Journal of Econometrics 249 (2025) 105857

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Efficient quantile covariate adjusted response adaptive experiments

Zhonghua Li [a], Lan Luo [b], Jingshen Wang [c,*], Long Feng [a,*]

[a] School of Statistics and Data Science, LPMC, LEBPS and KLMDASR, Nankai University, China
[b] Department of Biostatistics and Epidemiology, Rutgers University, USA
[c] Division of Biostatistics, University of California, Berkeley, USA

ARTICLE INFO

JEL classification:
C12
C13

Keywords:
Semiparametric statistics
Adaptive design
Heterogeneity

ABSTRACT

In program evaluation studies, understanding the heterogeneous distributional impacts of a program beyond the average effect is crucial. Quantile treatment effect (QTE) provides a natural measure to capture such heterogeneity. While much of the existing work for estimating QTE has focused on analyzing observational data based on untestable causal assumptions, little work has gone into designing randomized experiments specifically for estimating QTE. In this manuscript, we propose two covariate-adjusted response adaptive design strategies–fully adaptive designs and multi-stage designs–to efficiently estimate the QTE. We demonstrate that the QTE estimator obtained from our designs attains the optimal variance lower bound from a semiparametric theory perspective, which does not impose any parametric assumptions on underlying data distributions. Moreover, we show that using continuous covariates in multi-stage designs can improve the precision of the estimated QTE compared to the classical fully adaptive setting. We illustrate the finite-sample performance of our designs through Monte Carlo experiments and one synthetic case study on charitable giving. Our proposed designs offer a new approach to conducting randomized experiments to estimate QTE, which can have important implications for policy and program evaluation.

1. Introduction

1.1. Motivation

Randomized experiments have become increasingly important in evaluating the effects of drugs, social policies, or economic theories. Results obtained from well-designed randomized experiments are considered the most reliable evidence and have become the gold standard for assessing the effectiveness of interventions. As randomized experiments can be costly and time-consuming, it is crucial to design them carefully to achieve specific objectives, such as reducing the variance of the estimated average treatment effect (ATE) or identifying which subgroups benefit from medication.

While existing literature often measures treatment effectiveness through the ATE, it is also important to capture the distributional impact of treatment beyond the ATE. One approach to achieve this is by estimating the quantile treatment effect (QTE), which measures the quantile differences between the distributions of treated and control outcomes (Firpo, 2007). In particular, when outcome variables have skewed distributions, such as income and donation amount, the QTE provides a more informative measure of the treatment than the classical ATE. For example, in studying the effectiveness of a job training program, it may be more relevant for public policy to examine the treatment effect on the lower tail of the earnings distribution rather than the center. Another example is studying the effectiveness of a matching grant on charitable giving, where the mean treatment effect tends to

* Corresponding authors.
E-mail addresses: jingshenwang@berkeley.edu (J. Wang), flnankai@gmail.com (L. Feng).

https://doi.org/10.1016/j.jeconom.2024.105857
Received 5 April 2023; Received in revised form 29 January 2024; Accepted 29 July 2024
Available online 20 September 2024

be marginal, but the difference between the upper quantiles is significant (Karlan and List, 2007). In such cases, a QTE-targeted randomized experimental design may be more desirable.

In addition to careful selection of an experimental target parameter, efficient randomized experimental designs often make efficient use of the evidence collected during the experiment. This motivates us to deviate from classical randomized experiments, in which the treatment is assigned with a constant probability throughout the experiment. Instead, we allow randomization probabilities to be adapted during the experiment based on the accrued participant responses. These designs are referred to as response adaptive randomization (RAR) designs and have been shown to have several advantages over simple randomization and balanced allocation in sequential biomedical trials or experiments.

