

REGRESSION-ADJUSTED ESTIMATION OF QUANTILE TREATMENT EFFECTS
UNDER COVARIATE-ADAPTIVE RANDOMIZATIONS

By

Liang Jiang, Peter C.B. Phillips, Yubo Tao, and Yichong Zhang

May 2021

COWLES FOUNDATION DISCUSSION PAPER NO. 2288



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

Regression-Adjusted Estimation of Quantile Treatment Effects under Covariate-Adaptive Randomizations*

Liang Jiang[†] Peter C.B. Phillips[‡] Yubo Tao[§] Yichong Zhang[¶]

May 31, 2021

Abstract

This paper examines regression-adjusted estimation and inference of unconditional quantile treatment effects (QTEs) under covariate-adaptive randomizations (CARs). Datasets from field experiments usually contain extra baseline covariates in addition to the strata indicators. We propose to incorporate these extra covariates via auxiliary regressions in the estimation and inference of unconditional QTEs. We establish the consistency, limit distribution, and validity of the multiplier bootstrap of the QTE estimator under CARs. The auxiliary regression may be estimated parametrically, nonparametrically, or via regularization when the data are high-dimensional. Even when the auxiliary regression is misspecified, the proposed bootstrap inferential procedure still achieves the nominal rejection probability in the limit under the null. When the auxiliary regression is correctly specified, the regression-adjusted estimator achieves the minimum asymptotic variance. We also derive the optimal pseudo true values for the potentially misspecified parametric model that minimize the asymptotic variance of the corresponding QTE estimator. Our estimation and inferential methods can be implemented without tuning parameters and they allow for common choices of auxiliary regressions such as linear, probit and logit regressions despite the fact that these regressions may be misspecified. Finite-sample performance of the new estimation and inferential methods is assessed in simulations and an empirical application studying the impact of child health and nutrition on educational outcomes is included.

Keywords: Covariate-adaptive randomization, high-dimensional data, regression adjustment, quantile treatment effects.

JEL codes: C14, C21, I21

*Zhang acknowledges financial support from the Singapore Ministry of Education under Tier 2 grant No. MOE2018-T2-2-169 and a Lee Kong Chian fellowship. Phillips acknowledges support from NSF Grant No. SES 18-50860, a Kelly Fellowship at the University of Auckland, and a Lee Kong Chian Fellowship. Any errors are our own.

[†]Fudan University. E-mail address: jiangliang@fudan.edu.cn.

[‡]Yale University, University of Auckland, University of Southampton, and Singapore Management University. E-mail address: peter.phillips@yale.edu

[§]Singapore Management University. E-mail address: ybtao@smu.edu.sg. Corresponding author.

[¶]Singapore Management University. E-mail address: yczhang@smu.edu.sg.

1 Introduction

Covariate-adaptive randomizations (CARs) have recently seen growing use in a wide variety of randomized experiments in economic research. In CAR modeling, units are first stratified using some baseline covariates, and then, within each stratum, the treatment status is assigned (independent of covariates) to achieve the balance between the numbers of treated and control units. Researchers often use the randomized experiments to estimate quantile treatment effects (QTEs). The QTE has a useful role as a robustness check for the average treatment effect (ATE) and characterizes any heterogeneity that may be present in the sign and magnitude of the treatment effects according to their position within the distribution of outcomes. See, for example, Bitler, Gelbach, and Hoynes (2006), Muralidharan and Sundararaman (2011), Duflo, Greenstone, Pande, and Ryan (2013), Banerjee, Duflo, Glennerster, and Kinnan (2015), Crépon, Devoto, Duflo, and Parienté (2015), and Campos, Frese, Goldstein, Iacovone, Johnson, McKenzie, and Mensmann (2017).

Two practical issues arise in estimation and inference concerning QTEs under CARs. First, other covariates in addition to the strata indicators are collected during the experiment. In the estimation of ATE, the usual practice is to run a simple linear OLS regression of the outcome on treatment status, strata indicators, additional covariates, and interaction terms, as in the analysis of covariance (ANCOVA). However, because the quantile function is a nonlinear operator, even when the treatment status is completely randomly assigned, a similar linear quantile regression is unable to consistently estimate the *unconditional* QTE. Second, in order to achieve balance in the respective number of treated and control units within each stratum, treatment statuses under CARs usually exhibit a (negative) cross-sectional dependence. Standard inference procedures that rely on cross-sectional independence are therefore conservative and lack power. These two issues raise questions of how to use the additional covariates in the estimation of QTE and how to conduct valid statistical procedures that mitigate conservatism in inference.

The present paper addresses these issues by including additional covariates via auxiliary regressions, deriving the limit theory, and establishing the validity of multiplier bootstrap inference for the corresponding regression-adjusted QTE estimator under CARs. Even under potential misspecification of the auxiliary regressions, the QTE estimator is shown to maintain consistency and the multiplier bootstrap procedure to have asymptotic size equal to the nominal level under the null. When the auxiliary regression is correctly specified, the QTE estimator achieves minimum asymptotic variance. These results are built on high-level conditions concerning the estimates of the auxiliary regressions. The conditions are then verified for auxiliary regressions that are estimated (1) parametrically, (2) nonparametrically, or (3) via regularization in high dimensional cases. Specifically, for (1) we derive the pseudo true values of the parameters that minimize the

asymptotic variance of the QTE estimator and verify the validity of quasi maximum likelihood estimation of the auxiliary regressions. For (2) and (3), we consider logistic sieve regression and logistic regression under ℓ_1 penalization when the additional covariates are high-dimensional. The limit theory holds uniformly over a compact set of quantile indexes, implying that our multiplier bootstrap procedure can be used to conduct inference on QTEs involving single, multiple, or a continuum of quantile indexes.

From a practical perspective, our estimation and inferential methods have four advantages. First, they allow for common choices of auxiliary regressions such as linear probability, logit, and probit regressions, even though these regressions may be misspecified. Second, the methods can be implemented without tuning parameters. Third, our (bootstrap) estimator can be directly computed via the subgradient condition, and the auxiliary regressions need not be re-estimated in the bootstrap procedure, both of which save considerable computation time. Last, our estimation and inference methods can be implemented without the knowledge of the exact treatment assignment rule used in the experiment. Such information may not be available when researchers are using an experiment that was run in the past and the randomization procedure may not have been fully described. It also occurs in subsample analysis, where sub-groups are defined using variables that may have not been used to form the strata and the treatment assignment rule for each sub-group becomes unknown.

The contributions in the present paper relate to two strands of the present literature. The first is causal inference under CARs. Shao, Yu, and Zhong (2010) showed that the usual two-sample t-test for the ATE is conservative under CARs and proposed a covariate-adaptive bootstrap to conduct inference. Bugni, Canay, and Shaikh (2018) extended the framework to the nonparametric setup, proposed an adjusted standard error for the ATE estimator, and considered a permutation test. Zhang and Zheng (2020) showed that, under CARs, the standard bootstrap inference for the QTE is conservative under the null and lacks power under the alternative. Instead, they suggested multiplier bootstrapping the inverse propensity score weighted (IPW) estimator of the QTE and showed that the distribution of the bootstrap estimator can mimic the original one under CARs. Additional studies in this literature include Hu and Hu (2012); Shao and Yu (2013); Ma, Hu, and Zhang (2015); Ma, Qin, Li, and Hu (2018); Ye (2018); Bugni, Canay, and Shaikh (2019); Ye and Shao (2020); Bugni and Gao (2021); Olivares (2021). Our work complements these studies by considering estimation and bootstrap inference of the QTE with regression adjustments.

A second strand of research deals with regression adjustment of treatment effect estimation in randomized experiments. Freedman (2008a,b) pointed out that ordinary least squares regression adjustment in randomized experiments can degrade the precision of the ATE estimator. Lin (2013) reexamined Freedman’s critique and pointed out that linear regression adjustment does not lead

to loss of precision when a full set of interactions between the treatment status and covariates is included. Lei and Ding (2021) extended Lin (2013) to situations where the dimension of the covariates can grow with the sample size at a certain rate. Bloniarz, Liu, Zhang, Sekhon, and Yu (2016) provided a further extension to cases where the dimension of the covariates exceeds the sample size, proposing the use of Lasso in the auxiliary regression. Ye, Yi, and Shao (2020) and Ma, Tu, and Liu (2020) studied regression adjustment of ATE estimates under CARs. Liu, Tu, and Ma (2020) established a general theory for regression-adjusted ATE estimation under CARs, incorporating OLS and Lasso estimation as special cases. The other works in this branch include, but are limited to, Lu (2016), Fogarty (2018), Li and Ding (2020), Liu and Yang (2020), Negi and Wooldridge (2020), and Zhao and Ding (2020). All of these papers focus on ATE. Zhang, Tsiatis, and Davidian (2008) considered the regression-adjustment for a semiparametric model and linked the estimation method with the theory of semiparametric efficiency. The present paper complements the above work by studying QTE regression, which is nonparametrically specified, with both linear (linear probability model) and nonlinear (logit and probit models) regression adjustments. Similar to Liu et al. (2020), we establish a general theory for regression adjustment that allows for parametric, nonparametric, and regularized estimation of the auxiliary regressions. A further contribution to the literature is to establish the validity of a multiplier bootstrap inferential procedure that does not suffer from conservatism under CARs.

The present paper also comes under the umbrella of a growing literature that studies estimation and inference in randomized experiments. In this connection, we mention the work of Hahn, Hirano, and Karlan (2011); Athey and Imbens (2017); Abadie, Chingos, and West (2018); Tabord-Meehan (2018); Bai, Shaikh, and Romano (2019); Bai (2019); Jiang, Liu, Phillips, and Zhang (2021) among many others.

The remainder of the paper is organized as follows. Section 2 describes the model setup and notation. Section 3 develops the asymptotic properties of our regression-adjusted QTE estimator. Section 4 studies the validity of the multiplier bootstrap inference. Section 5 considers parametric, nonparametric, and regularized estimation of the auxiliary regressions. Section 6 provides computational details and recommendations for practitioners. Section 7 reports simulation results, and an empirical application of our methods to the impact of child health and nutrition on educational outcomes is provided in Section 8. Section 9 concludes. Proofs of all results and some additional simulations are given in the Online Supplement.

2 Setup and Notation

Potential outcomes for treated and control groups are denoted by $Y(1)$ and $Y(0)$, respectively. Treatment status is denoted by A , with $A = 1$ indicating treated and $A = 0$ untreated. The stratum indicator is denoted by S and is based on how the researcher implements the covariate-adaptive randomization. The support of S is denoted by \mathcal{S} , a finite set. After randomization, the researcher can observe the data $\{Y_i, S_i, A_i, X_i\}_{i=1}^n$ where $Y_i = Y_i(1)A_i + Y_i(0)(1 - A_i)$ is the observed outcome and X_i contains extra covariates besides S_i in the dataset. The support of X is denoted $\text{Supp}(X)$. In this paper, we allow X_i and S_i to be dependent. For $1 \leq i \leq n$ and $[n] = \{1, 2, \dots, n\}$, let $p(s) = \mathbb{P}(S_i = s)$, $n(s) = \sum_{i \in [n]} 1\{S_i = s\}$, $n_1(s) = \sum_{i \in [n]} A_i 1\{S_i = s\}$, and $n_0(s) = n(s) - n_1(s)$. We make the following assumptions on the data generating process (DGP) and the treatment assignment rule.

Assumption 1. (i) $\{Y_i(1), Y_i(0), S_i, X_i\}_{i=1}^n$ is *i.i.d.*

(ii) $\{Y_i(1), Y_i(0), X_i\}_{i=1}^n \perp\!\!\!\perp \{A_i\}_{i=1}^n \mid \{S_i\}_{i=1}^n$.

(iii) Suppose $p(s)$ is fixed w.r.t. n and is positive for every $s \in \mathcal{S}$.

(iv) Let $\pi(s)$ denote the target fraction of treated units in stratum s . Then, $c < \min_{s \in \mathcal{S}} \pi(s) \leq \max_{s \in \mathcal{S}} \pi(s) < 1 - c$ for some constant $c \in (0, 0.5)$ and $\frac{D_n(s)}{n(s)} = o_p(1)$ for $s \in \mathcal{S}$, where $D_n(s) = \sum_{i=1}^n (A_i - \pi(s)) 1\{S_i = s\}$.

Several remarks are in order. First, Assumption 1(i) allows for the cross-sectional dependence among treatment statuses ($\{A_i\}_{i \in [n]}$), thereby accommodating many covariate-adaptive randomization schemes as discussed below. Second, although treatment statuses are cross-sectionally dependent, they are independent of the potential outcomes and additional covariates conditional on the stratum indicator S . Therefore, data are still experimental rather than observational. Third, Assumption 1(iii) requires the size of each stratum to be proportional to the sample size. Fourth, we call $\pi(s)$ the target fraction of treated units in stratum s . Similar to Bugni et al. (2019), we allow the target fractions to differ across strata. Just as for the overlapping support condition in an observational study, the target fractions are assumed to be bounded away from zero and one. In randomized experiments, this condition usually holds because investigators can determine $\pi(s)$ in the design stage. In fact, in most CARs, $\pi(s) = 0.5$ for $s \in \mathcal{S}$. Fifth, $D_n(s)$ represents the degree of imbalance between the real and target fractions of treated units in the s th stratum. Bugni et al. (2018) show that Assumption 1(iv) holds under several covariate-adaptive treatment assignment rules such as simple random sampling (SRS), biased-coin design (BCD), adaptive biased-coin design (WEI), and stratified block randomization (SBR). For completeness, we briefly repeat their descriptions below. Note we only require $D_n(s)/n(s) = o_p(1)$, which is weaker than the assumption

imposed by Bugni et al. (2018) but the same as that imposed by Bugni et al. (2019) and Zhang and Zheng (2020).

Example 1 (SRS). Let $\{A_i\}_{i=1}^n$ be drawn independently across i and of $\{S_i\}_{i=1}^n$ as Bernoulli random variables with success rate π , i.e., for $k = 1, \dots, n$,

$$\mathbb{P}\left(A_k = 1 \mid \{S_i\}_{i=1}^n, \{A_j\}_{j=1}^{k-1}\right) = \mathbb{P}(A_k = 1) = \pi(S_i).$$

Then, Assumption 1(iv) holds.

Example 2 (WEI). This design was first proposed by Wei (1978). Let $n_{k-1}(S_k) = \sum_{i=1}^{k-1} 1\{S_i = S_k\}$, $D_{k-1}(s) = \sum_{i=1}^{k-1} (A_i - \frac{1}{2}) 1\{S_i = s\}$, and

$$\mathbb{P}\left(A_k = 1 \mid \{S_i\}_{i=1}^k, \{A_i\}_{i=1}^{k-1}\right) = \phi\left(\frac{D_{k-1}(S_k)}{n_{k-1}(S_k)}\right),$$

where $\phi(\cdot) : [-1, 1] \mapsto [0, 1]$ is a pre-specified non-increasing function satisfying $\phi(-x) = 1 - \phi(x)$. Here, $\frac{D_0(S_1)}{0}$ is understood to be zero. Then, Bugni et al. (2018) show that Assumption 1(iv) holds with $\pi(s) = \frac{1}{2}$. Recently, Hu (2016) generalized the adaptive biased-coin design to multiple treatment values and unequal target factions.

Example 3 (BCD). The treatment status is determined sequentially for $1 \leq k \leq n$ as

$$\mathbb{P}\left(A_k = 1 \mid \{S_i\}_{i=1}^k, \{A_i\}_{i=1}^{k-1}\right) = \begin{cases} \frac{1}{2} & \text{if } D_{k-1}(S_k) = 0 \\ \lambda & \text{if } D_{k-1}(S_k) < 0 \\ 1 - \lambda & \text{if } D_{k-1}(S_k) > 0, \end{cases}$$

where $D_{k-1}(s)$ is defined as above and $\frac{1}{2} < \lambda \leq 1$. Then, Bugni et al. (2018) show that Assumption 1(iv) holds with $\pi = \frac{1}{2}$.

Example 4 (SBR). For each stratum, $\lfloor \pi(s)n(s) \rfloor$ units are assigned to treatment and the rest are assigned to control. Then obviously Assumption 1(iv) holds as $n(s) \rightarrow \infty$.

Denote the τ th quantile of $Y(a)$ by $q_a(\tau)$ for $a = 0, 1$. We are interested in estimating and inferring the τ th quantile treatment effect defined as $q(\tau) = q_1(\tau) - q_0(\tau)$. The testing problems of interest involve single, multiple, or even a continuum of quantile indexes, as in the following null hypotheses

$$\begin{aligned} \mathcal{H}_0 : q(\tau) &= \underline{q} \quad \text{versus} \quad q(\tau) \neq \underline{q}, \\ \mathcal{H}_0 : q(\tau_1) - q(\tau_2) &= \underline{q} \quad \text{versus} \quad q(\tau_1) - q(\tau_2) \neq \underline{q}, \text{ and} \\ \mathcal{H}_0 : q(\tau) &= \underline{q}(\tau) \quad \forall \tau \in \Upsilon \quad \text{versus} \quad q(\tau) \neq \underline{q}(\tau) \text{ for some } \tau \in \Upsilon, \end{aligned}$$

for some pre-specified value \underline{q} or function $\underline{q}(\tau)$, where Υ is some compact subset of $(0, 1)$.

