o := \operatorname*{argmin}_{k:k\in J^C} E\big[A\widetilde{W}(\alpha_0 + \boldsymbol{X}^{\mathsf{T}}\boldsymbol{\beta}_0 - \widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\delta}_k - X_k\psi_k)^2\big] = \operatorname*{argmax}_{k:k\in J^C}\Big[E\big(A\widetilde{W}L_k^2\big)\psi_k^2\Big].$$

In addition, the hypothesis of interest is

$$H_0 : \psi_0 = 0 \text{ vs. } H_a : \psi_0 \neq 0, \tag{4.2}$$

where $\psi_0 := \psi_{k_0^o}$. We can estimate ψ_0 using $\hat{\psi}_n := \hat{\psi}_{\hat{k}_n^o}$, where $\hat{\psi}_k = \mathbb{P}_n[\widehat{W}(Y - \hat{h}_n(\boldsymbol{X}))\hat{L}_k]/\mathbb{P}_n(A\widehat{W}\hat{L}_k^2)$ and $\hat{k}_n^o = \operatorname{argmax}_{k\in J^C}[\hat{\psi}_k^2\mathbb{P}_n(A\widehat{W}\hat{L}_k^2)]$, with $\widehat{W} = A - \hat{q}_n(\boldsymbol{X})$, $\hat{L}_k = X_k - \widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\hat{\boldsymbol{\eta}}_k$, and $\hat{\boldsymbol{\eta}}_k = \operatorname{argmin}_{\boldsymbol{\eta}}\mathbb{P}_n[A\widehat{W}(X_k - \widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\boldsymbol{\eta})^2]$.

The asymptotic distribution of $\hat{\psi}_n$ depends on the limiting behavior of $\hat{q}_n(\boldsymbol{X})$ and $\hat{h}_n(\boldsymbol{X})$, and is difficult to estimate. Thus, the m-out-of-n bootstrap plays an important role in obtaining a valid inference for ψ_0. Let $\hat{\psi}_m^*$ denote the bootstrap analog of $\hat{\psi}_n$. Below, we give the asymptotic distribution of $\hat{\psi}_n$, and prove the bootstrap consistency. A complete list of assumptions is given in Appendix A.

Theorem 4. *Suppose Assumptions* (A1)–(A5) *in Appendix A hold. Then, under*

model (3.1),

$$n^{1/2}(\hat{\psi}_n - \psi_0) \xrightarrow{d} \frac{\tilde{Z}_{k_0^o}}{E(\widetilde{AW}L_{k_0^o}^2)} 1_{\boldsymbol{\beta}_{0,J^C} \neq \mathbf{0}} + \frac{\tilde{Z}_K}{E(\widetilde{AW}L_K^2)} 1_{\boldsymbol{\beta}_{0,J^C} = \mathbf{0}},$$

where $K = \text{argmax}_{k \in J^C} \left[\tilde{Z}_k^2 / E(\widetilde{AW}L_k^2)\right]$, *and* $\{\tilde{Z}_k : k \in J^C\}$ *is the normal random vector defined in* (A.1).

Furthermore, suppose Assumption (A6) *in Appendix* A *holds, and that* $m/n = O(1)1_{\boldsymbol{\beta}_{0,J^C} \neq \mathbf{0}} + o(1)1_{\boldsymbol{\beta}_{0,J^C} = \mathbf{0}}$ *and* $m \to \infty$ *as* $n \to \infty$. *Then, under model* (3.1), $m^{1/2}(\hat{\psi}_m^* - \hat{\psi}_n)$ *converges to the same limiting distribution as that of* $n^{1/2}(\hat{\psi}_n - \psi_0)$ *conditionally (on the data), in probability.*

Note that although the doubly robust method can be used for randomized trials, it may cause extra dispersion in the variance in our setting. This issue is discussed in Section S4 of the Supplementary Material and using simulations.

5. Numerical Studies

In this section, we examine the performance of the proposed sequential testing procedure using simulated data, and apply the approach to two real-data examples.

5.1. Simulations

Below, we briefly summarize the simulation studies. See Section S5 in the Supplementary Material for further detail.