In practice, experimenters often have access to covariate information about participants, such as their preferences or past activities, prior to participant arrival. This information can be used to allow the randomization probability to vary depending on the observed covariate values. Consider an example in A/B tests, where an internet company is interested in comparing two promotion strategies, A and B, that are displayed to users on the company website. The comparison could be based on the performance (such as click or not click) of a strategy that is taken expectation over a randomly arrived user. As an extreme case, a user group with some specific covariate information never chooses Strategy A, and hence, their measured outcomes have no variation. This suggests that, ideally, the allocation probability of Strategy A in this user group should be set to zero. In such cases, the availability of covariate information and the variability of treatment effect variances across user subgroups provide a natural opportunity to improve RAR design efficiency. This type of design, known as covariate-adjusted response adaptive (CARA) randomization (see Hu et al., 2015, for example) design, takes advantage of the variability in treatment effects across participant subgroups to improve design efficiency.

Existing literature on CARA designs typically incorporates covariate information by either assuming a parametric model for the conditional mean of the potential outcome given the covariates (see Zhang et al., 2007; Zhu, 2015; Hu et al., 2015, for example), or by stratifying participants into pre-fixed strata based on their covariate information, whether it is continuous or discrete (see Rosenberger and Sverdlov, 2008; Bandyopadhyay and Biswas, 2001; Zhao et al., 2022; Atkinson et al., 2011; Baldi Antognini and Zagoraiou, 2011; Antognini and Zagoraiou, 2012, for example). However, there are potential limitations with both approaches. For instance, misspecification of the parametric model can undermine the statistical validity of the proposed CARA design strategies, while the finite number of values of the covariate can lead to under-utilization of covariate information (Aletti et al., 2018). Furthermore, when the covariate takes a continuum of values, it is not possible to accumulate a sufficient number of observations under each covariate value with a pre-fixed sample size, which can also lead to a loss of efficiency in existing CARA designs. As such, it is necessary to explore nonparametric methods for estimating the conditional mean of the outcome variable and to develop new CARA algorithms that can handle continuous covariate information efficiently.

In this manuscript, in response to the concerns raised regarding existing CARA designs, we propose two CARA design frameworks addressing the following research questions: First, when we do not assume a parametric model for the conditional outcome given the covariates, how much improvement in efficiency can be achieved for estimating the QTE by using covariate information to assist the experimental design, compared to using RAR without the covariate information? Second, when covariates take continuous values, can we design a CARA strategy that maximizes the use of available covariate information and produces a causal effect estimator that achieves the best achievable variance lower bound for estimating the QTE asymptotically?

1.2. Our contribution

In this manuscript, we develop two new CARA experimental designs and theoretical insights toward efficiently learning quantile treatment effects. We elaborate our contribution as follows:

We introduce a design-based optimal variance lower bound for estimating the QTE from a semiparametric theory point of view, meaning that this variance lower bound does not rely on any distributional assumptions on the collected data. This design-based optimal variance lower bound is attainable if the treatment is assigned during the experiment accounting for both the fluctuations of treatment effect variances occurring across participants with different covariate information and the density of the potential outcomes. We propose two CARA strategies–the multi-stage design and the fully adaptive design–both achieving the optimal variance lower bound for estimating the QTE. Both of these designs allocate treatment assignments for incoming participants based on all available information from previous participants, including treatment assignments, responses, and covariates, while also incorporating the incoming participant's covariate information.

We demonstrate the trade-offs between the two designs in terms of practical implementation, theoretical justification, and simulation studies. From a theoretical justification perspective, when covariates are discrete, both designs achieve the same optimal variance lower bound as the sample size tends to infinity. Nevertheless, in finite samples, our simulation results in Section 5 show that fully adaptive designs tend to outperform multi-stage designs potentially due to their frequently updated randomization probabilities. In contrast, when covariates are continuous, multi-stage designs provide more efficient QTE estimators than fully adaptive designs, as classical fully adaptive designs need to stratify continuous covariates. Our theoretical investigation (Theorem 3) and simulation results (Section 5) both support this conclusion. From a practical standpoint, multi-stage designs may be preferred in clinical or social experiments where experimenters have limited opportunities to revise experimental effort allocation during the experiment, as demonstrated in studies such as Karlan and Zinman (2008) and Gertler et al. (2012). Fully adaptive designs are more suitable for digital experiments, such as online A/B testing or digital clinical trials, where sequentially allocating experimental efforts in a large number of stages is practical and less costly, as shown in studies such as Kharitonov et al. (2015) and Robertson et al. (2020). To demonstrate the practical performance of the proposed designs, we apply them in a synthetic case study in which the goal is to estimate the quantile treatment effect of matching grant on charitable giving.