3 Estimation

Our estimator is based on the doubly robust moments for the quantiles of $Y(1)$ and $Y(0)$. Define $m_a(\tau, s, x) = \tau - \mathbb{P}(Y_i(a) \leq q_a(\tau) | S_i = s, X_i = x)$. Firpo (2007), Belloni, Chernozhukov, Fernández-Val, and Hansen (2017), and Kallus, Mao, and Uehara (2019) showed that the doubly robust moment for $q_1(\tau)$ is

$$\mathbb{E} \left[\frac{A_i(\tau - 1\{Y_i(1) \leq q\})}{\bar{\pi}(S_i)} - \frac{A_i - \bar{\pi}(S_i)}{\bar{\pi}(S_i)} \bar{m}_1(\tau, S_i, X_i) \right] = 0, \quad (3.1)$$

where $\bar{\pi}(s)$ and $\bar{m}_1(\tau, s, x)$ are the working models for the propensity score ($\pi(s)$) and conditional probability ($m_1(\tau, s, x)$), respectively. In CAR, the propensity score is usually known or can be consistently estimated by $\hat{\pi}(s) = n_1(s)/n(s)$. Therefore, $\bar{\pi}(s)$ is correctly specified as $\pi(s)$. Then, due to the double robustness, even though the working model $\bar{m}_1(\tau, s, x)$ is misspecified, (3.1) still identifies the true unconditional quantile $q_1(\tau)$. Therefore, it is expected that the estimator of $q_1(\tau)$ based on the sample analogue of (3.1) is still consistent. In addition, (3.1) also satisfies the Neyman orthogonality condition w.r.t. the working models $\bar{\pi}(\cdot)$ and $\bar{m}_1(\cdot)$. This implies our QTE estimator still has the usual parametric rate of convergence even when the estimators of the working models converge at a nonparametric rate.

To proceed, we note that (3.1) is the first-order condition of the optimization problem

$$\min_q \mathbb{E} \left[\frac{A_i}{\pi(S_i)} \rho_\tau(Y_i - q) + \frac{A_i - \pi(S_i)}{\pi(S_i)} \bar{m}_1(\tau, S_i, X_i) q \right], \quad (3.2)$$

where $\rho_\tau(u) = u(\tau - 1\{u \leq 0\})$ is the usual check function. Denote $\hat{m}_1(\cdot)$ as a consistent estimator of $\bar{m}_1(\cdot)$ obtained via an auxiliary regression. By the sample analog of (3.2), our regression-adjusted estimator of $q_1(\tau)$, denoted as $\hat{q}_1^{adj}(\tau)$, can be defined as

$$\hat{q}_1^{adj}(\tau) = \arg \min_q \sum_{i \in [n]} \left[\frac{A_i}{\hat{\pi}(S_i)} \rho_\tau(Y_i - q) + \frac{(A_i - \hat{\pi}(S_i))}{\hat{\pi}(S_i)} \hat{m}_1(\tau, S_i, X_i) q \right]. \quad (3.3)$$

Here, we emphasize that $\hat{m}_1(\cdot)$ may not consistently estimate the true specification $m_1(\cdot)$. Similarly, we can define

$$\hat{q}_0^{adj}(\tau) = \arg \min_q \sum_{i \in [n]} \left[\frac{1 - A_i}{1 - \hat{\pi}(S_i)} \rho_\tau(Y_i - q) - \frac{(A_i - \hat{\pi}(S_i))}{1 - \hat{\pi}(S_i)} \hat{m}_0(\tau, S_i, X_i) q \right], \quad (3.4)$$

where $\hat{m}_0(\cdot)$ is a consistent estimator of $\bar{m}_0(\cdot)$. Then, our regression adjusted QTE estimator is

$$\hat{q}^{adj}(\tau) = \hat{q}_1^{adj}(\tau) - \hat{q}_0^{adj}(\tau). \quad (3.5)$$

We use the estimated propensity score $\hat{\pi}(s)$ even when $\pi(s)$ is known. Consider the special case of our setup for which there is no regression adjustment ($\widehat{m}_a(\cdot) = 0$) and $\pi(s) = 1/2$ for $s \in \mathcal{S}$. In the design stage, researchers can use simple random sampling (SRS) or CAR. In the analysis stage, researchers can choose the difference in quantile estimator (DIQ) which is just the difference of the τ th quantiles of the treated and control groups or the IPW estimator with $\hat{\pi}(s)$ defined in (3.3)–(3.5). Note when $\pi(s) = 1/2$, the DIQ estimator can be viewed as an IPW estimator in which $\hat{\pi}(s)$ is replaced by the true propensity score $\pi(s) = 1/2$. In this simplified setting, Zhang and Zheng (2020) derived the asymptotic variances of two estimators under SRS and CAR, which can be summarized below.

Table 1: Asymptotic Variance of DIQ and IPW estimators

	DIQ	IPW
SRS	$\Sigma(1/4)$	$\Sigma(0)$
CAR(δ)	$\Sigma(\delta)$	$\Sigma(0)$

In Table 1, the asymptotic variance $\Sigma(\delta)$ is a monotone increasing function of δ and $\delta \in [0, 1/4]$ indicates the balance level achieved by the CAR scheme. When $\delta = 0$, the CAR achieves the strong balance. For the four examples of CARs mentioned earlier, we have $\delta(\text{SRS}) = \frac{1}{4}$, $\delta(\text{WEI}) = \frac{1}{4}(1 - \phi'(0))^{-1}$, $\delta(\text{BCD}) = 0$, and $\delta(\text{SBR}) = 0$. From Table 1, we can make two observations. First, the asymptotic variance of the QTE estimator with $\hat{\pi}(s)$ (i.e., IPW) is no larger than that with $\pi(s)$ (i.e., DIQ). Second, when we conduct multiplier bootstrap inference, the bootstrap weights are independent cross-sectionally. Therefore, conditionally on the data, the observations in the bootstrap sample mimic the SRS scheme. The bootstrap variance for the DIQ estimator is conservative because $\Sigma(1/4) \geq \Sigma(\delta)$ while it is not for the IPW estimator. This motivates us to use the IPW estimator with the estimated propensity score $\hat{\pi}(s)$ for both estimation and bootstrap inference, even when $\pi(s)$ is known.

In this and the next sections, we establish the limit distribution and validity of bootstrap inference for $\hat{q}^{adj}(\tau)$ under high-level conditions on $(\bar{m}_a(\cdot), \widehat{m}_a(\cdot))_{a=0,1}$. In Section 5, we verify these conditions when $(\bar{m}_a(\cdot))_{a=0,1}$ are correctly or incorrectly specified and $(\widehat{m}_a(\cdot))_{a=0,1}$ are computed parametrically, nonparametrically, or via regularization when the covariate X is high-dimensional.

Assumption 2. For $a = 0, 1$, denote $f_a(\cdot)$, $f_a(\cdot|s)$, and $f_a(\cdot|x, s)$ as the PDFs of $Y_i(a)$, $Y_i(a)|S_i = s$, and $Y_i(a)|S_i = s, X_i = x$, respectively.

- (i) $f_a(q_a(\tau))$ and $f_a(q_a(\tau)|s)$ are bounded and bounded away from zero uniformly over $\tau \in \Upsilon$ and $s \in \mathcal{S}$, where Υ is a compact subset of $(0, 1)$.
- (ii) $f_a(\cdot)$ and $f_a(\cdot|s)$ are Lipschitz over $\{q_j(\tau) : \tau \in \Upsilon\}$.
- (iii) $\sup_{y \in \mathbb{R}, x \in \text{Supp}(X), s \in \mathcal{S}} f_a(y|x, s) < \infty$.

Assumption 3. (i) For $a = 0, 1$, there exists a function $\bar{m}_a(\tau, x, s)$ such that for $\bar{\Delta}_a(\tau, s, X_i) = \hat{m}_a(\tau, s, X_i) - \bar{m}_a(\tau, s, X_i)$, we have

$$\sup_{\tau \in \Upsilon, s \in \mathcal{S}} \left| \frac{\sum_{i \in I_1(s)} \bar{\Delta}_a(\tau, s, X_i)}{n_1(s)} - \frac{\sum_{i \in I_0(s)} \bar{\Delta}_a(\tau, s, X_i)}{n_0(s)} \right| = o_p(n^{-1/2}),$$

where $I_a(s) = \{i \in [n] : A_i = a, S_i = s\}$.

(ii) For $a = 0, 1$, let $\mathcal{F}_a = \{\bar{m}_a(\tau, S_i, X_i) : \tau \in \Upsilon\}$ with an envelope $F_{a,i}$. Then, $\max_{s \in \mathcal{S}} \mathbb{E}(|F_{a,i}|^q | S_i = s) < \infty$ for $q > 2$ and there exist fixed constants $(\alpha, v) > 0$ such that

$$\sup_Q N(\mathcal{F}_a, e_Q, \varepsilon \|F_a\|_{Q,2}) \leq \left(\frac{\alpha}{\varepsilon}\right)^v, \quad \forall \varepsilon \in (0, 1],$$

where $N(\cdot)$ denotes the covering number, $e_Q(f, g) = \|f - g\|_{Q,2}$, and the supremum is taken over all finitely discrete probability measures Q .

(iii) For any $\tau_1, \tau_2 \in \Upsilon$, there exists a constant $C > 0$ such that $\mathbb{E}((\bar{m}_a(\tau_2, S_i, X_i) - \bar{m}_a(\tau_1, S_i, X_i))^2 | S_i = s) \leq C|\tau_2 - \tau_1|$.

Several remarks are in order. First, Assumption 2 is standard in the quantile regression literature. We do not need $f_a(y|x, s)$ to be bounded away from zero because we are interested in the unconditional quantile $q_a(\tau)$, which is uniquely defined as long as the unconditional density $f_a(q_a(\tau))$ is positive. Second, Assumption 3(i) is high-level. If we consider a linear probability model such that $\bar{m}_a(\tau, s, X_i) = \tau - X_i^\top \theta_{a,s}(\tau)$ and $\hat{m}_a(\tau, s, X_i) = \tau - X_i^\top \hat{\theta}_{a,s}(\tau)$, then Assumption 3(i) is equivalent to

$$\sup_{\tau \in \Upsilon, a=0,1, s \in \mathcal{S}} \left| \left(\frac{\sum_{i \in I_1(s)} X_i}{n_1(s)} - \frac{\sum_{i \in I_0(s)} X_i}{n_0(s)} \right)^\top (\hat{\theta}_{a,s}(\tau) - \theta_{a,s}(\tau)) \right| = o_p(n^{-1/2}),$$

which is similar to Liu et al. (2020, Assumption 3) and holds intuitively if $\hat{\theta}_{a,s}(\tau)$ is a consistent estimator of the pseudo true value $\theta_{a,s}(\tau)$. Third, Assumptions 3(ii) and 3(iii) impose mild regularity conditions on $\bar{m}_a(\cdot)$. They hold if

$$\sup_{a=0,1, s \in \mathcal{S}, x \in \text{Supp}(X)} |\bar{m}_a(\tau_2, s, x) - \bar{m}_a(\tau_1, s, x)| \leq L|\tau_2 - \tau_1|.$$

for some constant $L > 0$. Such Lipschitz continuity holds for the true specification ($\bar{m}_a(\cdot) = m_a(\cdot)$) under Assumption 2. Fourth, we provide primitive sufficient conditions for Assumption 3 in Section 5.

Theorem 3.1. *Suppose Assumptions 1–3 hold. Then, uniformly over $\tau \in \Upsilon$,*

$$\sqrt{n}(\hat{q}^{adj}(\tau) - q(\tau)) \rightsquigarrow \mathcal{B}(\tau),$$

where $\mathcal{B}(\tau)$ is a tight Gaussian process with covariance kernel $\Sigma(\tau, \tau')$ defined in Section F of the Online Supplement. In addition, for any finite set of quantile indexes (τ_1, \dots, τ_K) , the asymptotic covariance matrix of $(\hat{q}^{adj}(\tau_1), \dots, \hat{q}^{adj}(\tau_K))$ is denoted as $[\Sigma(\tau_k, \tau_l)]_{k, l \in [K]}$, where we use $[A_{kl}]_{k, l \in [K]}$ to denote a $K \times K$ matrix whose (k, l) th entry is A_{kl} . Then, $[\Sigma(\tau_k, \tau_l)]_{k, l \in [K]}$ is minimized in the matrix sense¹ when the auxiliary regressions are correctly specified at (τ_1, \dots, τ_K) , i.e., $\bar{m}_a(\tau_k, s, x) = m_a(\tau_k, s, x)$ for $a = 0, 1$, $k \in [K]$, and all (s, x) in the joint support of (S_i, X_i) .

Two remarks are in order. First, the expression for the asymptotic variance of $\hat{q}^{adj}(\tau)$ can be found in the proof of Theorem 3.1. It is the same whether the randomization scheme achieves strong balance or not. This robustness is due to the use of the estimated propensity score $(\hat{\pi}(s))$. The same phenomenon was discovered in the simplified setting by Zhang and Zheng (2020) as shown in the (2,2) entry of Table 1. Second, the asymptotic variance of $\hat{q}^{adj}(\tau)$ depends on $(f_a(q_a(\tau)), m_a(\tau, s, x))_{a=0,1}$, which are infinite-dimensional nuisance parameters. To conduct analytic inference, it is necessary to nonparametrically estimate these nuisance parameters, which requires tuning parameters. Nonparametric estimation can be sensitive to the choice of tuning parameters and rule-of-thumb tuning parameter selection may not be appropriate for every data generating process (DGP) or every quantile. Use of cross-validation in selecting the tuning parameters is possible in principle but in practice time-consuming. These practical difficulties of analytic methods of inference provide strong motivation to investigate bootstrap inference procedures that are much less reliant on tuning parameters.

4 Multiplier Bootstrap Inference

We approximate the asymptotic distributions of $\hat{q}^{adj}(\tau)$ via the multiplier bootstrap. Let $\{\xi_i\}_{i=1}^n$ be a sequence of bootstrap weights which will be specified later. Define $n_1^w(s) = \sum_{i=1}^n \xi_i A_i 1\{S_i = s\}$, $n_0^w(s) = \sum_{i=1}^n \xi_i (1 - A_i) 1\{S_i = s\}$, $n^w(s) = \sum_{i=1}^n \xi_i 1\{S_i = s\} = n_1^w(s) + n_0^w(s)$, and $\hat{\pi}^w(s) = n_1^w(s)/n^w(s)$. The multiplier bootstrap counterpart of $\hat{q}^{adj}(\tau)$ is denoted by $\hat{q}^w(\tau)$ and defined as

$$\hat{q}^w(\tau) = \hat{q}_1^w(\tau) - \hat{q}_0^w(\tau),$$

where

$$\hat{q}_1^w(\tau) = \arg \min_q \sum_{i \in [n]} \xi_i \left[\frac{A_i}{\hat{\pi}^w(S_i)} \rho_\tau(Y_i - q) + \frac{(A_i - \hat{\pi}^w(S_i))}{\hat{\pi}^w(S_i)} \hat{m}_1(\tau, S_i, X_i) q \right], \quad (4.1)$$

and

$$\hat{q}_0^w(\tau) = \arg \min_q \sum_{i \in [n]} \xi_i \left[\frac{1 - A_i}{1 - \hat{\pi}^w(S_i)} \rho_\tau(Y_i - q) - \frac{(A_i - \hat{\pi}^w(S_i))}{1 - \hat{\pi}^w(S_i)} \hat{m}_0(\tau, S_i, X_i) q \right]. \quad (4.2)$$

¹For two symmetric matrices A and B , we say A is greater than or equal to B if $A - B$ is positive semidefinite.

Two comments on implementation are noted here: (i) we do not re-estimate $\widehat{m}_a(\cdot)$ in the bootstrap sample, which is similar to the multiplier bootstrap procedure proposed by Belloni et al. (2017); and (ii) in Section 6 we propose a way to directly compute $(\hat{q}_a^w(\tau))_{a=0,1}$ from the subgradient conditions of (4.1) and (4.2), thereby avoiding the optimization. Both features considerably reduce computation time of our bootstrap procedure.

Next, we specify the bootstrap weights.

Assumption 4. *Suppose $\{\xi_i\}_{i=1}^n$ is a sequence of nonnegative i.i.d. random variables with unit expectation and variance and a sub-exponential upper tail.*

Assumption 5. *Recall $\bar{\Delta}_a(\tau, s, x)$ defined in Assumption 3. We have, for $a = 0, 1$,*

$$\sup_{\tau \in \Upsilon, s \in \mathcal{S}} \left| \frac{\sum_{i \in I_1(s)} \xi_i \bar{\Delta}_a(\tau, s, X_i)}{n_1^w(s)} - \frac{\sum_{i \in I_0(s)} \xi_i \bar{\Delta}_a(\tau, s, X_i)}{n_0^w(s)} \right| = o_p(n^{-1/2}).$$

We require the bootstrap weights to be nonnegative so that the objective functions in (4.1) and (4.2) are convex. In practice, we generate ξ_i independently from the standard exponential distribution. Assumption 5 is the bootstrap counterpart of Assumption 3. Continuing with the linear model example considered after Assumption 3, Assumption 5 requires

$$\sup_{\tau \in \Upsilon, a=0,1, s \in \mathcal{S}} \left| \left(\frac{\sum_{i \in I_1(s)} \xi_i X_i}{n_1(s)} - \frac{\sum_{i \in I_0(s)} \xi_i X_i}{n_0(s)} \right)^\top \left(\hat{\theta}_{a,s}(\tau) - \theta_{a,s}(\tau) \right) \right| = o_p(n^{-1/2}),$$

which holds if $\hat{\theta}_{a,s}(\tau)$ is a uniformly consistent estimator of $\theta_{a,s}(\tau)$.