In the randomized trial setting, we compare the proposed sampling from null (NULL), m-out-of-n bootstrap ($\hat{m}$-boot), and doubly robust m-out-of-n bootstrap ($\hat{m}$-boot-DR) procedures with four competing methods: the likelihood ratio test (LRT); multiple testing with a Bonferroni correction (BONF); the n-out-of-n bootstrap (n-boot); and the m-out-of-n bootstrap, with m chosen using Bickel and Sakov's method ($\hat{m}^{BS}$-boot). Three data-generating models are considered: i) a null model; ii) a model with one active interaction term; and iii) a model with two equally active interaction terms. The sequential testing procedure is carried out to evaluate the power (when there is at least one active predictor in the candidate set) or the type-I error rate (when there is no active predictor remaining in the candidate set) at each step. The two proposed methods for randomized trials (NULL and $\hat{m}$-boot) provide good control of type I error rate and good power in all cases. The $\hat{m}$-boot-DR, $\hat{m}^{BS}$-boot, and LRT methods are less powerful than the NULL and $\hat{m}$-boot methods, and n-boot fails to control

the type-I error rate. When the components of $\boldsymbol{X}$ are uncorrelated, BONF is as good as our proposed methods in terms of the type-I error rate control and power. However, when the components of $\boldsymbol{X}$ are highly correlated, BONF is less powerful for large p (see Tables S1 and S2 in the Supplementary Material).

In the observational study setting, we compare the $\hat{m}$-boot-DR method with the $\hat{m}^{BS}$-boot and n-boot methods. We consider four data-generating models: two with correctly specified $q_0(\boldsymbol{X})$ and misspecified $h_0(\boldsymbol{X})$, and two with misspecified $q_0(\boldsymbol{X})$ and correctly specified $h_0(\boldsymbol{X})$. A linear logistic regression model with adaptive Lasso is used to estimate the propensity score model $q_0(\boldsymbol{X})$, and a linear regression with adaptive Lasso is used to estimate the main effect $h_0(\boldsymbol{X})$. The proposed $\hat{m}$-boot method provides good control of the type-I error rate and good power in all cases. The $\hat{m}^{BS}$-boot method lacks power relative to the $\hat{m}$-boot method, and n-boot fails to control the type-I error rate (see Table S3 in the Supplementary Material).

We also compare our methods with the kernel-based method proposed by Shen and Cai (2016) (KM_l) and the GESAT method proposed by Lin et al. (2013) for the global test of no treatment-by-covariate interactions. Both KM_l and GESAT are designed to test for the integrated effect of all covariates, with KM_l based on randomized trials, and GESAT based on a correct specification of both the main effect of the covariates and the interaction effects. As a result, we expect KM_l and GESAT to be more powerful in the case of weak dense signals, whereas our method will perform better in the case of strong sparse signals. This is indeed the case (see Table S4 in the Supplementary Material). All methods perform better when the covariates are correlated. In addition, KM_l fails to control the type-I error rate in the observational studies, and GESAT fails when the main effect is misspecified (see Table S5 in the Supplementary Material).

5.2. Nefazodone-CBASP trial example

The Nefazodone-CBASP trial was conducted to compare the efficacy of three treatments for chronic depression (Keller et al. (2000)). In this trial, 681 patients were randomly assigned to 12 weeks of outpatient treatment with nefazodone, the cognitive behavioral analysis system of psychotherapy (CBASP), or a combination of the two treatments. Various assessments were taken throughout the study, among which the score on the 24-item Hamilton Rating Scale for Depression (HRSD) was the primary outcome. Low HRSD scores are desirable. The primary analysis showed that, on average, the combination treatment was significantly better than any single treatment (p-value < 0.001 for both comparisons), and there was no overall difference between the two single treatments.

Table 1. p-values of sequentially selected covariates in the Nefazodone CBASP trial example. (PD: Panic Disorder; PsyPD: Psychotherapy for Past Depression; GAD: Generalized Anxiety Disorder; MDD: Major Depression Disorder; OCD: Obsessive Compulsive Disorder; AD-NOS: Anxiety Disorder not otherwise specified.)

Steps	Covariates selected	NULL	$\hat{m}$-boot	n-boot	LRT	BONF
Combination vs CBASP						
1	Subthreshold PD	0.628	0.578	0.016	0.792	1.000
2	PsyPD	0.622	0.754	0.034	0.903	1.000
3	Alcohol abuse	0.764	0.302	0.050	0.966	1.000
4	Threshold GAD	0.983	0.946	0.832	0.990	1.000
5	moderate MDD	0.993	0.564	0.622	0.994	1.000
Combination vs Nefazodone						
1	Past Psychotherapy	0.096	0.088	0.000	0.147	0.086
2	Alcohol dependence	0.369	0.122	0.020	0.386	0.458
3	OCD	0.562	0.758	0.030	0.604	0.862
4	AD-NOS	0.907	0.296	0.122	0.774	1.000
5	Atypical MDD	0.925	0.564	0.140	0.852	1.000

In the current analysis, we compare the combination treatment to nefazodone alone and the combination to CBASP alone to determine whether there are any covariate-by-treatment interactions. The outcome Y is a reduction in the 24-item HRSD score from the baseline, and there are 50 baseline covariates. We consider $n = 656$ patients, for which the final 24-item HRSD were observed. In both comparisons, we estimate $\tilde{\phi}(\boldsymbol{X}) = E(Y|\boldsymbol{X})$ using ridge regressions, and carry out the sequential tests in five steps. The selected covariates and p-values from the various methods are presented in Table 1.

For the comparison of the combination treatment with CBASP alone ($n = 438$), the regular n-boot suggests three important covariates: subthreshold panic disorder ($p = 0.016$), psychotherapy for past depression ($p = 0.034$), and alcohol abuse ($p = 0.050$). However, this is not supported by any other methods. For the comparison of the combination treatment with Nefazodone alone ($n = 440$), the regular n-boot again identifies three covariates: psychotherapy for past depression ($p < 0.001$), alcohol dependence ($p = 0.020$), and obsessive compulsive disorder ($p = 0.030$). For the LRT, all p-values are greater than 0.1. Falling between the two, the NULL, $\hat{m}$-boot, and BONF methods show that only "psychotherapy for past depression" has a p-value less than 0.1, which indicates that this may be worth further investigation.

5.3. COPES and CODIACS example

The COPES and CODIACS studies were conducted to compare the stepped care approach to standard care for patients with post-ACS (acute coronary syndrome) depression (Davidson et al. (2010, 2013)). Each trial enrolled about 150 patients. In both trials, patients were randomly assigned to six months of stepped care or usual care. Stepped-care participants were assigned to psychotherapy and/or an antidepressant, based on their preferences and a team of clinicians' recommendations. Usual-care participants received psychotherapy and/or an antidepressant from their current physicians. Depressive symptom changes, assessed using the Beck Depression Inventory (BDI) score, was the primary outcome for the CODIACS and a secondary outcome for the COPES.

As an example, using the combined data from the above two studies, we consider treatment A as a patient having received psychotherapy more than half the time during the six-month study period. Note that although the original two studies were randomized trials, the "treatment received" variable A is observational in nature. There are 26 baseline covariates, including patient demographics, symptoms, and the severity of symptoms in different domains (cardiac, depression, etc.). The outcome is a reduction in the BDI at six months from the baseline. The sample size $n = 265$, for which the final BDI and treatment information are available.

For the propensity score model, the initial analysis shows that treatment A depends strongly on the treatment arm and patient preference. Thus, we estimate the propensity score using a logistic regression with the treatment preference and whether the patient was assigned to stepped care or usual care. The main effect $h_0(\boldsymbol{X}) = E(Y|\boldsymbol{X}, A = 0)$ is estimated using a ridge regression with 26 baseline covariates using patients in the $A = 0$ group. We conduct five-step sequential tests; the results from $\hat{m}$-boot and n-boot are presented in Table 2. The n-boot method identifies two important treatment-by-covariate interactions: NEMC role limitation emotional T-score ($p = 0.004$), and GRACE score ($p = 0.024$); our proposed $\hat{m}^{ad}$-boot method suggests that only the NEMC role limitation emotional T-score may be worth further investigation.