Theorem 4.1. *Suppose Assumptions 1–5 hold. Then, uniformly over $\tau \in \Upsilon$ and conditionally on data,*

$$\sqrt{n}(\hat{q}^w(\tau) - \hat{q}^{adj}(\tau)) \rightsquigarrow \mathcal{B}(\tau),$$

where $\mathcal{B}(\tau)$ is the same Gaussian process defined in Theorem 3.1.²

Theorem 4.1 shows the limit distribution of the bootstrap estimator conditional on data can approximate that of the original estimator uniformly over $\tau \in \Upsilon$. This provides the theoretical foundation for our bootstrap confidence intervals and bands described in Section 6. We also emphasize

²We view $\sqrt{n}(\hat{q}^w(\tau) - \hat{q}^{adj}(\tau))$ and $\mathcal{B}(\tau)$ as two processes indexed by $\tau \in \Upsilon$ and denote them as G_n and G , respectively. Then, following van der Vaart and Wellner (1996, Chapter 2.9), we say G_n weakly converges to G conditionally on data and uniformly over $\tau \in \Upsilon$ if

$$\sup_{h \in \text{BL}_1} |\mathbb{E}_\xi h(G_n) - \mathbb{E}h(G)| \xrightarrow{P} 0,$$

where BL_1 is the set of all functions $h : \ell^\infty(\Upsilon) \mapsto [0, 1]$ such that $|h(z_1) - h(z_2)| \leq |z_1 - z_2|$ for every $z_1, z_2 \in \ell^\infty(\Upsilon)$, and \mathbb{E}_ξ denotes expectation with respect to the bootstrap weights $\{\xi\}_{i=1}^n$.

that the distribution of the bootstrap estimator conditional on data can consistently approximate that of the original estimator under CAR because the estimated propensity score, i.e., $\hat{\pi}(s)$, is used in (4.1) and (4.2).

5 Auxiliary Regressions

In this section, we consider three approaches to estimation for the auxiliary regressions: (1) a parametric method, (2) a nonparametric method, and (3) a regularization method. For the parametric method, we do not require the model to be correctly specified. We propose ways to estimate the pseudo true value of the auxiliary regression. For the other two methods, we (nonparametrically) estimate the true model so that the asymptotic variance of $\hat{q}^{adj}(\tau)$ achieves its minimum based on Theorem 3.1. For all three methods, we verify Assumptions 3 and 5.

5.1 Parametric Method

In this section, we consider the case where X_i is finite-dimensional. Recall $m_a(\tau, S_i, X_i) \equiv \tau - \mathbb{P}(Y_i(a) \leq q_a(\tau) | X_i, S_i)$ for $a = 0, 1$. We propose to model $\mathbb{P}(Y_i(a) \leq q_a(\tau) | X_i, S_i = s)$ as $\Lambda_{a,s}(X_i, \theta_{a,s}(\tau))$ so that our model for $m_a(\tau, S_i, X_i)$ is

$$\bar{m}_a(\tau, S_i, X_i) = \tau - \sum_{s \in \mathcal{S}} 1\{S_i = s\} \Lambda_{a,s}(X_i, \theta_{a,s}(\tau)) \quad (5.1)$$

We propose to estimate $\theta_{a,s}(\tau)$ by $\hat{\theta}_{a,s}(\tau)$. The corresponding $\hat{m}(\tau, S_i, X_i)$ can be written as

$$\hat{m}_a(\tau, S_i, X_i) = \tau - \sum_{s \in \mathcal{S}} 1\{S_i = s\} \Lambda_{a,s}(X_i, \hat{\theta}_{a,s}(\tau)) \quad (5.2)$$

Assumption 6. (i) *Suppose there exist a positive random variable L_i and a positive constant $C > 0$ such that*

$$\begin{aligned} \sup_{\tau \in \Upsilon, a=0,1, s \in \mathcal{S}} \|\partial_{\theta} \Lambda_{a,s}(X_i, \theta_{a,s}(\tau))\|_2 &\leq L_i, \\ \sup_{\tau \in \Upsilon, a=0,1, s \in \mathcal{S}} |\Lambda_{a,s}(X_i, \theta_{a,s}(\tau))| &\leq L_i, \end{aligned}$$

and $\mathbb{E}(L_i^q | S_i = s) < \infty$.

(ii) *Suppose $\sup_{\tau_1, \tau_2 \in \Upsilon, a=0,1, s \in \mathcal{S}} |\theta_{a,s}(\tau_1) - \theta_{a,s}(\tau_2)| \leq C|\tau_1 - \tau_2|$.*

(iii) *There exists $\hat{\theta}_{a,s}(\tau)$ such that $\sup_{\tau \in \Upsilon, a=0,1, s \in \mathcal{S}} \|\hat{\theta}_{a,s}(\tau) - \theta_{a,s}(\tau)\|_2 \xrightarrow{p} 0$.*

Theorem 5.1. *Denote $\hat{q}^{par}(\tau)$ and $\hat{q}^{par,w}(\tau)$ as the τ th QTE estimator and its multiplier bootstrap counterpart defined in Sections 3 and 4, respectively, with $\bar{m}_a(\tau, S_i, X_i)$ and $\hat{m}_a(\tau, S_i, X_i)$ defined in (5.1) and (5.2), respectively. Suppose Assumptions 1, 2, 4, and 6 hold. Then, Assumptions 3*

and 5 hold, which further implies Theorems 3.1 and 4.1 hold for $\hat{q}^{par}(\tau)$ and $\hat{q}^{par,w}(\tau)$, respectively.

Several remarks are in order. First, common choices for auxiliary regressions are linear probability, logistic, and probit regressions, corresponding to $\Lambda_{a,s}(X_i, \theta_{a,s}(\tau)) = X_i^\top \theta_{a,s}(\tau)$, $\lambda(\bar{X}_i^\top \theta_{a,s}(\tau))$, and $\Phi(\bar{X}_i^\top \theta_{a,s}(\tau))$, respectively, where $\lambda(\cdot)$ and $\Phi(\cdot)$ are the logistic and standard normal CDFs, respectively, and $\bar{X}_i = (1, X_i^\top)^\top$. For the linear regression case, we do not include the intercept because our regression adjusted estimators ((3.3) and (3.4)) and their bootstrap counterparts ((4.1) and (4.2)) are numerically invariant to location shift of the auxiliary models. Second, Theorem 5.1 shows that, as long as the estimator of the pseudo true value ($\hat{\theta}_{a,s}(\tau)$) is uniformly consistent, under mild regularity conditions, all the general estimation and bootstrap inference results established in Sections 3 and 4 hold. Third, it is still in question what pseudo true value one should target. We note that the asymptotic variance (denoted as σ^2) of the $\hat{q}^{adj}(\tau)$ is a function of the working model ($\bar{m}_a(\tau, s, \cdot)$), which is further indexed by its parameters (denoted as $\{t_{a,s}(\tau)\}_{a=0,1,s \in \mathcal{S}}$), i.e., $\sigma^2 = \sigma^2(\{\bar{m}_a(\tau, s, \cdot; t_{a,s})\}_{a=0,1,s \in \mathcal{S}})$. One option is to target the pseudo true value that minimizes $\sigma^2(\{\bar{m}_a(\tau, s, \cdot; t_{a,s})\}_{a=0,1,s \in \mathcal{S}})$, i.e.,

$$\{\theta_{a,s}(\tau)\}_{a=0,1,s \in \mathcal{S}} \in \arg \min_{t_{a,s}: a=0,1,s \in \mathcal{S}} \sigma^2(\{\bar{m}_a(\tau, s, \cdot; t_{a,s})\}_{a=0,1,s \in \mathcal{S}}).$$

Theorem 5.2 below characterizes such $\{\theta_{a,s}(\tau)\}_{a=0,1,s \in \mathcal{S}}$. We then consider the consistent estimator of these pseudo true values for linear and logistic models in Propositions 5.1 and 5.2, respectively. Fourth, as the working model may be misspecified, the pseudo true value that minimizes $\sigma^2(\{\bar{m}_a(\tau, s, \cdot; t_{a,s})\}_{a=0,1,s \in \mathcal{S}})$ may not be the best parameter value that fits the data. On the other hand, the minimum of $\sigma^2(\{\bar{m}_a(\tau, s, \cdot; t_{a,s})\}_{a=0,1,s \in \mathcal{S}})$ w.r.t. $\{t_{a,s}\}_{a=0,1,s \in \mathcal{S}}$ is still no smaller than the one attained by minimizing the whole working model, i.e., $\min_{\bar{m}_a(\tau, s, \cdot), a=0,1,s \in \mathcal{S}} \sigma^2(\{\bar{m}_a(\tau, s, \cdot)\}_{a=0,1,s \in \mathcal{S}})$. The latter minimum is attained if $\bar{m}_a(\tau, s, x) = m_a(\tau, s, x)$ for $a = 0, 1$ and (s, x) in the joint support of (S_i, X_i) , as shown in Theorem 3.1. Therefore, it is still meaningful to pursue a working model that can closely approximate the true model. In Proposition 5.3, we propose to estimate a parametric but flexible distribution regression model via quasi maximum likelihood estimation (QMLE). We give our practical recommendation after that.

Theorem 5.2. *Suppose Assumptions 1, 2, 4, 6 hold, and $\Lambda_{a,s}(X_i, \theta_{a,s}(\tau))$ is differentiable in $\theta_{a,s}(\tau)$. Then, the asymptotic variance of $\hat{q}^{par}(\tau)$ is minimized at $(\theta_{1,s}(\tau), \theta_{0,s}(\tau)) \in \Theta_s(\tau)$, where for $s \in \mathcal{S}$ and $\tau \in \Upsilon$,*

$$\Theta_s(\tau) = \arg \min_{\theta_1, \theta_0} Q(s, \tau, \theta_1, \theta_0), \text{ where}$$

$$Q(s, \tau, \theta_1, \theta_0) = \mathbb{E} \left\{ \left(\frac{g_{1,s}(X_i, \theta_1)}{f_1(q_1(\tau))} + \frac{\pi(s)}{1 - \pi(s)} \frac{g_{0,s}(X_i, \theta_0)}{f_0(q_0(\tau))} \right)^2 \right\}$$

$$\begin{aligned}
& -2 \left(\frac{g_{1,s}(X_i, \theta_1)}{f_1(q_1(\tau))} + \frac{\pi(s)}{1-\pi(s)} \frac{g_{0,s}(X_i, \theta_0)}{f_0(q_0(\tau))} \right) \left(\frac{m_1(\tau, s, X_i) - m_1(\tau, s)}{f_1(q_1(\tau))} \right) \\
& - \frac{2\pi(s)}{1-\pi(s)} \left(\frac{g_{1,s}(X_i, \theta_1)}{f_1(q_1(\tau))} + \frac{\pi(s)}{1-\pi(s)} \frac{g_{0,s}(X_i, \theta_0)}{f_0(q_0(\tau))} \right) \left(\frac{m_0(\tau, s, X_i) - m_0(\tau, s)}{f_0(q_0(\tau))} \right) \Big|_{S_i = s} \Big\}, \\
& \hspace{15em} (5.3)
\end{aligned}$$

and $g_{a,s}(X_i, \theta_a) = \mathbb{E}(\Lambda_{a,s}(X_i, \theta_a) | S_i = s) - \Lambda_{a,s}(X_i, \theta_a)$. If we further assume a linear probability model, i.e., $\Lambda_{a,s}(X_i, \theta_{a,s}(\tau)) = X_i^\top \theta_{a,s}(\tau)$, and denote the asymptotic covariance matrix of $(\hat{q}^{par}(\tau_1), \dots, \hat{q}^{par}(\tau_K))$ for any finite set of quantile indexes (τ_1, \dots, τ_K) as $[\Sigma^{LP}(\tau_k, \tau_l)]_{k,l \in [K]}$. Then, $[\Sigma^{LP}(\tau_k, \tau_l)]_{k,l \in [K]}$ is minimized in the matrix sense when the pseudo true values $(\theta_{1,s}(\tau_k), \theta_{0,s}(\tau_k))_{k \in [K]}$ satisfy

$$\frac{\theta_{1,s}(\tau_k)}{f_1(q_1(\tau_k))} + \frac{\pi(s)\theta_{0,s}(\tau_k)}{(1-\pi(s))f_0(q_0(\tau_k))} = \frac{\theta_{1,s}^{LP}(\tau_k)}{f_1(q_1(\tau_k))} + \frac{\pi(s)\theta_{0,s}^{LP}(\tau_k)}{(1-\pi(s))f_0(q_0(\tau_k))}, \quad k \in [K],$$

where $[K] = \{1, 2, \dots, K\}$ and

$$\begin{aligned}
\theta_{1,s}^{LP}(\tau_k) &= \left[\mathbb{E}(\tilde{X}_{i,s} \tilde{X}_{i,s}^\top | S_i = s, A_i = 1) \right]^{-1} \mathbb{E} \left[\tilde{X}_{i,s} 1\{Y_i \leq q_1(\tau_k)\} | S_i = s, A_i = 1 \right], \\
\theta_{0,s}^{LP}(\tau_k) &= \left[\mathbb{E}(\tilde{X}_{i,s} \tilde{X}_{i,s}^\top | S_i = s, A_i = 0) \right]^{-1} \mathbb{E} \left[\tilde{X}_{i,s} 1\{Y_i \leq q_0(\tau_k)\} | S_i = s, A_i = 0 \right],
\end{aligned}$$

and $\tilde{X}_{i,s} = X_i - \mathbb{E}(X_i | S_i = s) = X_i - \mathbb{E}(X_i | S_i = s, A_i = 1) = X_i - \mathbb{E}(X_i | S_i = s, A_i = 0)$.

Two remarks are in order. First, we propose consistent estimators of $\theta_{a,s}(\tau)$ for linear probability and logistic models below. The corresponding results for the probit model are similar to those of the logistic model and are therefore omitted. Second, when $\mathbb{E}(\|X_i\|_2^q | S_i = s) < \infty$ for some $q > 2$, Assumption 6(i) holds for these two models. We maintain Assumption 6(ii) as it usually holds under the regularity conditions imposed in Assumption 2. In the following, we aim to establish Assumption 6(iii).

5.1.1 Linear Probability Model

Consider the linear probability model $\Lambda_{a,s}(X_i, \theta_a) = X_i^\top \theta_a$. From Theorem 5.2, we see that $\Theta_s(\tau)$ is not a singleton. In order to achieve the minimal variance, we only need to consistently estimate one point in the set $\Theta_s(\tau)$. We choose

$$(\theta_{1,s}(\tau), \theta_{0,s}(\tau)) = (\theta_{1,s}^{LP}(\tau), \theta_{0,s}^{LP}(\tau)), \quad s \in \mathcal{S},$$

as this choice avoids estimation of the densities $f_1(q_1(\tau))$ and $f_0(q_0(\tau))$.

Note $\theta_{a,s}^{LP}(\tau)$ is the projection coefficient of $1\{Y_i \leq q_a(\tau)\}$ on $\tilde{X}_{i,s}$ for the sub-population with $S_i = s$ and $A_i = a$. We estimate them by the sample analog. The parameter $q_a(\tau)$ is unknown, and thus, is replaced by some \sqrt{n} -consistent estimator denoted as $\hat{q}_a(\tau)$.

Assumption 7. Assume that $\sup_{\tau \in \Upsilon, a=0,1} |\hat{q}_a(\tau) - q_a(\tau)| = O_p(n^{-1/2})$.

In practice, we compute $\{\hat{q}_a(\tau)\}_{a=0,1}$ based on (3.3) and (3.4) by setting $\hat{m}_a(\tau, S_i, X_i) \equiv 0$. Then, Assumption 7 holds automatically by Theorem 3.1 with $\hat{m}_a(\tau, S_i, X_i) = \bar{m}_a(\tau, S_i, X_i) = 0$.

Next, we can define the estimator of $\theta_{a,s}^{LP}(\tau)$. Let

$$\bar{m}_a(\tau, S_i, X_i) = \tau - \sum_{s \in \mathcal{S}} 1\{S_i = s\} X_i^\top \theta_{a,s}^{LP}(\tau), \quad (5.4)$$

$$\hat{m}_a(\tau, S_i, X_i) = \tau - \sum_{s \in \mathcal{S}} 1\{S_i = s\} X_i^\top \hat{\theta}_{a,s}^{LP}(\tau), \quad (5.5)$$

$$\hat{X}_{i,a,s} = X_i - \frac{1}{n_a(s)} \sum_{i \in I_a(s)} X_i, \quad (5.6)$$

and

$$\hat{\theta}_{a,s}^{LP}(\tau) = \left[\frac{1}{n_a(s)} \sum_{i \in I_a(s)} \hat{X}_{i,a,s} \hat{X}_{i,a,s}^\top \right]^{-1} \left[\frac{1}{n_a(s)} \sum_{i \in I_a(s)} \hat{X}_{i,a,s} 1\{Y_i \leq \hat{q}_a(\tau)\} \right]. \quad (5.7)$$

Assumption 8. There exist constants $0 < c < C < \infty$ such that

$$c < \min_{s \in \mathcal{S}} \lambda_{\min}(\mathbb{E} \tilde{X}_{i,s} \tilde{X}_{i,s}^\top | S_i = s) \leq \max_{s \in \mathcal{S}} \lambda_{\max}(\mathbb{E} \tilde{X}_{i,s} \tilde{X}_{i,s}^\top | S_i = s) \leq C$$

and $\mathbb{E}(\|\tilde{X}_{i,s}\|_2^q | S_i = s) \leq C$ for some $q > 2$, where for a generic symmetric matrix A , $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ denote the minimal and maximal eigenvalues of A , respectively.