6. Conclusion

This study develops a novel inference procedure to sequentially identify treatment-by-covariate interactions based on data from randomized trials and observational studies. The proposed method guarantees rigorous control of the type-I error rate at each step, and has greater power than competing testing procedures.

Table 2. p-values of sequentially selected covariates in the COPES and CODIACS example.

Steps	Covariates selected	$\hat{m}$-boot-DR	n-boot
1	NEMC Role limitation emotional T-score	0.028	0.004
2	GRACE score	0.330	0.024
3	NEMC General health T-score	0.504	0.112
4	BDI$\geq$ 29	0.672	0.186
5	Prefer psychotherapy only	0.690	0.232

Although the derivation of the asymptotic results assumes a fixed dimension p, numerical studies show that the proposed test continues to work when p is large. A theoretical investigation of the diverging p case is a challenging and interesting topic for future research.

Another challenge is to provide theoretical support to effectively control the false positive rate over the whole sequence of the forward stepwise path. As discussed in Tibshirani et al. (2016), sequential testing is typically only validated at each step. G'Sell et al. (2016) proposed stopping rules for exact control of the ordered false discovery rate under the assumption that the null p-values are independent. It would be worthwhile extending their methods to our setting so that the stepwise guarantees can be converted into stopping rules with desired inferential properties.

We have considered the case of binary treatments. When there are more than two treatments (e.g., L arms), the interaction effect between a treatment and a covariate includes $L-1$ terms, and a test statistic can be built based on the sum of the squares of the $L-1$ standardized parameter estimates. This would be an interesting extension for future research.

We have restricted our attention to a single-stage treatment-by-covariate interaction problem. However, time-varying treatments are common, and are needed by, for example, individuals with a chronic disease who experience a waxing and waning course of illness. The goal then is to identify informative covariates of the treatment effect at each stage. Q-learning and A-learning extend regressions to a multi-stage setting (Murphy (2003, 2005); Moodie, Richardson and Stephens (2007); Schulte et al. (2014)). It is well known that the regression coefficients for variables at stages prior to the last are nonregular (Robins (2004); Moodie, Richardson and Stephens (2010); Chakraborty, Murphy and Strecher (2010); Chakraborty, Laber and Zhao (2013); Laber et al. (2014); Song et al. (2015)). In that case, the selection procedure adds another layer of nonregularity to this already nonregular problem. It would be interesting, albeit challenging,

to extend our approach to the multi-stage setting. We view this as an important future work.

Supplementary Material

The online Supplementary Material contains the proofs of the theorems, an extension of Theorem 1 to non-unique k_0, a discussion on using the doubly robust method in randomized trials, and details of the simulation studies.

Acknowledgments

Min Qian's research was supported in part by NIH grants R21MH108999, 2R01GM095722, and R01MH109496. Bibhas Chakraborty and Raju Maiti's research was supported in part by grant MOE2015-T2-2-056 from Singapore's Ministry of Education and a start-up grant from Duke-NUS Medical School. Ying-Kuen Cheung's research was supported in part by NIH grants R21MH108999 and R01MH109496.

A. Assumptions of Theorems 4

Theorem 4 requires the following assumptions.

(A1) $EX_k^4 < \infty$ for $k = 1, \ldots, p$.

(A2) There exist functions $\tilde{h}(\boldsymbol{X})$ and $\tilde{q}(\boldsymbol{X})$ such that $n^{1/2}[\hat{h}_n(\boldsymbol{x}) - \tilde{h}(\boldsymbol{x})] = \Delta_h(\boldsymbol{x})\hat{S}_h + o_P(1)$ and $n^{1/2}[\hat{q}_n(\boldsymbol{x}) - \tilde{q}(\boldsymbol{x})] = \Delta_q(\boldsymbol{x})\hat{S}_q + o_P(1)$, where $\Delta_h(\boldsymbol{x})$ and $\Delta_q(\boldsymbol{x})$ are vector-valued deterministic functions of $\boldsymbol{x}$, and $\hat{S}_h$ and $\hat{S}_q$ are data dependent random vectors satisfying

i). $\Delta_h(\boldsymbol{X})$ and $\Delta_q(\boldsymbol{X})$ are square integrable random vectors; and

ii).

$$\begin{aligned}&\Bigg(\Big\{\mathbb{G}_n\Big[\widetilde{W}L_k\Big(Y - \tilde{h}(\boldsymbol{X}) - \psi_k AL_k - E\Big[\widetilde{W}\Big(Y - \tilde{h}(\boldsymbol{X})\Big)\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\Big] \\ &\qquad \times \Big[E\Big(A\widetilde{W}\widetilde{\boldsymbol{X}}_J\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\Big)\Big]^{-1}A\widetilde{\boldsymbol{X}}_J\Big)\Big]\Big\}_{k\in J^C}, \hat{S}_h, \hat{S}_q\Bigg)^{\mathsf{T}} \\ &\xrightarrow{d} (\{Z_k^o : k \in J^C\}, S_h, S_q)^{\mathsf{T}} \sim N(0, \Sigma^o)\end{aligned}$$

for some variance-covariance matrix Σ^o assumed to exist.

(A3) The error term ϵ in model (2.1) has mean zero, finite variance, and is uncorrelated with $(\widetilde{W}, \widetilde{W}\boldsymbol{X})$, where $\widetilde{W} = A - \tilde{q}(\boldsymbol{X})$.

(A4) k_0^o is unique when $\boldsymbol{\beta}_{0,J^C} \neq 0$.

(A5) $\tilde{q}(\boldsymbol{X}) = q_0(\boldsymbol{X})$ or $\tilde{h}(\boldsymbol{X}) = h_0(\boldsymbol{X})$ a.s.

(A6) Let $\hat{q}_m^*(\boldsymbol{X})$ and $\hat{h}_m^*(\boldsymbol{X})$ be estimates of $q_0(\boldsymbol{X})$ and $h_0(\boldsymbol{X})$ based on the bootstrap sample of size m. Assume $m^{1/2}[\hat{h}_m^*(\boldsymbol{x}) - \hat{h}_n(\boldsymbol{x})] = \Delta_h(\boldsymbol{x})\hat{S}_h^* + o_{P_M}(1)$ and $m^{1/2}[\hat{q}_m^*(\boldsymbol{x}) - \hat{q}_n(\boldsymbol{x})] = \Delta_q(\boldsymbol{x})\hat{S}_q^* + o_{P_M}(1)$ conditionally on the data (in probability), where $\Delta_h(\boldsymbol{x})$ and $\Delta_q(\boldsymbol{x})$ are defined in Assumption (A2), and $\hat{S}_h^*$ and $\hat{S}_q^*$ are bootstrap sample dependent random vectors satisfying

$$\begin{aligned}&\Bigg(\Big\{\mathbb{G}_m^*\Big[\widetilde{W}L_k\Big(Y - \tilde{h}(\boldsymbol{X}) - \psi_k AL_k - E\Big[\widetilde{W}\Big(Y - \tilde{h}(\boldsymbol{X})\Big)\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\Big] \\ &\qquad \times \Big[E\Big(A\widetilde{W}\widetilde{\boldsymbol{X}}_J\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\Big)\Big]^{-1}A\widetilde{\boldsymbol{X}}_J\Big)\Big]\Big\}_{k\in J^C}, \hat{S}_h^*, \hat{S}_q^*\Bigg)^{\mathsf{T}} \\ &\quad \xrightarrow{d} (\{Z_k^o : k \in J^C\}, S_h, S_q)^{\mathsf{T}} \sim N(0, \Sigma^o) \text{ conditionally, in probability.}\end{aligned}$$