Proposition 5.1. Suppose Assumptions 1, 2, 7, and 8 hold. Then Assumptions 6(ii) and 6(iii) hold for $(\theta_{a,s}(\tau), \hat{\theta}_{a,s}(\tau)) = (\theta_{a,s}^{LP}(\tau), \hat{\theta}_{a,s}^{LP}(\tau))$, $a = 0, 1, s \in \mathcal{S}, \tau \in \Upsilon$.

5.1.2 Logistic Model

Next consider the logistic model. Set $\lambda(u) = \exp(u)/(1 + \exp(u))$ to be the logistic CDF and then $\Lambda_{a,s}(X_i, \theta_{a,s}(\tau)) = \lambda(\vec{X}_i^\top \theta_{a,s}(\tau))$ and

$$\hat{m}_a(\tau, S_i, X_i) = \tau - \sum_{s \in \mathcal{S}} 1\{S_i = s\} \lambda(\vec{X}_i^\top \hat{\theta}_{a,s}^{LG}(\tau)), \quad (5.8)$$

where $\hat{\theta}_{a,s}^{LG}(\tau)$ is computed as the minimizer of the sample analogue of the objective function in (5.3) and $\vec{X}_i = (1, X_i^\top)^\top$. Specifically, let $I(s) = I_1(s) \cup I_0(s)$, $\hat{g}_{a,s}(X_i, \theta_a) = \frac{1}{n_a(s)} \sum_{i \in I_a(s)} \lambda(\vec{X}_i^\top \theta_a) - \lambda(\vec{X}_i^\top \theta_a)$, and let $\hat{f}_a(\cdot)$ be an estimator of the density of $Y(a)$ (i.e., $f_a(\cdot)$). Define $\Theta = \Theta_1 \times \Theta_0$, where Θ_a is a compact set in \mathfrak{R}^{d_x} for $a = 0, 1$, and then

$$(\hat{\theta}_{1,s}^{LG}(\tau), \hat{\theta}_{0,s}^{LG}(\tau)) \in \arg \min_{(\theta_1, \theta_0) \in \Theta} Q_n(s, \tau, \theta_1, \theta_0), \text{ with}$$

$$\begin{aligned}
Q_n(s, \tau, \theta_1, \theta_0) &= \frac{1}{n(s)} \sum_{i \in I(s)} \left(\frac{\hat{g}_{1,s}(X_i, \theta_1)}{\hat{f}_1(\hat{q}_1(\tau))} + \frac{\hat{\pi}(s)}{1 - \hat{\pi}(s)} \frac{\hat{g}_{0,s}(X_i, \theta_0)}{\hat{f}_0(\hat{q}_0(\tau))} \right)^2 \\
&+ \frac{2}{n_1(s)} \sum_{i \in I_1(s)} \left(\frac{\hat{g}_{1,s}(X_i, \theta_1)}{\hat{f}_1(\hat{q}_1(\tau))} + \frac{\hat{\pi}(s)}{1 - \hat{\pi}(s)} \frac{\hat{g}_{0,s}(X_i, \theta_0)}{\hat{f}_0(\hat{q}_0(\tau))} \right) \frac{1\{Y_i \leq \hat{q}_1(\tau)\}}{\hat{f}_1(\hat{q}_1(\tau))} \\
&+ \frac{2\hat{\pi}(s)}{n_0(s)(1 - \hat{\pi}(s))} \sum_{i \in I_0(s)} \left(\frac{\hat{g}_{1,s}(X_i, \theta_1)}{\hat{f}_1(\hat{q}_1(\tau))} + \frac{\hat{\pi}(s)}{1 - \hat{\pi}(s)} \frac{\hat{g}_{0,s}(X_i, \theta_0)}{\hat{f}_0(\hat{q}_0(\tau))} \right) \frac{1\{Y_i \leq \hat{q}_0(\tau)\}}{\hat{f}_0(\hat{q}_0(\tau))}. \tag{5.9}
\end{aligned}$$

Assumption 9. Suppose $\Theta_s(\tau) \cap \Theta$ is a singleton, denoted as $(\theta_{1,s}^{LG}(\tau), \theta_{0,s}^{LG}(\tau))$, and

$$\sup_{\tau \in \Upsilon} |f_a(\hat{q}_a(\tau)) - f_a(q_a(\tau))| \xrightarrow{p} 0.$$

Proposition 5.2. Suppose Assumptions 1, 2, 7–9 hold. Then Assumption 6(iii) holds for $(\theta_{a,s}(\tau), \hat{\theta}_{a,s}(\tau)) = (\theta_{a,s}^{LG}(\tau), \hat{\theta}_{a,s}^{LG}(\tau))$, $a = 0, 1, s \in \mathcal{S}, \tau \in \Upsilon$.

In Section D of the Online Supplement, we propose a way to consistently estimate one point in $\Theta_s(\tau) \cap \Theta$ when it is not a singleton.

5.1.3 Quasi Maximum Likelihood Estimation

It is also common to estimate the logistic regression by maximum likelihood. The main goal for the working model is to approximate the true model as closely as possible. It is therefore useful to include additional technical regressors such as interactions in the logistic regression. The set of regressors used is defined as $H_i = H(X_i)$. Let $\hat{\theta}_{a,s}^{ML}(\tau)$ and $\theta_{a,s}^{ML}(\tau)$ be the quasi-ML estimator and its corresponding pseudo true value, respectively, i.e.,

$$\hat{\theta}_{a,s}^{ML}(\tau) = \arg \max_{\theta_a} \frac{1}{n_a(s)} \sum_{i \in I_a(s)} [1\{Y_i \leq \hat{q}_a(\tau)\} \log(\lambda(H_i^\top \theta_a)) + 1\{Y_i > \hat{q}_a(\tau)\} \log(1 - \lambda(H_i^\top \theta_a))] \tag{5.10}$$

and

$$\theta_{a,s}^{ML}(\tau) = \arg \max_{\theta_a} \mathbb{E}[1\{Y_i(a) \leq q_a(\tau)\} \log(\lambda(H_i^\top \theta_a)) + 1\{Y_i(a) > q_a(\tau)\} \log(1 - \lambda(H_i^\top \theta_a)) | S_i = s]. \tag{5.11}$$

We then define

$$\hat{m}_a(\tau, S_i, X_i) = \tau - \sum_{s \in \mathcal{S}} 1\{S_i = s\} \lambda(H_i^\top \hat{\theta}_{a,s}^{ML}(\tau)). \tag{5.12}$$

In addition to the inclusion of technical regressors we allow the pseudo true value to vary across quantiles τ , giving another layer of flexibility of the model. Such a model is called distribution regression and was first proposed by Chernozhukov, Fernández-Val, and Melly (2013). We emphasize here that, although we aim to make the regression model as flexible as possible, our theoretical results do not require the model to be correctly specified.

Assumption 10. *Suppose $\theta_{a,s}^{ML}(\tau)$ is the unique minimizer defined in (5.11) for $a = 0, 1$.*

Theorem 5.3. *Suppose Assumptions 1, 2, 7, 8, 10 hold, then Assumption 6(iii) holds for $(\theta_{a,s}(\tau), \hat{\theta}_{a,s}(\tau)) = (\theta_{a,s}^{ML}(\tau), \hat{\theta}_{a,s}^{ML}(\tau))$, $a = 0, 1, s \in \mathcal{S}, \tau \in \Upsilon$.*

5.1.4 Remarks

From a practical perspective we prefer the linear probability (LP) and quasi maximum likelihood (ML) estimators to the logistic model estimator that minimizes the asymptotic variance of $\hat{q}^{adj}(\tau)$ (LG) for two reasons. First, the LP and ML estimators are free of tuning parameters whereas estimator LG requires nonparametric estimation of the densities, which involves tuning parameters. Second, the estimators LP and ML are easy to compute: the LP estimator has a closed-form expression and the ML estimator is just simple logistic regression which can be implemented in standard statistical softwares. On the other hand, the LG estimator requires non-convex optimization and is therefore computationally costly.

From a theory perspective, we continue to prefer the LP and ML estimators. The LP estimator minimizes (over the class of linear models) not only the asymptotic variance of $\hat{q}^{par}(\tau)$ but also the variance matrix of $(\hat{q}^{par}(\tau_1), \dots, \hat{q}^{par}(\tau_K))$ for any finite-dimension quantile indexes (τ_1, \dots, τ_K) . This implies we can use the same LP estimator for hypothesis testing involving single, multiple, or even a continuum of quantile indexes. In contrast, the LG estimator only aims to minimize the asymptotic variance for a single quantile index, and may not be optimal for tests that involve multiple quantile indexes.³ Although the ML estimator is not guaranteed to be optimal, it is near optimal if the logistic model is close to the true model $m_a(\tau, S_i, X_i)$. To achieve this benefit we suggest including additional technical terms in the regression and allowing the regression coefficients to vary across τ . The next section justifies the use of the ML estimator with a flexible logistic model by letting the number of technical terms (or equivalently, the dimension of H_i) diverge to infinity, showing by this means that the ML estimator can indeed consistently estimate the true model and thereby achieve the minimum covariance matrix of the adjusted QTE estimator, as in Theorem 5.1.

³Section C of the supplement derives the logistic model estimator that minimizes the asymptotic variance of $\hat{q}^{par}(\tau_1) - \hat{q}^{par}(\tau_2)$.

5.2 Nonparametric Method

This section considers nonparametric estimation of $m_a(\tau, s, X_i)$ when the dimension of X_i is fixed as d_x . For ease of notation we assume all coordinates of X_i are continuously distributed. If in an application some elements of X are discrete the dimension d_x is interpreted as the dimension of the continuous covariates. All results in this section can then be extended in a conceptually straightforward manner by using the continuous covariates only within samples that are homogeneous in discrete covariates.

As $m_a(\tau, s, X_i)$ is nonparametrically estimated, we have $\bar{m}_a(\tau, s, X_i) = m_a(\tau, s, X_i) = \tau - \mathbb{P}(Y_i(a) \leq q_a(\tau) | S_i = s, X_i)$. We estimate $\mathbb{P}(Y_i(a) \leq q_a(\tau) | S_i = s, X_i)$ by the sieve method of fitting a logistic model, as studied in Hirano, Imbens, and Ridder (2003). Specifically, recall $\lambda(\cdot)$ is the logistic CDF and denote the number of sieve bases by h_n , which depends on the sample size n and can grow to infinity as $n \rightarrow \infty$. Let $H_{h_n}(x) = (b_{1n}(x), \dots, b_{h_n n}(x))^\top$ where $\{b_{hn}(x)\}_{h \in [h_n]}$ is an h_n dimensional basis of a linear sieve space. We provide more details on the sieve space in Section 6. Denote

$$\hat{m}_a(\tau, s, X_i) = \lambda(H_{h_n}^\top(X_i)\hat{\theta}_{a,s}^{NP}(\tau)) \quad \text{and} \quad (5.13)$$

$$\hat{\theta}_{a,s}^{NP}(\tau) = \arg \max_{\theta_a} \frac{1}{n_a(s)} \sum_{i \in I_a(s)} \left[1\{Y_i \leq \hat{q}_a(\tau)\} \log(\lambda(H_{h_n}^\top(X_i)\theta_a)) + 1\{Y_i > \hat{q}_a(\tau)\} \log(1 - \lambda(H_{h_n}^\top(X_i)\theta_a)) \right] \quad (5.14)$$

Assumption 11. (i) *There exist constants $0 < \kappa_1 < \kappa_2 < \infty$ such that with probability approaching one,*

$$\kappa_1 \leq \lambda_{\min} \left(\frac{1}{n_a(s)} \sum_{i \in I_a(s)} H_{h_n}(X_i) H_{h_n}^\top(X_i) \right) \leq \lambda_{\max} \left(\frac{1}{n_a(s)} \sum_{i \in I_a(s)} H_{h_n}(X_i) H_{h_n}^\top(X_i) \right) \leq \kappa_2$$

and

$$\kappa_1 \leq \lambda_{\min} \left(\mathbb{E}(H_{h_n}(X_i) H_{h_n}^\top(X_i) | S_i = s) \right) \leq \lambda_{\max} \left(\mathbb{E}(H_{h_n}(X_i) H_{h_n}^\top(X_i) | S_i = s) \right) \leq \kappa_2.$$

(ii) *There exists an $h_n \times 1$ vector $\theta_{a,s}^{NP}(\tau)$ such that for $R_a(\tau, s, x) = \mathbb{P}(Y_i(a) \leq q_a(\tau) | S_i = s, X_i = x) - \lambda(H_{h_n}^\top(x)\theta_{a,s}^{NP}(\tau))$, we have $\sup_{a=0,1, s \in \mathcal{S}, \tau \in \Upsilon, x \in \text{Supp}(X)} |R_a(\tau, s, x)| = o(1)$,*

$$\sup_{a=0,1, \tau \in \Upsilon, s \in \mathcal{S}} \frac{1}{n_a(s)} \sum_{i \in I_a(s)} R_a^2(\tau, s, X_i) = O_p \left(\frac{h_n \log(n)}{n} \right),$$

and

$$\sup_{a=0,1,\tau \in \Upsilon, s \in \mathcal{S}} \mathbb{E}(R_a^2(\tau, s, X_i) | S_i = s) = O\left(\frac{h_n \log(n)}{n}\right).$$

(iii) There exists a constant $c \in (0, 0.5)$ such that

$$\begin{aligned} c &\leq \inf_{a=0,1,s \in \mathcal{S}, \tau \in \Upsilon, x \in \text{Supp}(X)} \mathbb{P}(Y_i(a) \leq q_a(\tau) | S_i = s, X_i = x) \\ &\leq \sup_{a=0,1,s \in \mathcal{S}, \tau \in \Upsilon, x \in \text{Supp}(X)} \mathbb{P}(Y_i(a) \leq q_a(\tau) | S_i = s, X_i = x) \leq 1 - c. \end{aligned}$$

(iv) Suppose $\mathbb{E}(H_{h_n, h}^2(X_i) | S_i = s) \leq C < \infty$ for some constant $C > 0$, $\sup_{x \in \text{Supp}(X)} \|H_{h_n}(x)\|_2 \leq \zeta(h_n)$, $\zeta^2(h_n)h_n \log(n) = o(n)$, and $h_n^2 \log^2(n) = o(n)$, where $H_{h_n, h}(X_i)$ denotes the h th coordinate of $H_{h_n}(X_i)$.

Theorem 5.4. Denote $\hat{q}^{NP}(\tau)$ and $\hat{q}^{NP, w}(\tau)$ as the τ th QTE estimator and its multiplier bootstrap counterpart defined in Sections 3 and 4, respectively, with $\bar{m}_a(\tau, S_i, X_i) = m_a(\tau, S_i, X_i)$ and $\hat{m}_a(\tau, S_i, X_i)$ defined in (5.13). Further suppose Assumptions 1, 2, 4, 7, and 11 hold. Then, Assumptions 3 and 5 hold, which further implies that Theorems 3.1 and 4.1 hold for $\hat{q}^{NP}(\tau)$ and $\hat{q}^{NP, w}(\tau)$, respectively. In addition, for any finite-dimensional quantile indexes (τ_1, \dots, τ_K) , the covariance matrix of $(\hat{q}^{NP}(\tau_1), \dots, \hat{q}^{NP}(\tau_K))$ achieves the minimum (in the matrix sense) as characterized in Theorem 3.1.

5.3 Regularization Method

This section considers estimation of $m_a(\tau, s, X)$ in a high-dimensional environment. Let $H_{p_n}(X_i)$ be the regressors with dimension p_n , which may exceed the sample size. When the number of raw controls is comparable to or exceeds the sample size, we can just let $H_{p_n}(X_i) = X_i$. On the other hand, $H_{p_n}(X_i)$ may be composed of a large dictionary of sieve bases derived from a fixed dimensional vector X_i through suitable transformations such as powers and interactions. Thus, high dimensionality in $H_{p_n}(X_i)$ can arise from a desire flexibly approximate nuisance functions. In our approach we follow Belloni et al. (2017) and implement a logistic regression with ℓ_1 -penalization. In their notation we view $m_a(\tau, s, x)$ as a function of $q_a(\tau)$, i.e., $m_a(\tau, s, x) = \tau - \mathcal{M}_a(q_a(\tau), s, x)$, where $\mathcal{M}_a(q, s, x) = \mathbb{P}(Y_i(a) \leq q | S_i = s, X_i = x)$. We estimate $\mathcal{M}_a(q_a(\tau), s, x)$ as $\lambda(H_{p_n}(X_i)^\top \hat{\theta}_{a, s}^{HD}(\hat{q}_a(\tau)))$, where

$$\begin{aligned} \hat{\theta}_{a, s}^{HD}(q) = \arg \min_{\theta_a} \frac{-1}{n_a(s)} \sum_{i \in I_a(s)} &\left[1\{Y_i \leq q\} \log(\lambda(H_{p_n}(X_i)^\top \theta_a)) \right. \\ &\left. + 1\{Y_i > q\} \log(1 - \lambda(H_{p_n}(X_i)^\top \theta_a)) \right] + \frac{\varrho_{n, a}(s)}{n_a(s)} \|\hat{\Omega} \theta_a\|_1, \end{aligned}$$

$\varrho_{n,a}(s)$ is a tuning parameter, and $\hat{\Omega} = \text{diag}(\hat{\omega}_1, \dots, \hat{\omega}_{p_n})$ is a diagonal matrix of data-dependent penalty loadings. We specify $\varrho_{n,a}(s)$ and $\hat{\Omega}$ in Section 6. Post-Lasso estimation is also considered. Let $\hat{\mathbb{S}}_{a,s}(q)$ be the support of $\hat{\theta}_{a,s}^{HD}(q) = \{h \in [p_n] : \hat{\theta}_{a,s,h}^{HD}(q) \neq 0\}$, where $\hat{\theta}_{a,s,h}^{HD}(q)$ is the h th coordinate of $\hat{\theta}_{a,s}^{HD}(q)$. We can complement $\hat{\mathbb{S}}_{a,s}(q)$ with additional variables in $\hat{\mathbb{S}}_{a,s}^+(q)$ that researchers want to control for and define the enlarged set of variables as $\tilde{\mathbb{S}}_{a,s}(q) = \hat{\mathbb{S}}_{a,s}(q) \cup \hat{\mathbb{S}}_{a,s}^+(q)$. We compute the post-Lasso estimator $\hat{\theta}_{a,s}^{post}(q)$ as

$$\hat{\theta}_{a,s}^{post}(q) = \arg \min_{\theta_a \in \tilde{\mathbb{S}}_{a,s}(q)} \frac{-1}{n_a(s)} \sum_{i \in I_a(s)} \left[1\{Y_i \leq q\} \log(\lambda(H_{p_n}(X_i)^\top \theta_a)) + 1\{Y_i > q\} \log(1 - \lambda(H_{p_n}(X_i)^\top \theta_a)) \right]. \quad (5.15)$$

Finally, we compute the auxiliary model as

$$\hat{m}_a(\tau, s, X_i) = \lambda(H_{p_n}^\top(X_i) \hat{\theta}_{a,s}^{HD}(\tau)) \quad \text{or} \quad \hat{m}_a(\tau, s, X_i) = \lambda(H_{p_n}^\top(X_i) \hat{\theta}_{a,s}^{post}(\tau)). \quad (5.16)$$

Assumption 12. (i) Let $\mathcal{Q}_a^\varepsilon = \{q : \inf_{\tau \in \Upsilon} |q - q_a(\tau)| \leq \varepsilon\}$. Suppose $\mathbb{P}(Y_i(a) \leq q | X, S_i = s) = \lambda(H_{p_n}(X_i)^\top \theta_{a,s}^{HD}(q)) + r_a(q, s, X)$ such that $\sup_{a=0,1, q \in \mathcal{Q}_a^\varepsilon, s \in \mathcal{S}} \|\theta_{a,s}^{HD}(q)\|_0 \leq h_n$.

(ii) Suppose $\sup_{i \in [n]} \|H_{p_n}(X_i)\|_\infty \leq \zeta_n$ and $\sup_{h \in [p_n]} \mathbb{E}(H_{p_n,h}^q(X_i) | S_i = s) < \infty$ for $q > 2$.

(iii) Suppose

$$\sup_{a=0,1, q \in \mathcal{Q}_a^\varepsilon, s \in \mathcal{S}} \frac{1}{n_a(s)} \sum_{i \in I_a(s)} r_a^2(q, s, X_i) = O_p(h_n \log(p_n)/n),$$

$$\sup_{a=0,1, q \in \mathcal{Q}_a^\varepsilon, s \in \mathcal{S}} \mathbb{E}(r_a^2(q, s, X_i) | S_i = s) = O(h_n \log(p_n)/n),$$

and

$$\sup_{a=0,1, q \in \mathcal{Q}_a^\varepsilon, s \in \mathcal{S}, x \in \mathcal{X}} |r_a(q, s, X)| = O(\sqrt{\xi_n^2 h_n^2 \log(p_n)/n}).$$

(iv) $\frac{\log(p_n) \xi_n^2 h_n^2}{n} \rightarrow 0$, $\frac{\log^2(p_n) \log^2(n) h_n^2}{n} \rightarrow 0$, $\sup_{a=0,1, q \in \mathcal{Q}_a^\varepsilon, s \in \mathcal{S}} |\hat{\mathbb{S}}_{a,s}^+(q)| = O_p(h_n)$, where $|\hat{\mathbb{S}}_{a,s}^+(q)|$ denotes the number of elements in $\hat{\mathbb{S}}_{a,s}^+(q)$.

(v) There exists a constant $c \in (0, 0.5)$ such that

$$c \leq \inf_{a=0,1, s \in \mathcal{S}, \tau \in \Upsilon, x \in \text{Supp}(X)} \mathbb{P}(Y_i(a) \leq q_a(\tau) | S_i = s, X_i = x)$$

$$\leq \sup_{a=0,1, s \in \mathcal{S}, \tau \in \Upsilon, x \in \text{Supp}(X)} \mathbb{P}(Y_i(a) \leq q_a(\tau) | S_i = s, X_i = x) \leq 1 - c.$$

(vi) Let ℓ_n be a sequence that diverges to infinity. Then, there exist two constants κ_1 and κ_2 such that with probability approaching one,

$$0 < \kappa_1 \leq \inf_{a=0,1, s \in \mathcal{S}, |v|_0 \leq h_n \ell_n} \frac{v^T \left(\frac{1}{n_a(s)} \sum_{i \in I_a(s)} H_{p_n}(X_i) H_{p_n}(X_i)^\top \right) v}{\|v\|_2^2}$$

$$\leq \sup_{a=0,1, s \in \mathcal{S}, |v|_0 \leq h_n \ell_n} \frac{v^T \left(\frac{1}{n_a(s)} \sum_{i \in I_a(s)} H_{p_n}(X_i) H_{p_n}(X_i)^\top \right) v}{\|v\|_2^2} \leq \kappa_2 < \infty,$$

and

$$\begin{aligned} 0 < \kappa_1 &\leq \inf_{a=0,1, s \in \mathcal{S}, |v|_0 \leq h_n \ell_n} \frac{v^T \mathbb{E}(H_{p_n}(X_i) H_{p_n}(X_i)^\top | S_i = s) v}{\|v\|_2^2} \\ &\leq \sup_{a=0,1, s \in \mathcal{S}, |v|_0 \leq h_n \ell_n} \frac{v^T \mathbb{E}(H_{p_n}(X_i) H_{p_n}(X_i)^\top | S_i = s) v}{\|v\|_2^2} \leq \kappa_2 < \infty, \end{aligned}$$

where $|v|_0$ denotes the number of nonzero components in v .

(vii) $\varrho_{n,a}(s) = c\sqrt{n_a(s)}\Phi^{-1}(1 - 0.1/(\log(n_a(s))4p_n))$ where $\Phi(\cdot)$ is the standard normal CDF and $c > 0$ is a constant.

Assumption 12 is standard in the literature and we refer interested readers to Belloni et al. (2017) for more discussion. Assumption 12(i) implies the logistic model is approximately correctly specified. As the approximation is assumed to be sparse, the condition is not innocuous in the high-dimensional setting. As our method is valid even when the auxiliary model is misspecified, we conjecture that Assumption 12(i) can be relaxed, which links to the recent literature on the study of regularized estimation in the high-dimensional setting under misspecification: see, for example, Bradic, Wager, and Zhu (2019) and Tan (2020) and the references therein. An interesting topic for future work is to study misspecification-robust high-dimensional estimators of the conditional probability model and their use to adjust the QTE estimator under CAR based on (3.3) and (3.4).

Theorem 5.5. Denote $\hat{q}^{HD}(\tau)$ and $\hat{q}^{HD,w}(\tau)$ as the τ th QTE estimator and its multiplier bootstrap counterpart defined in Sections 3 and 4, respectively, with $\bar{m}_a(\tau, S_i, X_i) = m_a(\tau, S_i, X_i)$ and $\hat{m}_a(\tau, S_i, X_i)$ defined in (5.16). Further suppose Assumptions 1, 2, 4, 7, and 12 hold. Then, Assumptions 3 and 5 hold, which further imply Theorems 3.1 and 4.1 hold for $\hat{q}^{HD}(\tau)$ and $\hat{q}^{HD,w}(\tau)$, respectively. In addition, for any finite-dimensional quantile indexes (τ_1, \dots, τ_K) , the covariance matrix of $(\hat{q}^{HD}(\tau_1), \dots, \hat{q}^{HD}(\tau_K))$ achieves the minimum (in the matrix sense) as characterized in Theorem 3.1.

6 Practical Guidance and Computation

6.1 Procedures for Estimation and Bootstrap Inference

We can compute $(\hat{q}_1^{adj}, \hat{q}_0^{adj})$ by solving the subgradient conditions of (3.3) and (3.4), respectively. Specifically, we have $(\hat{q}_1^{adj}, \hat{q}_0^{adj}) = (Y_{i_1}, Y_{i_0})$ such that $A_{i_1} = 1, A_{i_0} = 0$,

$$\tau \left(\sum_{i \in [n]} \frac{A_i}{\hat{\pi}(S_i)} \right) - \sum_{i \in [n]} \left(\frac{(A_i - \hat{\pi}(S_i))}{\hat{\pi}(S_i)} \hat{m}_1(\tau, S_i, X_i) \right)$$

$$\begin{aligned}
&\geq \sum_{i \in [n]} \frac{A_i}{\hat{\pi}(S_i)} 1\{Y_i < Y_{i_1}\} \\
&\geq \tau \left(\sum_{i \in [n]} \frac{A_i}{\hat{\pi}(S_i)} \right) - \frac{1}{\hat{\pi}(S_{i_1})} - \sum_{i \in [n]} \left(\frac{(A_i - \hat{\pi}(S_i))}{\hat{\pi}(S_i)} \hat{m}_1(\tau, S_i, X_i) \right), \tag{6.1}
\end{aligned}$$

and

$$\begin{aligned}
&\tau \left(\sum_{i \in [n]} \frac{1 - A_i}{1 - \hat{\pi}(S_i)} \right) + \sum_{i \in [n]} \left(\frac{(A_i - \hat{\pi}(S_i))}{1 - \hat{\pi}(S_i)} \hat{m}_0(\tau, S_i, X_i) \right) \\
&\geq \sum_{i \in [n]} \frac{1 - A_i}{1 - \hat{\pi}(S_i)} 1\{Y_i < Y_{i_0}\} \\
&\geq \tau \left(\sum_{i \in [n]} \frac{1 - A_i}{1 - \hat{\pi}(S_i)} \right) - \frac{1}{1 - \hat{\pi}(S_{i_0})} + \sum_{i \in [n]} \left(\frac{(A_i - \hat{\pi}(S_i))}{1 - \hat{\pi}(S_i)} \hat{m}_0(\tau, S_i, X_i) \right). \tag{6.2}
\end{aligned}$$

We note (i_1, i_0) are uniquely defined as long as all the inequalities in (6.1) and (6.2) are strict, which is usually the case. If we have

$$\tau \left(\sum_{i \in [n]} \frac{A_i}{\hat{\pi}(S_i)} \right) - \sum_{i \in [n]} \left(\frac{(A_i - \hat{\pi}(S_i))}{\hat{\pi}(S_i)} \hat{m}_1(\tau, S_i, X_i) \right) = \sum_{i \in [n]} \frac{A_i}{\hat{\pi}(S_i)} 1\{Y_i \leq Y_{i_1}\},$$

then both i_1 and i'_1 satisfy (6.1), where i'_1 is the index such that $A_{i'_1} = 1$ and $Y_{i'_1}$ is the smallest observation in the treatment group that is larger than Y_{i_1} . In this case, we let $\hat{q}_1^{adj} = Y_{i_1}$. Similarly, by solving the subgradient conditions of (4.1) and (4.2), we have $(\hat{q}_1^w, \hat{q}_0^w) = (Y_{i_1^w}, Y_{i_0^w})$ such that $A_{i_1^w} = 1, A_{i_0^w} = 0$,

$$\begin{aligned}
&\tau \left(\sum_{i \in [n]} \frac{\xi_i A_i}{\hat{\pi}^w(S_i)} \right) - \sum_{i \in [n]} \left(\frac{\xi_i (A_i - \hat{\pi}^w(S_i))}{\hat{\pi}^w(S_i)} \hat{m}_1(\tau, S_i, X_i) \right) \\
&\geq \sum_{i \in [n]} \frac{\xi_i A_i}{\hat{\pi}^w(S_i)} 1\{Y_i < Y_{i_1^w}\} \\
&\geq \tau \left(\sum_{i \in [n]} \frac{\xi_i A_i}{\hat{\pi}^w(S_i)} \right) - \frac{\xi_{i_1^w}}{\hat{\pi}^w(S_{i_1^w})} - \sum_{i \in [n]} \left(\frac{\xi_i (A_i - \hat{\pi}^w(S_i))}{\hat{\pi}^w(S_i)} \hat{m}_1(\tau, S_i, X_i) \right) \tag{6.3}
\end{aligned}$$

and

$$\begin{aligned}
&\tau \left(\sum_{i \in [n]} \frac{\xi_i (1 - A_i)}{1 - \hat{\pi}^w(S_i)} \right) + \sum_{i \in [n]} \left(\frac{\xi_i (A_i - \hat{\pi}^w(S_i))}{1 - \hat{\pi}^w(S_i)} \hat{m}_0(\tau, S_i, X_i) \right) \\
&\geq \sum_{i \in [n]} \frac{\xi_i (1 - A_i)}{1 - \hat{\pi}^w(S_i)} 1\{Y_i < Y_{i_0^w}\}
\end{aligned}$$

$$\geq \tau \left(\sum_{i \in [n]} \frac{\xi_i(1 - A_i)}{1 - \hat{\pi}^w(S_i)} \right) - \frac{\xi_{i_0}^w}{1 - \hat{\pi}^w(S_{i_0})} + \sum_{i \in [n]} \left(\frac{\xi_i(A_i - \hat{\pi}^w(S_i))}{1 - \hat{\pi}^w(S_i)} \hat{m}_0(\tau, S_i, X_i) \right). \quad (6.4)$$

The inequalities in (4.1) and (4.2) are strict with probability one if ξ_i is continuously distributed. In this case, $(\hat{q}_1^w, \hat{q}_0^w)$ is uniquely defined with probability one.

We summarize the steps in the bootstrap procedure as follows.

1. Let \mathcal{G} be a set of quantile indexes. For $\tau \in \mathcal{G}$, compute $\hat{q}_1(\tau)$ and $\hat{q}_0(\tau)$ following (6.1) and (6.2) with $\hat{m}_1(\tau, S_i, X_i)$ and $\hat{m}_0(\tau, S_i, X_i)$ replaced by zero.
2. Compute $\hat{m}_a(\tau, S_i, X_i)$ for $a = 0, 1$ and $\tau \in \mathcal{G}$ using $\hat{q}_1(\tau)$ and $\hat{q}_0(\tau)$.
3. Compute the original estimator $\hat{q}^{adj}(\tau) = \hat{q}_1^{adj}(\tau) - \hat{q}_0^{adj}(\tau)$, following (6.1) and (6.2) for $\tau \in \mathcal{G}$.
4. Let B be the number of bootstrap replications. For $b \in [B]$, generate $\{\xi_i\}_{i \in [n]}$. Compute $\hat{q}^{w,b}(\tau) = \hat{q}_1^{w,b}(\tau) - \hat{q}_0^{w,b}(\tau)$ for $\tau \in \mathcal{G}$ following (6.3) and (6.4). Obtain $\{\hat{q}^{w,b}(\tau)\}_{\tau \in \mathcal{G}}$.
5. Repeat the above step for $b \in [B]$ and obtain B bootstrap estimators of the QTE, denoted as $\{\hat{q}^{w,b}(\tau)\}_{b \in [B], \tau \in \mathcal{G}}$.

6.2 Bootstrap Confidence Intervals

Given the bootstrap estimates, we discuss how to conduct bootstrap inference for the null hypotheses with single, multiple, and a continuum of quantile indexes.

Case (1). We test the single null hypothesis that $\mathcal{H}_0 : q(\tau) = \underline{q}$ vs. $q(\tau) \neq \underline{q}$. Set $\mathcal{G} = \{\tau\}$ in the procedures described above and let $\hat{\mathcal{C}}(\nu)$ and $\mathcal{C}(\nu)$ be the ν th empirical quantile of the sequence $\{\hat{q}^{w,b}(\tau)\}_{b \in [B]}$ and the ν th standard normal critical value, respectively. Let $\alpha \in (0, 1)$ be the significance level. We suggest using the bootstrap estimator to construct the standard error of $\hat{q}^{adj}(\tau)$ as $\hat{\sigma} = \frac{\hat{\mathcal{C}}(0.975) - \hat{\mathcal{C}}(0.025)}{\mathcal{C}(0.975) - \mathcal{C}(0.025)}$. Then the valid confidence interval and Wald test using this standard error are

$$CI(\alpha) = (\hat{q}^{adj}(\tau) + \mathcal{C}(\alpha/2)\hat{\sigma}, \hat{q}^{adj}(\tau) + \mathcal{C}(1 - \alpha/2)\hat{\sigma}),$$

and $1\left\{\left|\frac{\hat{q}^{adj}(\tau) - \underline{q}}{\hat{\sigma}}\right| \geq \mathcal{C}(1 - \alpha/2)\right\}$, respectively.⁴

Case (2). We test the null hypothesis that $\mathcal{H}_0 : q(\tau_1) - q(\tau_2) = \underline{q}$ vs. $q(\tau_1) - q(\tau_2) \neq \underline{q}$. In this case, we have $\mathcal{G} = \{\tau_1, \tau_2\}$ in the procedure described in Section 6.1. Further, let $\hat{\mathcal{C}}(\nu)$ be the ν th empirical quantile of the sequence $\{\hat{q}^{w,b}(\tau_1) - \hat{q}^{w,b}(\tau_2)\}_{b \in [B]}$, and let $\alpha \in (0, 1)$ be the significance

⁴It is asymptotically valid to use standard and percentile bootstrap confidence intervals. In simulations we found that the confidence interval proposed in the paper has better finite-sample performance.

level. We suggest using the bootstrap standard error to construct the valid confidence interval and Wald test as

$$CI(\alpha) = (\hat{q}^{adj}(\tau_1) - \hat{q}^{adj}(\tau_2) + \mathcal{C}(\alpha/2)\hat{\sigma}, \hat{q}^{adj}(\tau_1) - \hat{q}^{adj}(\tau_2) + \mathcal{C}(1 - \alpha/2)\hat{\sigma}),$$

and $1\left\{\left|\frac{\hat{q}^{adj}(\tau_1) - \hat{q}^{adj}(\tau_2) - q}{\hat{\sigma}}\right| \geq \mathcal{C}(1 - \alpha/2)\right\}$, respectively, where $\hat{\sigma} = \frac{\widehat{\mathcal{C}}(0.975) - \widehat{\mathcal{C}}(0.025)}{\mathcal{C}(0.975) - \mathcal{C}(0.025)}$.