Remark 3. $\{\tilde{Z}_k : k \in J^C\}$ in Theorem 4 is defined as, for $k \in J^C$,

$$\begin{aligned}\tilde{Z}_k = Z_k^o - E\Big[\widetilde{W}L_k\Delta_h(\boldsymbol{X})\Big]S_h - E\Bigg\{L_k\Big(Y - \tilde{h}(\boldsymbol{X}) - \psi_k AL_k \\ - E\Big[\widetilde{W}\Big(Y - \tilde{h}(\boldsymbol{X})\Big)\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\Big]\Big[E\Big(A\widetilde{W}\widetilde{\boldsymbol{X}}_J\widetilde{\boldsymbol{X}}_J^{\mathsf{T}}\Big)\Big]^{-1}A\widetilde{\boldsymbol{X}}_J\Big)\Delta_q(\boldsymbol{X})\Bigg\}S_q \quad \text{(A.1)}\end{aligned}$$

Remark 4. Assumptions (A2) and (A6) require that the original sample and bootstrap sample estimates of $h_0(\boldsymbol{x})$ and $q_0(\boldsymbol{x})$ are well behaved. One can verify that, under appropriate conditions, the assumptions hold when $h_0(\boldsymbol{x})$ and $q_0(\boldsymbol{x})$ are estimated using linear/logistic regression, ridge regression, or variable selection methods with oracle properties (e.g. SCAD (Fan and Li (2001)), adaptive Lasso (Zou and Hastie (2005))).

References

Barron, A. and Cohen, A. (2008). Approximation and learning by greedy algorithms. *The Annals of Statistics* **36**, 64–94.

Bickel, P., Gotze, F. and van Zwet, W. (1997). Resampling fewer than n observations: Gains, losses and remedies for losses. *Statistica Sinica* **7**, 1–31.

Bickel, P. and Sakov, A. (2008). On the choice of m in the m out of n bootstrap and confidence bounds for extrema. *Statistica Sinica* **18**, 967–985.

Bühlmann, P. (2013). Statistical significance in high-dimensional linear models. *Bernoulli* **19**, 1212–1242.

Chakraborty, B., Laber, E. B. and Zhao, Y. (2013). Inference for optimal dynamic treatment regimes using an adaptive m-out-of-n bootstrap scheme. *Biometrics* **69**, 714–723.

Chakraborty, B., Murphy, S. and Strecher, V. (2010). Inference for nonregular parameters in optimal dynamic treatment regimes. *Statistical Methods in Medical Research* **19**, 317–343.

Chen, S., Tian, L., Cai, T. and Yu, M. (2017). A general statistical framework for subgroup identification and comparative treatment scoring. *Biometrics* **73**, 1199–1209.

Ciarleglio, A., Petkova, E., Ogden, R. and Tarpey, T. (2015). Treatment decisions based on scalar and functional baseline covariates. *Biometrics* **71**, 884–894.

Davidson, K., Bigger, J., Burg, M., Carney, R., Chaplin, W., Czajkowski, S. et al. (2013). Centralized, stepped, patient preference-based treatment for patients with post-acute coronary syndrome depression: CODIACS vanguard randomized controlled trial. *The Journal of the American Medical Association* **173**, 997–1004.

Davidson, K., Rieckmann, N., Clemow, L., Schwartz, J., Shimbo, D., Medina, V. et al. (2010). Enhanced depression care for patients with acute coronary syndrome and persistent depressive symptoms: coronary psychosocial evaluation studies randomized controlled trial. *Archives of Internal Medicine* **170**, 600–608.

Donoho, D. and Stodden, V. (2006). Breakdown point of model selection when the number of variables exceeds the number of observations. In *The 2006 IEEE International Joint Conference on Neural Network Proceedings*, 1916–1921.

Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *The Annals of Statistics* **7**, 1–26.

Fan, A., Lu, W. and Song, R. (2016). Sequential advantage selection for optimal treatment regime. *The Annals of Applied Statistics* **10**, 32–53.

Fan, A., Song, R. and Lu, W. (2017). Change-plane analysis for subgroup detection and sample size calculation. *Journal of the American Statistical Association* **112**, 769–778.

Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association* **96**, 1348–1360.

G'Sell, M. G., Wager, S., Chouldechova, A. and Tibshirani, R. (2016). Sequential selection procedures and false discovery rate control. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **78**, 423–444.

Gunter, L., Zhu, J. and Murphy, S. (2011). Variable selection for qualititative interactions. *Statistical Methodology* **8**, 42–55.

Hall, P., Horowitz, J. and Jing, B. (1995). On blocking rules for the bootstrap with dependent data. *Biometrika* **82**, 561–574.

Huang, Y. and Fong, Y. (2014). Identifying optimal biomarker combinations for treatment selection via a robust kernel method. *Biometrics* **70**, 891–901.

Ing, C.-K. and Lai, T. L. (2011). A stepwise regression method and consistent model selection for high-dimensional sparse linear models. *Statistica Sinica* **21**, 1473–1513.

Keller, M., McCullough, J., Klein, D., Arnow, B., Dunner, D., Gelenberg, A. et al. (2000). A comparison of nefazodone, the cognitive behavioral-analysis system of psychotherapy and their combination for the treatment of chronic depression. *The New England Journal of Medicine* **342**, 1462–1470.

Laber, E. B., Lizotte, D. J., Qian, M., Pelham, W. E. and Murphy, S. A. (2014). Dynamic treatment regimes: Technical challenges and applications. *Electronic Journal of Statistics* **8**, 1225.

Laber, E. and Staicu, A. (2018). Functional feature construction for individualized treatment regimes. *Journal of the American Statistical Association* **113**, 1219–1227.

Laber, E. and Zhao, Y. (2015). Tree-based methods for estimating individualized treatment

regimes. *Biometrika* **102**, 501–514.

Lee, S. (1999). On a class of m out of n bootstrap confidence intervals. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **61**, 901–911.

Lin, X., Lee, S., Christiani, D. C. and Lin, X. (2013). A test for the interaction between a genetic marker set and environment in generalized linear models. *Biostatistics* **14**, 667–681.

Liu, Y., Wang, Y., Kosorok, M., Zhao, Y. and Zeng, D. (2018). Augmented outcome-weighted learning for estimating optimal dynamic treatment regimens. *Statistics in Medicine* **37**, 3776–3788.

Lockhart, R., Taylor, J., Tibshirani, R. J. and Tibshirani, R. (2014). A significance test for the Lasso. *The Annals of Statistics* **42**, 413–468.

Lu, W., Zhang, H. H. and Zeng, D. (2013). Variable selection for optimal treatment decision. *Statistical Methods in Medical Research* **22**, 493–504.

Luedtke, A. and van der Laan, M. (2018). Parametric-rate inference for one-sided differentiable parameters. *Journal of the American Statistical Association* **113**, 780–788.

Marceau, R., Lu, W., Holloway, S., Sale, M., Worrall, B., Williams, S. et al. (2015). A fast multiple-kernel method with applications to detect gene-environment interaction. *Genetic Epidemiology* **39**, 456–468.

McKeague, I. W. and Qian, M. (2014). Estimation of treatment policies based on functional predictors. *Statistica Sinica* **24**, 1461–1485.

McKeague, I. W. and Qian, M. (2015). An adaptive resampling test for detecting the presence of significant predictors. *Journal of the American Statistical Association* **110**, 1422–1433.

Moodie, E., Richardson, T. and Stephens, D. (2007). Demystifying optimal dynamic treatment regimes. *Biometrics* **63**, 447–455.

Moodie, E., Richardson, T. and Stephens, D. (2010). Estimating optimal dynamic regimes: Correcting bias under the null. *Scandinavian Journal of Statistics* **37**, 126–146.

Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **65**, 331–366.

Murphy, S. A. (2005). A generalization error for Q-learning. *Journal of Machine Learning Research* **6**, 1073–1097.

Ning, Y. and Liu, H. (2017). A general theory of hypothesis tests and confidence regions for sparse high dimensional models. *Psychological Methods* **45**, 158–195.

Orellana, L., Rotnitzky, A. and Robins, J. (2010). Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part I: Main content. *The International Journal of Biostatistics* **6**, 1–49.

Qian, M. and Murphy, S. (2011). Performance guarantees for individualized treatment rules. *The Annals of Statistics* **39**, 1180–1210.

Robins, J. (1994). Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics - Theory and Methods* **23**, 2379–2412.