Case (3). We test the null hypothesis that

$$\mathcal{H}_0 : q(\tau) = \underline{q}(\tau) \quad \forall \tau \in \Upsilon \text{ vs. } q(\tau) \neq \underline{q}(\tau) \quad \exists \tau \in \Upsilon.$$

In theory, we should let $\mathcal{G} = \Upsilon$. In practice, we let $\mathcal{G} = \{\tau_1, \dots, \tau_G\}$ be a fine grid of Υ where G should be as large as computationally possible. Further, let $\widehat{\mathcal{C}}_\tau(\nu)$ denote the ν th empirical quantile of the sequence $\{\hat{q}^{w,b}(\tau)\}_{b \in [B]}$ for $\tau \in \mathcal{G}$. Compute the standard error of $\hat{q}^{adj}(\tau)$ as

$$\hat{\sigma}_\tau = \frac{\widehat{\mathcal{C}}_\tau(0.975) - \widehat{\mathcal{C}}_\tau(0.025)}{\mathcal{C}(0.975) - \mathcal{C}(0.025)}.$$

The uniform confidence band with an α significance level is constructed as

$$CB(\alpha) = \{\hat{q}^{adj}(\tau) - \widetilde{\mathcal{C}}_\alpha \hat{\sigma}_\tau, \hat{q}^{adj}(\tau) + \widetilde{\mathcal{C}}_\alpha \hat{\sigma}_\tau : \tau \in \mathcal{G}\},$$

where the critical value $\widetilde{\mathcal{C}}_\alpha$ is computed as

$$\widetilde{\mathcal{C}}_\alpha = \inf \left\{ z : \frac{1}{B} \sum_{b=1}^B 1 \left\{ \sup_{\tau \in \mathcal{G}} \left| \frac{\hat{q}^{w,b}(\tau) - \tilde{q}(\tau)}{\hat{\sigma}_\tau} \right| \leq z \right\} \geq 1 - \alpha \right\},$$

and $\tilde{q}(\tau)$ is first-order equivalent to $\hat{q}^{adj}(\tau)$ in the sense that $\sup_{\tau \in \Upsilon} |\tilde{q}(\tau) - \hat{q}^{adj}(\tau)| = o_p(1/\sqrt{n})$. We suggest choosing $\tilde{q}(\tau) = \widehat{\mathcal{C}}_\tau(0.5)$ over other choices such as $\tilde{q}(\tau) = \hat{q}^{adj}(\tau)$ due to its better finite-sample performance. We reject \mathcal{H}_0 at significance level α if $q(\cdot) \notin CB(\alpha)$.

6.3 Computation of Auxiliary Regressions

Parametric Regressions. For the linear probability model, we compute the LP estimator via (5.7). For the logistic model, we consider both LG and ML estimators. First, we compute the LG estimator $\hat{\theta}_{a,s}^{LG}(\tau)$ as in (5.9), which minimizes the sample variance of our adjusted QTE estimator $\hat{q}^{adj}(\tau)$.⁵ Second, we propose to compute the ML estimator $\hat{\theta}_{a,s}^{ML}(\tau)$ as in (5.11), which is the quasi maximum likelihood estimator of a flexible distribution regression.

⁵When the null hypothesis in Case (2) is tested, the optimal pseudo true value is defined to minimize the asymptotic variance of $\hat{q}^{par}(\tau_1) - \hat{q}^{par}(\tau_2)$, which is not necessarily equal to the pseudo true values that minimize the asymptotic variances of $\hat{q}^{par}(\tau_1)$ and $\hat{q}^{par}(\tau_2)$, i.e., $(\theta_{a,s}^{LG}(\tau_1), \theta_{a,s}^{LG}(\tau_2))$. We propose a way to compute the new pseudo true value in Section C.

Sieve Logistic Regressions. We provide more details on the sieve basis. Recall $H_{h_n}(x) \equiv (b_{1n}(x), \dots, b_{h_n n}(x))^\top$, where $\{b_{hn}(\cdot)\}_{h \in [h_n]}$ are h_n basis functions of a linear sieve space, denoted as \mathbb{B} . Given that all d_x elements of X are continuously distributed, the sieve space \mathbb{B} can be constructed as follows.

1. For each element $X^{(l)}$ of X , $l = 1, \dots, d_x$, let \mathcal{B}_l be the univariate sieve space of dimension J_n . One example of \mathcal{B}_l is the linear span of the J_n dimensional polynomials given by

$$\mathbb{B}_l = \left\{ \sum_{k=0}^{J_n} \alpha_k x^k, x \in \text{Supp}(X^{(l)}), \alpha_k \in \mathfrak{R} \right\};$$

Another is the linear span of r -order splines with J_n nodes given by

$$\mathbb{B}_l = \left\{ \sum_{k=0}^{r-1} \alpha_k x^k + \sum_{j=1}^{J_n} b_j [\max(x - t_j, 0)]^{r-1}, x \in \text{Supp}(X^{(l)}), \alpha_k, b_j \in \mathfrak{R} \right\},$$

where the grid $-\infty = t_0 \leq t_1 \leq \dots \leq t_{J_n} \leq t_{J_n+1} = \infty$ partitions $\text{Supp}(X^{(l)})$ into $J_n + 1$ subsets $I_j = [t_j, t_{j+1}) \cap \text{Supp}(X^{(l)})$, $j = 1, \dots, J_n - 1$, $I_0 = (t_0, t_1) \cap \text{Supp}(X^{(l)})$, and $I_{J_n} = (t_{J_n}, t_{J_n+1}) \cap \text{Supp}(X^{(l)})$.

2. Let \mathbb{B} be the tensor product of $\{\mathcal{B}_l\}_{l=1}^{d_x}$, which is defined as a linear space spanned by the functions $\prod_{l=1}^{d_x} g_l$, where $g_l \in \mathcal{B}_l$. The dimension of \mathbb{B} is then $K \equiv d_x J_n$.

We refer interested readers to Hirano et al. (2003) and Chen (2007) for more details about the implementation of sieve estimation. Given the sieve basis, we can compute the $\hat{m}_a(\tau, s, X_i)$ following (5.13) and (5.14).

Logistic Regressions with an ℓ_1 Penalization. We follow the estimation procedure and the choice of tuning parameter proposed by Belloni et al. (2017). We provide details below for completeness. Recall $\varrho_{n,a}(s) = c\sqrt{n_a(s)}\Phi^{-1}(1 - 1/(p_n \log(n_a(s))))$. We set $c = 1.1$ following Belloni et al. (2017). We then implement the following algorithm to estimate $\hat{\theta}_{a,s}^{HD}(\tau)$ for $\tau \in \mathfrak{T}$:

- (i) Let $\hat{\sigma}_h^{(0)} = \frac{1}{n_a(s)} \sum_{i \in I_a(s)} (1\{Y_i \leq \hat{q}_a(\tau)\} - \bar{Y}_{a,s}(\tau))^2 H_{p_n, h}^2$ for $h \in [p_n]$, where $\bar{Y}_{a,s}(\tau) = \frac{1}{n_a(s)} \sum_{i \in I_a(s)} 1\{Y_i \leq \hat{q}_a(\tau)\}$. Estimate

$$\begin{aligned} \hat{\theta}_{a,s}^{HD,0}(\tau) = \arg \min_{\theta_a} \frac{-1}{n_a(s)} \sum_{i \in I_a(s)} \left[1\{Y_i \leq q\} \log(\lambda(H_{p_n}(X_i)^\top \theta_a)) \right. \\ \left. + 1\{Y_i > q\} \log(1 - \lambda(H_{p_n}(X_i)^\top \theta_a)) \right] + \frac{\varrho_{n,a}(s)}{n_a(s)} \sum_{h \in [p_n]} \hat{\sigma}_h^{(0)} |\theta_{a,h}|. \end{aligned}$$

- (ii) For $k = 1, \dots, K$, obtain $\hat{\sigma}_h^{(k)} = \sqrt{\frac{1}{n} \sum_{i \in [n]} (H_{p_n, h} \hat{\varepsilon}_i^{(k)})^2}$, where $\hat{\varepsilon}_i^{(k)} = 1\{Y_i \leq \hat{q}_a(\tau)\} -$

$\lambda(H_{p_n}^T \hat{\theta}_{a,s}^{HD,k-1}(\tau))$. Estimate

$$\hat{\theta}_{a,s}^{HD,k}(\tau) = \arg \min_{\theta_a} \frac{-1}{n_a(s)} \sum_{i \in I_a(s)} \left[1\{Y_i \leq q\} \log(\lambda(H_{p_n}(X_i)^\top \theta_a)) + 1\{Y_i > q\} \log(1 - \lambda(H_{p_n}(X_i)^\top \theta_a)) \right] + \frac{\varrho_{n,a}(s)}{n_a(s)} \sum_{h \in [p_n]} \hat{\sigma}_h^{(k)} |\theta_{a,h}|.$$

(iii) Let $\hat{\theta}_{a,s}^{HD}(\tau) = \hat{\theta}_{a,s}^{HD,K}(\tau)$.

(iv) Repeat the above procedure for $\tau \in \mathcal{G}$.

7 Simulation

7.1 Data Generating Processes

Three data generating processes (DGPs) are used to assess the finite-sample performance of the estimation and inference methods introduced in the paper. We consider the outcomes equation

$$Y_i = \alpha(X_i) + \gamma Z_i + \mu(X_i) A_i + \eta_i, \quad (7.1)$$

where $\gamma = 4$ for all cases while $\alpha(X_i)$, $\mu(X_i)$, and η_i are separately specified as follows.

- (i) Let Z be standardized Beta(2, 2) distributed, $S_i = \sum_{j=1}^4 1\{Z_i \leq g_j\}$, and $(g_1, \dots, g_4) = (-0.25\sqrt{20}, 0, 0.25\sqrt{20}, 0.5\sqrt{20})$. X_i contains two covariates $(X_{1i}, X_{2i})^\top$, where X_{1i} follows a uniform distribution on $[-2, 2]$, X_{2i} follows a standard normal distribution, and X_{1i} and X_{2i} are independent. Further define $\alpha(X_i) = 1 + X_{2i}$, $\mu(X_i) = 1 + X_i^\top \beta$, $\beta = (3, 3)^\top$, and $\eta_i = (0.25 + X_{1i}^2) A_i \varepsilon_{1i} + (1 - A_i) \varepsilon_{2i}$, where $(\varepsilon_{1i}, \varepsilon_{2i})$ are jointly standard normal.
- (ii) Let Z be uniformly distributed on $[-2, 2]$, $S_i = \sum_{j=1}^4 1\{Z_i \leq g_j\}$, and $(g_1, \dots, g_4) = (-1, 0, 1, 2)$. Let $X_i = (X_{1i}, X_{2i})^\top$ be the same as defined in DGP (i). Further define $\alpha(X_i) = 1 + X_{1i} + X_{2i}$, $\mu(X_i) = 1 + X_{1i} + X_{2i} + \frac{1}{4}(X_i^\top \beta)^2$ with $\beta = (2, 2)^\top$, and $\eta_i = 2(1 + Z_i^2) A_i \varepsilon_{1i} + (1 + Z_i^2)(1 - A_i) \varepsilon_{2i}$, where $(\varepsilon_{1i}, \varepsilon_{2i})$ are mutually independently $T(5)/\sqrt{5}$ distributed.
- (iii) Let Z be standardized Beta(2, 2) distributed, $S_i = \sum_{j=1}^4 1\{Z_i \leq g_j\}$, and $(g_1, \dots, g_4) = (-0.25\sqrt{20}, 0, 0.25\sqrt{20}, 0.5\sqrt{20})$. Further suppose that X_i contains twenty covariates $(X_{1i}, \dots, X_{20,i})^\top$, where $X = \Phi(Z)$ with $Z \sim N(0_{20 \times 1}, \Omega)$ and the variance-covariance matrix Ω is the Toeplitz

matrix

$$\Omega = \begin{pmatrix} 0.5 & 0.5^2 & 0.5^3 & \dots & 0.5^{19} \\ 0.5^2 & 0.5 & 0.5^2 & \dots & 0.5^{18} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0.5^{19} & 0.5^{18} & 0.5^{17} & \dots & 0.5 \end{pmatrix}$$

Further define $\alpha(X_i) = 1$, $\mu(X_i) = 1 + \sum_{k=1}^{20} X_{ki}\beta_k$ with $\beta_k = 4/k^2$, and $\eta_i = 2A_i\varepsilon_{1i} + (1 - A_i)\varepsilon_{2i}$, where $(\varepsilon_{1i}, \varepsilon_{2i})$ are jointly standard normal.

For each data generating process, we consider the following four randomization schemes as in Zhang and Zheng (2020) with $\pi(s) = 0.5$ for $s \in \mathcal{S}$:

- (i) SRS: Treatment assignment is generated as in Example 1.
- (ii) WEI: Treatment assignment is generated as in Example 2 with $\phi(x) = (1 - x)/2$.
- (iii) BCD: Treatment assignment is generated as in Example 3 with $\lambda = 0.75$.
- (iv) SBR: Treatment assignment is generated as in Example 4.

We assess the empirical size and power of the tests for $n = 200$ and $n = 400$. All simulations are replicated 10,000 times with bootstrap sample size being 1,000. We compute the true QTEs or QTE differences by simulations with 10,000 sample size and 1,000 replications. To compute power, we perturb the true values by 1.5 for DGPs (i) and (ii) and 2 for DGP (iii). As discussed in Section 6, we examine three null hypotheses:

- (i) Pointwise test

$$H_0 : q(\tau) = \text{truth} \quad \text{v.s.} \quad H_1 : q(\tau) = \text{truth} + \text{perturbation}, \quad \tau = 0.25, 0.5, 0.75;$$

- (ii) Test for the difference

$$H_0 : q(0.75) - q(0.25) = \text{truth} \quad \text{v.s.} \quad H_1 : q(0.75) - q(0.25) = \text{truth} + \text{perturbation};$$

- (iii) Uniform test

$$H_0 : q(\tau) = \text{truth}(\tau) \quad \text{v.s.} \quad H_1 : q(\tau) = \text{truth}(\tau) + \text{perturbation}, \quad \tau \in [0.25, 0.75].$$

For the pointwise test, we report the results for the median ($\tau = 0.5$) in the main text and give the cases $\tau = 0.25$ and $\tau = 0.75$ to the Online Supplement.

7.2 Estimation Methods

For DGPs (i) and (ii), we consider the following estimation methods of the auxiliary regression.

- (i) NA: the estimator without any adjustments, i.e., setting $\hat{m}_a(\cdot) = \overline{m}_a(\cdot) = 0$.
- (ii) LP: the linear probability model with regressors X_i and the pseudo true value is estimated by $\hat{\theta}_{a,s}^{LP}(\tau)$ defined in (5.7).
- (iii) LG: the logistic model with regressors $(1, X_{1i}, X_{2i})$ and the pseudo true value is estimated by $\hat{\theta}_{a,s}^{LG}(\tau)$ defined in (5.9).
- (iv) ML: the logistic model with regressor $H_i = (1, X_{1i}, X_{2i}, X_{1i}X_{2i})^\top$ and the pseudo true value is estimated by $\hat{\theta}_{a,s}^{ML}(\tau)$ defined in (5.10).
- (v) NP: the logistic model with regressor $H_{h_n}(X_i) = (1, X_{1i}, X_{2i}, X_{1i}X_{2i}, X_{1i}1\{X_{1i} > t_1\}X_{2i}1\{X_{2i} > t_2\})^\top$ where t_1 and t_2 are the sample medians of $\{X_{1i}\}_{i \in [n]}$ and $\{X_{2i}\}_{i \in [n]}$, respectively. The pseudo true value is estimated by $\hat{\theta}_{a,s}^{NP}(\tau)$ defined in (5.14).

For DGP (iii), we consider the post-Lasso estimator $\hat{\theta}_{a,s}^{post}(\tau)$ as defined in (5.15) with $H_{p_n}(X_i) = (1, X_i^\top)^\top$ and $\hat{S}_{a,s}^+(q) = \{2\}$. The choice of tuning parameter and the estimation procedure are detailed in Section 6.3.

7.3 Simulation Results

Table 2 presents the empirical size and power for the pointwise test with $\tau = 0.5$ under DGPs (i) and (ii). We make five observations. First, none of the auxiliary models is correctly specified but test sizes are all close to the nominal level 5%, confirming that estimation and inference are robust to misspecification. Second, the inclusion of auxiliary regressions improves efficiency of the QTE estimator as the powers for method “NA” are the smallest among all the methods for both DGPs and all randomization schemes. This finding is consistent with theory because methods “LP”, “LG”, “NP” are guaranteed to be weakly more efficient than “NA”. Method “ML” fits a flexible distribution regression which can approximate the true DGP well. Therefore, its performance is close to the method “NP” and, thus, better than “NA”. Third, the powers of method “NP” are better than those for “LP”, “LG”, and “ML” because this method estimates the true specification and achieves the minimum asymptotic variance of $\hat{q}^{adj}(\tau)$ as shown in Theorem 5.1. Fourth, from a theory perspective the performance of methods “LP” and “LG” is not comparable as both are misspecified. In simulations the powers of method “LG” are worse for two reasons. First, computation of the objective function requires nonparametric estimation of the densities, which requires tuning

parameters⁶ and has a slow convergence rate. Second, since the objective function is not convex, the computation of the optimal values $\hat{\theta}_{a,s}^{LG}(\tau)$ is sometimes unstable.⁷ Fifth, when the sample size is 200, the method “NP” slightly overrejects but size becomes closer to nominal when the sample size increases to 400.

Table 2: Pointwise Test ($\tau = 0.5$)

Methods	$N = 200$				$N = 400$			
	SRS	WEI	BCD	SBR	SRS	WEI	BCD	SBR
Panel A: DGP (i)								
<i>A.1: Size</i>								
NA	0.051	0.052	0.050	0.054	0.055	0.055	0.051	0.054
LP	0.050	0.047	0.055	0.053	0.050	0.052	0.053	0.054
LG	0.060	0.063	0.061	0.057	0.058	0.058	0.060	0.058
ML	0.060	0.058	0.060	0.055	0.051	0.055	0.050	0.052
NP	0.059	0.061	0.063	0.066	0.053	0.055	0.055	0.054
<i>A.2: Power</i>								
NA	0.403	0.402	0.405	0.402	0.681	0.675	0.679	0.681
LP	0.495	0.495	0.499	0.494	0.791	0.791	0.788	0.782
LG	0.466	0.478	0.462	0.466	0.760	0.757	0.769	0.756
ML	0.508	0.511	0.533	0.511	0.808	0.813	0.810	0.809
NP	0.526	0.537	0.535	0.530	0.813	0.815	0.813	0.809
Panel B: DGP (ii)								
<i>B.1: Size</i>								
NA	0.049	0.049	0.044	0.045	0.051	0.047	0.047	0.046
LP	0.049	0.049	0.048	0.049	0.053	0.052	0.051	0.045
LG	0.055	0.056	0.055	0.058	0.054	0.054	0.054	0.057
ML	0.062	0.062	0.057	0.058	0.055	0.053	0.055	0.056
NP	0.070	0.063	0.059	0.066	0.059	0.057	0.055	0.057
<i>B.2: Power</i>								
NA	0.490	0.494	0.497	0.485	0.779	0.772	0.768	0.762
LP	0.579	0.567	0.569	0.573	0.851	0.850	0.852	0.852
LG	0.524	0.526	0.524	0.521	0.819	0.823	0.821	0.819
ML	0.612	0.601	0.609	0.623	0.868	0.875	0.877	0.871
NP	0.612	0.618	0.612	0.618	0.878	0.876	0.880	0.876

Tables 3 and 4 present sizes and powers of inference on $q(0.75) - q(0.25)$ and on $q(\tau)$ uniformly over $\tau \in [0.25, 0.75]$, respectively, for DGPs (i) and (ii) and four randomization schemes. All the observations made above apply to these results.

⁶We use the Gaussian kernel and bandwidth $1.06\hat{\sigma}_{a,s}(n_a(s))^{1/5}$ where $\hat{\sigma}_{a,s}$ and $n_a(s)$ are the standard error and sample size of the observations $\{Y_i\}_{i \in I_a(s)}$, respectively, for $a = 0, 1$ and $s \in \mathcal{S}$.

⁷In simulations we added a ridge penalty $\log^{-1}(n(s))\|\theta\|_2^2$ to the objective function to stabilize the solution. Such a penalty term does not affect the proof of consistency of $\hat{\theta}_{a,s}^{LG}(\tau)$ because it will vanish as $n \rightarrow \infty$.

Table 3: Test for Differences ($\tau_1 = 0.25, \tau_2 = 0.75$)

Methods	$N = 200$				$N = 400$			
	SRS	WEI	BCD	SBR	SRS	WEI	BCD	SBR
Panel A: DGP (i)								
<i>A.1: Size</i>								
NA	0.039	0.039	0.041	0.043	0.043	0.045	0.042	0.045
LP	0.046	0.045	0.045	0.047	0.045	0.047	0.046	0.047
LG	0.047	0.052	0.042	0.046	0.045	0.050	0.046	0.046
ML	0.048	0.047	0.048	0.047	0.048	0.047	0.047	0.048
NP	0.050	0.048	0.049	0.050	0.050	0.049	0.046	0.053
<i>A.2: Power</i>								
NA	0.193	0.194	0.201	0.218	0.369	0.376	0.377	0.397
LP	0.211	0.201	0.207	0.207	0.400	0.395	0.401	0.398
LG	0.227	0.235	0.237	0.236	0.462	0.468	0.457	0.453
ML	0.245	0.242	0.242	0.241	0.464	0.454	0.455	0.455
NP	0.240	0.239	0.246	0.240	0.455	0.464	0.462	0.464
Panel B: DGP (ii)								
<i>B.1: Size</i>								
NA	0.040	0.038	0.040	0.042	0.042	0.043	0.044	0.043
LP	0.042	0.044	0.041	0.048	0.045	0.047	0.043	0.041
LG	0.047	0.046	0.046	0.048	0.050	0.046	0.052	0.052
ML	0.051	0.052	0.049	0.051	0.051	0.049	0.048	0.049
NP	0.060	0.057	0.053	0.050	0.050	0.053	0.056	0.051
<i>B.2: Power</i>								
NA	0.207	0.208	0.212	0.237	0.403	0.401	0.397	0.423
LP	0.236	0.241	0.244	0.233	0.444	0.444	0.443	0.433
LG	0.236	0.241	0.244	0.233	0.444	0.444	0.443	0.433
ML	0.255	0.261	0.267	0.252	0.471	0.475	0.467	0.475
NP	0.266	0.265	0.267	0.267	0.465	0.480	0.467	0.477

Table 4: Uniform Test ($\tau \in [0.25, 0.75]$)

Methods	$N = 200$				$N = 400$			
	SRS	WEI	BCD	SBR	SRS	WEI	BCD	SBR
Panel A: DGP (i)								
<i>A.1: Size</i>								
NA	0.044	0.047	0.043	0.048	0.047	0.047	0.046	0.046
LP	0.047	0.047	0.047	0.044	0.048	0.045	0.046	0.047
LG	0.049	0.048	0.051	0.051	0.053	0.056	0.053	0.053
ML	0.050	0.051	0.051	0.050	0.051	0.055	0.050	0.049
NP	0.059	0.054	0.058	0.052	0.052	0.053	0.053	0.054
<i>A.2: Power</i>								
NA	0.448	0.459	0.456	0.448	0.771	0.769	0.774	0.772
LP	0.594	0.580	0.599	0.596	0.895	0.897	0.904	0.899
LG	0.567	0.559	0.582	0.568	0.883	0.887	0.882	0.883
ML	0.614	0.621	0.616	0.613	0.918	0.916	0.923	0.918
NP	0.630	0.630	0.634	0.639	0.920	0.919	0.921	0.917
Panel B: DGP (ii)								
<i>B.1: Size</i>								
NA	0.035	0.039	0.039	0.038	0.044	0.044	0.042	0.044
LP	0.040	0.046	0.041	0.040	0.042	0.044	0.045	0.042
LG	0.050	0.049	0.047	0.044	0.051	0.056	0.050	0.055
ML	0.055	0.053	0.055	0.052	0.052	0.053	0.049	0.054
NP	0.066	0.061	0.055	0.062	0.053	0.056	0.053	0.055
<i>B.2: Power</i>								
NA	0.560	0.572	0.570	0.582	0.876	0.878	0.874	0.880
LP	0.700	0.705	0.707	0.707	0.948	0.955	0.954	0.951
LG	0.647	0.654	0.654	0.650	0.936	0.939	0.936	0.935
ML	0.765	0.770	0.770	0.771	0.973	0.974	0.974	0.977
NP	0.771	0.770	0.769	0.778	0.971	0.972	0.975	0.976

Next, we assess the empirical size and power of all three testing scenarios under the high-dimensional setting. In Table 5, we compare the methods “NA” and our post-Lasso estimator. In general, all sizes for both methods approach the nominal level as the sample size increases. The post-Lasso method dominates “NA” in all tests with superior power performance. This result is consistent with theory given in Theorem 5.5.

Table 5: Empirical Size and Power for High-dimensional Covariates (DGP 3)

	$N = 200$				$N = 400$			
	SRS	WEI	BCD	SBR	SRS	WEI	BCD	SBR
Panel A: Size								
<i>A.1. NA</i>								
$\tau = 0.25$	0.043	0.044	0.043	0.043	0.041	0.047	0.044	0.046
$\tau = 0.5$	0.045	0.046	0.049	0.046	0.051	0.043	0.046	0.045
$\tau = 0.75$	0.046	0.047	0.049	0.049	0.049	0.046	0.046	0.052
Diff(0.25-0.75)	0.042	0.041	0.038	0.042	0.043	0.045	0.043	0.041
Uniform	0.045	0.044	0.048	0.043	0.049	0.055	0.050	0.052
<i>A.2. Post-Lasso</i>								
$\tau = 0.25$	0.055	0.057	0.058	0.057	0.049	0.048	0.055	0.055
$\tau = 0.5$	0.061	0.056	0.060	0.058	0.057	0.047	0.050	0.051
$\tau = 0.75$	0.056	0.062	0.058	0.052	0.053	0.052	0.051	0.053
Diff(0.25-0.75)	0.055	0.051	0.052	0.053	0.051	0.047	0.045	0.048
Uniform	0.070	0.066	0.066	0.070	0.060	0.064	0.058	0.055
Panel B: Power								
<i>B.1. NA</i>								
$\tau = 0.25$	0.399	0.403	0.409	0.410	0.689	0.690	0.698	0.696
$\tau = 0.5$	0.342	0.342	0.363	0.367	0.601	0.602	0.599	0.622
$\tau = 0.75$	0.336	0.338	0.342	0.328	0.596	0.593	0.604	0.589
Diff(0.25-0.75)	0.183	0.186	0.183	0.207	0.340	0.347	0.349	0.365
Uniform	0.529	0.533	0.521	0.536	0.846	0.846	0.845	0.843
<i>B.2. Post-Lasso</i>								
$\tau = 0.25$	0.430	0.428	0.428	0.425	0.718	0.709	0.722	0.713
$\tau = 0.5$	0.375	0.359	0.370	0.359	0.619	0.621	0.626	0.635
$\tau = 0.75$	0.355	0.366	0.364	0.363	0.606	0.617	0.618	0.604
Diff(0.25-0.75)	0.210	0.207	0.211	0.202	0.365	0.363	0.364	0.368
Uniform	0.580	0.585	0.586	0.582	0.870	0.870	0.878	0.870

Last, Table 6 reports the size and power of our regression-adjusted estimator for the median QTE when we set $\hat{\pi}(s)$ as the true propensity score 1/2. Comparing it with Table 2, we see that using the true, instead of the estimated, propensity score, the multiplier bootstrap inference is conservative for randomization schemes “WEI”, “BCD”, and “SBR”. Specifically, the sizes are much smaller than the nominal rate (5%) while the powers are smaller than their counterparts in

Table 2. This is consistent with the findings in Bugni et al. (2018) and Zhang and Zheng (2020) that the naive inference methods under CARs are conservative. Similar results for $\tau = 0.25$, $\tau = 0.75$, the difference of the 0.25 and 0.75 QTEs, and QTEs uniformly over $\tau \in [0.25, 0.75]$ are reported in Section B in the Online Supplement. All the remarks here still apply.

Table 6: Pointwise Test with Naïve Estimator ($\tau = 0.5$, $\hat{\pi}(s) = 0.5$)

Methods	$N = 200$				$N = 400$			
	SRS	WEI	BCD	SBR	SRS	WEI	BCD	SBR
Panel A: DGP (i)								
<i>A.1: Size</i>								
NA	0.049	0.025	0.014	0.017	0.047	0.028	0.015	0.016
LP	0.046	0.008	0.000	0.000	0.050	0.008	0.000	0.000
LG	0.053	0.017	0.008	0.006	0.051	0.019	0.008	0.006
ML	0.052	0.015	0.009	0.004	0.050	0.015	0.006	0.006
NP	0.053	0.020	0.007	0.006	0.056	0.018	0.006	0.006
<i>A.2: Power</i>								
NA	0.287	0.257	0.249	0.245	0.495	0.495	0.511	0.496
LP	0.180	0.103	0.057	0.036	0.307	0.253	0.194	0.150
LG	0.260	0.212	0.185	0.151	0.476	0.446	0.433	0.414
ML	0.267	0.242	0.213	0.183	0.511	0.504	0.509	0.485
NP	0.293	0.254	0.219	0.191	0.502	0.505	0.507	0.486
Panel B: DGP (ii)								
<i>B.1: Size</i>								
NA	0.050	0.017	0.006	0.005	0.050	0.017	0.006	0.006
LP	0.050	0.005	0.000	0.000	0.048	0.006	0.000	0.000
LG	0.055	0.012	0.005	0.006	0.059	0.016	0.005	0.005
ML	0.050	0.021	0.013	0.016	0.046	0.025	0.022	0.020
NP	0.060	0.022	0.019	0.018	0.055	0.029	0.027	0.024
<i>B.2: Power</i>								
NA	0.290	0.245	0.220	0.221	0.491	0.496	0.494	0.501
LP	0.170	0.093	0.041	0.019	0.301	0.234	0.154	0.118
LG	0.260	0.181	0.137	0.125	0.476	0.446	0.413	0.399
ML	0.350	0.279	0.241	0.218	0.625	0.607	0.575	0.559
NP	0.365	0.300	0.263	0.244	0.645	0.621	0.599	0.577

7.4 Practical Recommendations

When X is finite-dimensional, we suggest using the ML estimator to construct the fitted auxiliary model when the logistic model includes additional technical terms and the regression coefficients are allowed to depend on (τ, a, s) . When X is high-dimensional, we suggest using the logistic post-Lasso estimation of the auxiliary model.

8 Empirical Application

Evaluating the impact of child health and nutrition on educational outcomes has been found to provide a key to understanding the links between health and economic development in less developed countries (Glewwe and Miguel, 2007; Dupas and Miguel, 2017). In particular, Chong, Cohen, Field, Nakasone, and Torero (2016) conducted a randomized experiment in Peru with a CAR to study the impact of an iron supplement encouragement program on schooling attainment. In their paper, the authors focused on the estimating the ATEs of that program. This section reports an application of our methods to the same dataset to examine the QTEs of the program on student academic achievement – a central outcome of interest in their study.

The sample consists of 215 students from a rural secondary school in Peru during the 2009 academic year. Within each of 5 secondary school grade levels, one-third of students were randomly assigned to each of three groups: two treatment arms and one control group. This is a block stratified randomization design with 5 strata, which satisfies Assumption 1 in Section 2.⁸ Students in the first treatment arm were shown a video in which a popular soccer player encouraged iron supplements; students in the second treatment arm were shown a video in which a doctor encouraged iron supplements; students in the control group were shown a video without mentioning iron supplements at all. In this section, we focus on the impact of iron supplements on the academic achievement measured by a standardized average of a student’s grades from five subjects (See Chong et al., 2016 for a detailed description). Throughout the analysis, students in the two treatment arms are grouped into the treatment group as in Chong et al. (2016).

One of the key findings in Chong et al. (2016) is that iron supplements have a significant positive impact on grades for anemic students, which is shown in Table 7 where we repeat their ATE estimates for easier comparison between ATEs and QTEs.⁹ The columns “Full 1x” and “Anemic 1x” refer to the ATE estimates obtained from the regressions on the full sample and the anemic subsample respectively, controlling for only one additional baseline variable (male) besides strata indicators. The columns “Full 4x” and “Anemic 4x” refer to the ATE estimates obtained from similar regressions while controlling for four additional baseline variables (male, monthly income, electricity in home, and mother’s years of schooling).

Tables 8 and 9 present the QTE estimates and their standard errors (in parentheses) estimated by different methods at quantile indexes 0.25, 0.5, and 0.75. The description of these estimators is

⁸It is trivial to see that statements (i), (ii), and (iii) in Assumption 1 are satisfied. Because $\sup_{s \in \mathcal{S}} |D_n(s)| \approx 0.03$ in both the full sample and the anemic subsample, it is plausible to claim that Assumption 1(iv) is also satisfied in our analysis.