Robins, J. M. (1989). The analysis of randomised and non-randomised aids treatment trials using a new approach to causal inference in longitudinal studies. *Health Service Research Methodology: A Focus on AIDS* (Edited by L. Sechrest, H. Freeman and A. Mulley), 113–159. US Public Health Service, National Center for Health Services Research.

Robins, J. M. (2004). Optimal structural nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium in Biostatitics* (Edited by D. Lin and P. J. Heagerty), 189–326. Springer.

Schulte, P., Tsiatis, A., Laber, E. and Davidian, M. (2014). Q- and a-learning methods for

estimating optimal dynamic treatment regimes. *Statistical Science* **29**, 640–661.

Shao, J. (1994). Bootstrap sample size in nonregular cases. In *Proceedings of the American Mathematical Society* **122**, 1251–1262.

Shen, J. and He, X. (2015). Inference for subgroup analysis with a structured logistic-normal mixture model. *Journal of the American Statistical Association* **110**, 303–312.

Shen, Y. and Cai, T. (2016). Identifying predictive markers for personalized treatment selection. *Biometrics* **72**, 1017–1025.

Shi, C., Song, R. and Lu, W. (2019). On testing conditional qualitative treatment effects. *The Annals of Statistics* **47**, 2348–2377.

Song, R., Wang, W., Zeng, D. and Kosorok, M. (2015). Penalized Q-learning for dynamic treatment regimes. *Statistica Sinica* **25**, 901–920.

Su, X., Zhou, T., Yan, X., Fan, F. and Yang, S. (2008). Interaction trees with censored survival data. *The International Journal of Biostatistics* **4**, Article 2.

Tian, L., Alizadeh, A., Gentles, A. J. and Tibshirani, R. (2014). A simple method for estimating interactions between a treatment and a large number of covariates. *Journal of the American Statistical Association* **109**, 1517–1532.

Tian, L. and Tibshirani, R. (2011). Adaptive index models for marker-based risk stratification. *Biostatistics* **12**, 68–86.

Tibshirani, R. J., Taylor, J., Lockhart, R. and Tibshirani, R. (2016). Exact post-selection inference for sequential regression procedures. *Journal of the American Statistical Association* **111**, 600–620.

Tsai, W.-M., Zhang, H., Buta, E., O'Malley, S. and Gueorguieva, R. (2016). A modified classification tree method for personalized medicine decisions. *Statistics and Its Interface* **9**, 239–253.

van de Geer, S., Bühlmann, P., Ritov, Y. and Dezeure, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics* **42**, 1166–1202.

Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* **113**, 1228–1242.

Wang, H. (2009). Forward regression for ultra-high dimensional variable screening. *Journal of the American Statistical Association* **104**, 1512–1524.

Zhang, B., Tsiatis, A. A., Laber, E. B. and Davidian, M. (2012). A robust method for estimating optimal treatment regimes. *Biometrics* **68**, 1010–1018.

Zhang, C.-H. and Zhang, S. S. (2014) Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **76**, 217–242.

Zhao, Q., Small, D. S. and Ertefaie, A. (2017). Selective inference for effect modification via the Lasso. *arXiv preprint arXiv:1705.08020*.

Zhao, Y., Zeng, D., Rush, A. J. and Kosorok, M. R. (2012). Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association* **107**, 1106–1118.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **67**, 301–320.

Min Qian
Department of Biostatistics, Columbia University, 722 West 168th St., NYC, NY 10032, USA.
E-mail: mq2158@cumc.columbia.edu
Bibhas Chakraborty
Centre for Quantitative Medicine, Duke-NUS Medical School, Singapore 169857.
E-mail: bibhas.chakraborty@duke-nus.edu.sg
Raju Maiti
Centre for Quantitative Medicine, Duke-NUS Medical School, Singapore 169857.
E-mail: raju.maiti@duke-nus.edu.sg
Ying Kuen Cheung
Department of Biostatistics, Columbia University, 722 West 168th St., NYC, NY 10032, USA.
E-mail: yc632@cumc.columbia.edu

(Received November 2018; accepted October 2019)