⁹The ATE results on the anemic subsample are directly retrieved from Table 5 in Chong et al. (2016). The ATE results on the full sample are obtained by applying the same Stata code provided by them to the full sample.

similar to that in Section 7.¹⁰ Following the lead of Chong et al. (2016), throughout the analysis, we focus on two sets of additional baseline variables: male only (one auxiliary regressor) and male, monthly income, electricity in home, and mother’s years of schooling (four auxiliary regressors). Table 8 reports the results with only one auxiliary regressor and Table 9 reports the results with four auxiliary regressors.

Table 7: ATEs of Iron Supplement Encouragement on Grades in Chong et al. (2016)

	Full 1x	Full 4x	Anemic 1x	Anemic 4x
ATE	0.091 (0.140)	0.081 (0.140)	0.455 (0.218)	0.434 (0.215)

Notes: The table repeats the ATE estimates of the effect of the iron supplement encouragement program on grades reported in Chong et al. (2016). Standard errors are in parentheses.

The results in Tables 8-9 lead to two observations. First, consistent with the theoretical and simulation results, the standard errors for the regression-adjusted QTEs are mostly lower than those for the QTE estimate without adjustment. This observation holds generally, irrespective of the specification and estimation methods of the auxiliary regressions. For example in Table 8, in the anemic subsample, the standard errors for the “ML” QTE estimates are 14.1% and 3.5% less than those for the QTE estimate without adjustment at the median and the 75th percentile, respectively. For another example in Table 9, at the 25th percentile, the standard errors for the “Lasso” QTE estimates are 24.8% and 14.1% less than those for the QTE estimate without adjustment in the full sample and the anemic subsample, respectively.

Second, there seems to be heterogeneity in the impact of the iron supplement encouragement.

¹⁰Specifically,

- (i) NA: the estimator without any adjustments.
- (ii) LP: linear probability model.
- (iii) LG: logistic probability model with regressors $(1, X_i^\top)^\top$ and coefficient estimator $\hat{\theta}_{a,s}^{LG}(\tau)$. In line with the simulation, we add a ridge penalty $\log^{-1}(n(s))\|\theta\|_2^2$ to the objective function to stabilize the solution.
- (iv) ML: logistic probability model with regressor H and coefficient estimator $\hat{\theta}_{a,s}^{ML}(\tau)$. When there is only one auxiliary regressor, $H_i = (1, X_{1i})^\top$, and when there are four auxiliary regressors, $H_i = (1, X_{1i}, X_{2i}, X_{3i}, X_{4i}, X_{3i}X_{4i})^\top$, where $X_{1i}, X_{2i}, X_{3i}, X_{4i}$, represent these variables for student i : male, monthly income, electricity in home, and mother’s years of schooling, respectively.
- (v) NP: logistic probability model with regressor H_{h_n} and coefficient estimator $\hat{\theta}_{a,s}^{NP}(\tau)$. NP is only applied to the case when there are four auxiliary regressors with $H_{h_n} = (1, X_{1i}, X_{2i}, X_{3i}, X_{4i}, X_{2i}1\{X_{2i} > t_1\}X_{4i}1\{X_{4i} > t_2\})^\top$ where t_1 and t_2 are the sample medians of $\{X_{2i}\}_{i \in [n]}$ and $\{X_{4i}\}_{i \in [n]}$, respectively.
- (vi) Lasso: logistic probability model with regressor H_{p_n} and post-Lasso coefficient estimator $\hat{\theta}_{a,s}^{post}(\tau)$. Lasso is only applied to the case when there are four auxiliary regressors with $H_{p_n}(X_i) = (1, X_i^\top)^\top$ and $\hat{S}_{a,s}^+(q) = \{2\}$. The post-Lasso estimator $\hat{\theta}_{a,s}^{post}(\tau)$ is defined in (5.15). The choice of tuning parameter and the estimation procedure are detailed in Section 6.3.

Table 8: QTEs of Iron Supplement Encouragement on Grades (one auxiliary regressor)

	NA	LP	LG	ML
<i>Panel A. full sample</i>				
25%	0.317 (0.121)	0.317 (0.142)	0.238 (0.121)	0.317 (0.121)
50%	0.079 (0.182)	0.000 (0.162)	0.000 (0.162)	0.000 (0.162)
75%	0.159 (0.243)	-0.079 (0.223)	0.000 (0.223)	-0.079 (0.223)
<i>Panel B. anemic subsample</i>				
25%	0.381 (0.170)	0.381 (0.146)	0.381 (0.158)	0.381 (0.146)
50%	0.761 (0.243)	0.761 (0.243)	0.666 (0.243)	0.761 (0.243)
75%	0.666 (0.340)	0.666 (0.316)	0.666 (0.340)	0.666 (0.328)

Notes: The table presents the QTE estimates of the effect of the iron supplement encouragement program across the distribution of grade when only one auxiliary regressor (male) is used in the regression adjustment models. Standard errors are in parentheses.

Table 9: QTEs of Iron Supplement Encouragement on Grades (four auxiliary regressors)

	NA	LP	LG	ML	NP	Lasso
<i>Panel A. full sample</i>						
25%	0.317 (0.121)	0.317 (0.142)	0.317 (0.121)	0.317 (0.101)	0.317 (0.101)	0.317 (0.091)
50%	0.079 (0.182)	-0.079 (0.182)	0.000 (0.162)	-0.079 (0.142)	0.000 (0.162)	0.000 (0.142)
75%	0.159 (0.243)	-0.159 (0.233)	0.079 (0.223)	-0.079 (0.223)	-0.079 (0.243)	-0.079 (0.223)
<i>Panel B. anemic subsample</i>						
25%	0.381 (0.170)	0.476 (0.381)	0.381 (0.146)	0.381 (0.146)	0.381 (0.146)	0.381 (0.146)
50%	0.761 (0.243)	0.666 (0.267)	0.666 (0.243)	0.761 (0.219)	0.666 (0.194)	0.666 (0.219)
75%	0.666 (0.340)	0.571 (0.303)	0.571 (0.340)	0.666 (0.316)	0.666 (0.291)	0.857 (0.291)

Notes: The table presents the QTE estimates of the effect of the iron supplement encouragement program on academic grades at quantiles 25%, 50%, and 75% when four auxiliary regressor (male, monthly income, electricity in home, and mother's years of schooling) are used in the regression adjustment models. Standard errors are in parentheses.



Figure 1: Quantile Treatment Effects on the Distribution of Grades (Anemic subsample)

Notes: The figure plots the QTE estimates of the effect of the iron supplement encouragement program on the distribution of grades in the anemic subsample when there are four auxiliary regressors: male, monthly income, electricity at home, mother education. The shadow areas represent the bounds of 95% confidence intervals.

Specifically, in the full sample, only the treatment effects at the 25th percentile are significantly positive; in the anemic subsample, the treatment effects increase by about 50% from the 25th percentile to the median and seem to stay constant afterwards.

How does the variation in the impact of the iron supplement encouragement appear across the distribution of grades? The QTEs on the distribution of grades for the anemic subsample are plotted in Fig 1, where the shaded areas represent the 95% confidence region. The figure shows that the QTEs seem insignificantly different from zero below about the 20% percentile. At higher levels to around 80% percentile the treatment group grades exceed the control group grades, yielding significant positive QTEs. Beyond the 80% percentile, the QTEs again become insignificantly different from zero. These findings point to notable heterogeneity in the QTEs of the impact of the iron supplement encouragement program on the distribution of grades.

9 Conclusion

This paper proposes the use of auxiliary regression to incorporate additional covariates into estimation and inference relating to unconditional QTEs under CARs. The auxiliary regression

model may be estimated parametrically, nonparametrically, or via regularization if there are high-dimensional covariates. Both estimation and bootstrap inference methods are robust to potential misspecification of the auxiliary model and do not suffer from the conservatism due to the CAR. It is shown that efficiency can be improved when including extra covariates and when the auxiliary regression is correctly specified the regression-adjusted estimator achieves the minimum asymptotic variance. In both the simulations and the empirical application, the proposed regression-adjusted QTE estimator performs well. These results and the robustness of the methods to auxiliary model misspecification reflect the aphorism widespread in scientific modeling that all models may be wrong, but some are useful.¹¹

¹¹The aphorism “all models are wrong but some are useful” is often attributed to the statistician George Box (1976). But the notion has many antecedents, including a particularly apposite remark made in 1947 by John von Neumann (2019) in an essay on the empirical origins of mathematical ideas to the effect that “truth ... is much too complicated to allow anything but approximations”.

References

- Abadie, A., M. M. Chingos, and M. R. West (2018). Endogenous stratification in randomized experiments. *The Review of Economics and Statistics* 100(4), 567–580.
- Athey, S. and G. Imbens (2017). Chapter 3 - the econometrics of randomized experimentsa. In A. V. Banerjee and E. Duflo (Eds.), *Handbook of Field Experiments*, Volume 1 of *Handbook of Economic Field Experiments*, pp. 73 – 140. North-Holland.
- Bai, Y. (2019). Optimality of matched-pair designs in randomized controlled trials. *Available at SSRN 3483834*.
- Bai, Y., A. Shaikh, and J. P. Romano (2019). Inference in experiments with matched pairs. *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2019-63).
- Banerjee, A., E. Duflo, R. Glennerster, and C. Kinnan (2015). The miracle of microfinance? evidence from a randomized evaluation. *American Economic Journal: Applied Economics* 7(1), 22–53.
- Belloni, A., V. Chernozhukov, I. Fernández-Val, and C. Hansen (2017). Program evaluation with high-dimensional data. *Econometrica* 85(1), 233–298.
- Bitler, M. P., J. B. Gelbach, and H. W. Hoynes (2006). What mean impacts miss: Distributional effects of welfare reform experiments. *American Economic Review* 96(4), 988–1012.
- Bloniarz, A., H. Liu, C.-H. Zhang, J. S. Sekhon, and B. Yu (2016). Lasso adjustments of treatment effect estimates in randomized experiments. *Proceedings of the National Academy of Sciences* 113(27), 7383–7390.
- Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association* 71(356), 791–799.
- Bradic, J., S. Wager, and Y. Zhu (2019). Sparsity double robust inference of average treatment effects. *arXiv preprint arXiv:1905.00744*.
- Bugni, F. A., I. A. Canay, and A. M. Shaikh (2018). Inference under covariate-adaptive randomization. *Journal of the American Statistical Association* 113(524), 1741–1768.
- Bugni, F. A., I. A. Canay, and A. M. Shaikh (2019). Inference under covariate-adaptive randomization with multiple treatments. *Quantitative Economics* 10(4), 1747–1785.

- Bugni, F. A. and M. Gao (2021). Inference under covariate-adaptive randomization with imperfect compliance. *arXiv preprint arXiv:2102.03937*.
- Campos, F., M. Frese, M. Goldstein, L. Iacovone, H. C. Johnson, D. McKenzie, and M. Mensmann (2017). Teaching personal initiative beats traditional training in boosting small business in west africa. *Science* 357(6357), 1287–1290.
- Chen, X. (2007). Large sample sieve estimation of semi-nonparametric models. *Handbook of econometrics* 6, 5549–5632.
- Chernozhukov, V., I. Fernández-Val, and B. Melly (2013). Inference on counterfactual distributions. *Econometrica* 81(6), 2205–2268.
- Chong, A., I. Cohen, E. Field, E. Nakasone, and M. Torero (2016). Iron deficiency and schooling attainment in Peru. *American Economic Journal: Applied Economics* 8(4), 222–55.
- Crépon, B., F. Devoto, E. Duflo, and W. Parienté (2015). Estimating the impact of microcredit on those who take it up: Evidence from a randomized experiment in Morocco. *American Economic Journal: Applied Economics* 7(1), 123–50.
- Duflo, E., M. Greenstone, R. Pande, and N. Ryan (2013). Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India. *The Quarterly Journal of Economics* 128(4), 1499–1545.
- Dupas, P. and E. Miguel (2017). Impacts and Determinants of Health Levels in Low-Income Countries. In *Handbook of Economic Field Experiments*, Volume 2, pp. 3–93. Elsevier.
- Firpo, S. (2007). Efficient semiparametric estimation of quantile treatment effects. *Econometrica* 75(1), 259–276.
- Fogarty, C. B. (2018). Regression-assisted inference for the average treatment effect in paired experiments. *Biometrika* 105(4), 994–1000.
- Freedman, D. A. (2008a). On regression adjustments in experiments with several treatments. *Annals of Applied Statistics* 2(1), 176–196.
- Freedman, D. A. (2008b). On regression adjustments to experimental data. *Advances in Applied Mathematics* 40(2), 180–193.
- Glewwe, P. and E. A. Miguel (2007). Chapter 56 The Impact of Child Health and Nutrition on Education in Less Developed Countries. In *Handbook of Development Economics*, Volume 4, pp. 3561–3606. Elsevier.

- Hahn, J., K. Hirano, and D. Karlan (2011). Adaptive experimental design using the propensity score. *Journal of Business & Economic Statistics* 29(1), 96–108.
- Hirano, K., G. W. Imbens, and G. Ridder (2003). Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica* 71(4), 1161–1189.
- Hu, Y. (2016). Generalized Efron’s biased coin design and its theoretical properties. *Journal of Applied Probability* 53(2), 327–340.
- Hu, Y. and F. Hu (2012). Asymptotic properties of covariate-adaptive randomization. *The Annals of Statistics* 40(3), 1794–1815.
- Jiang, L., X. Liu, P. C. Phillips, and Y. Zhang (forthcoming 2021). Bootstrap inference for quantile treatment effects in randomized experiments with matched pairs. *Review of Economics and Statistics*.
- Kallus, N., X. Mao, and M. Uehara (2019). Localized debiased machine learning: Efficient inference on quantile treatment effects and beyond. *arXiv preprint arXiv:1912.12945*.
- Lei, L. and P. Ding (2021+). Regression adjustment in completely randomized experiments with a diverging number of covariates. *Biometrika*. forthcoming.
- Li, X. and P. Ding (2020). Rerandomization and regression adjustment. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82(1), 241–268.
- Lin, W. (2013). Agnostic notes on regression adjustments to experimental data: Reexamining freedman’s critique. *Annals of Applied Statistics* 7(1), 295–318.
- Liu, H., F. Tu, and W. Ma (2020). A general theory of regression adjustment for covariate-adaptive randomization: Ols, lasso, and beyond. *arXiv preprint arXiv:2011.09734*.
- Liu, H. and Y. Yang (2020). Regression-adjusted average treatment effect estimates in stratified randomized experiments. *Biometrika* 107(4), 935–948.
- Lu, J. (2016). Covariate adjustment in randomization-based causal inference for 2k factorial designs. *Statistics & Probability Letters* 119, 11–20.
- Ma, W., F. Hu, and L. Zhang (2015). Testing hypotheses of covariate-adaptive randomized clinical trials. *Journal of the American Statistical Association* 110(510), 669–680.
- Ma, W., Y. Qin, Y. Li, and F. Hu (2018). Statistical inference of covariate-adjusted randomized experiments. *arXiv preprint arXiv:1807.09678*.

- Ma, W., F. Tu, and H. Liu (2020). Regression analysis for covariate-adaptive randomization: A robust and efficient inference perspective. *arXiv preprint arXiv:2009.02287*.
- Muralidharan, K. and V. Sundararaman (2011). Teacher performance pay: Experimental evidence from India. *Journal of Political Economy* 119(1), 39–77.
- Negi, A. and J. M. Wooldridge (2020). Revisiting regression adjustment in experiments with heterogeneous treatment effects. *Econometric Reviews*, 1–31.
- Olivares, M. (2021). Robust permutation test for equality of distributions under covariate-adaptive randomization.
- Shao, J. and X. Yu (2013). Validity of tests under covariate-adaptive biased coin randomization and generalized linear models. *Biometrics* 69(4), 960–969.
- Shao, J., X. Yu, and B. Zhong (2010). A theory for testing hypotheses under covariate-adaptive randomization. *Biometrika* 97(2), 347–360.
- Tabord-Meehan, M. (2018). Stratification trees for adaptive randomization in randomized controlled trials. *arXiv preprint arXiv:1806.05127*.
- Tan, Z. (2020). Model-assisted inference for treatment effects using regularized calibrated estimation with high-dimensional data. *Annals of Statistics* 48(2), 811–837.
- van der Vaart, A. and J. A. Wellner (1996). *Weak Convergence and Empirical Processes*. Springer, New York.
- von Neumann, J. (2019). *The mathematician*. Chapman and Hall/CRC.
- Wei, L. (1978). An application of an urn model to the design of sequential controlled clinical trials. *Journal of the American Statistical Association* 73(363), 559–563.
- Ye, T. (2018). Testing hypotheses under covariate-adaptive randomisation and additive models. *Statistical Theory and Related Fields* 2(1), 96–101.
- Ye, T. and J. Shao (2020). Robust tests for treatment effect in survival analysis under covariate-adaptive randomization. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82(5), 1301–1323.
- Ye, T., Y. Yi, and J. Shao (2020). Inference on average treatment effect under minimization and other covariate-adaptive randomization methods. *arXiv preprint arXiv:2007.09576*.

- Zhang, M., A. A. Tsiatis, and M. Davidian (2008). Improving efficiency of inferences in randomized clinical trials using auxiliary covariates. *Biometrics* 64(3), 707–715.
- Zhang, Y. and X. Zheng (2020). Quantile treatment effects and bootstrap inference under covariate-adaptive randomization. *Quantitative Economics* 11(3), 957–982.
- Zhao, A. and P. Ding (2020). Covariate-adjusted fisher randomization tests for the average treatment effect. *arXiv preprint arXiv:2010.14555*.