Furthermore, we establish the asymptotic normality of the estimated QTE using the data collected from the proposed CARA designs. To ensure accurate inference on the QTE, we prove that the proposed variance estimator is a consistent estimate of the optimal variance lower bound. Importantly, neither the QTE estimator nor the variance estimator relies on a correctly specified model for the conditional distribution of the outcome given the covariates. To properly characterize the statistical properties of our proposed design, it is necessary to employ techniques that account for the reliance on adaptively estimated quantile treatment effects for both the estimated sequential randomization probabilities and the density function of the potential outcomes. Therefore, we need to develop novel techniques that go beyond the existing literature. Finally, we use semiparametric theory to demonstrate the efficiency gain of multi-stage designs that take continuous covariates into account when compared to fully adaptive designs (Theorem 3). This comparison helps to provide a more comprehensive understanding of the strengths and weaknesses of each design strategy.

1.3. Related literature

In economic research, quantile treatment effects are often explored to examine treatment heterogeneity. For example, under the unconfoundedness assumption, which states that there are no unmeasured confounding variables that affect both the treatment and the (potential) outcomes, Firpo (2007) develops a semiparametric efficient QTE estimator, and Cattaneo (2010) extends this estimator to multiple treatments. More generally, Chen et al. (2008) study the semiparametric efficiency bounds of parameters defined through general nonlinear and over-identified moment conditions. Chen et al. (2003) the consistency and asymptotic normality of a class of semiparametric optimization estimators, where the objective function does not adhere to standard smoothness conditions. When the unconfoundedness assumption does not hold in observational studies, Chen et al. (2019), Chernozhukov and Hansen (2005, 2006) and Abadie et al. (1998) illustrate how instrumental variables can be employed to identify QTE. In contrast, our proposed QTE estimators have a similar structure to Firpo (2007), as the treatment assignment mechanism in randomized experiments fulfills the unconfoundedness assumption during the data collection stage.

In contrast to our CARA design, another widely-used approach is covariate adaptive randomization (CAR). CAR involves randomizing treatment allocation based on participant covariate information and is commonly used in clinical trials to balance treatment assignments across important prognostic factors (Rosenberger et al., 2001; Zhang et al., 2007). Earlier variations of CAR include minimization (instead of randomization) (Zelen, 1974), and stratified block adaptive randomization (Pocock and Simon, 1975). Rerandomization has been studied to improve covariate balance by Morgan and Rubin (2012, 2015). Antognini and Zagoraiou (2015) provides asymptotic results for adaptive designs for treatment comparison, both with and without covariates. While the majority of the current CAR literature focuses on estimating ATE, recent works by Jiang et al. (2023) and Zhang and Zheng (2020) have explored the estimation of QTE under CAR. When the covariate information is incorporated in RAR, we have the CARA design (Hu et al., 2015). In such cases, the allocation probability of a participant is determined by the accumulated outcome information, the treatment assignment status and the covariates. Aletti et al. (2018) proposes nonparametric CARA design based on a functional urn model. Anderson and Mckenzie (2021) conducts a randomized experiment in Nigeria to test the relative effectiveness in improving business practices. Zhao et al. (2022) first classifies covariates into predictive and prognostic covariates, and proposes a new family of CARA designs. Lastly, Zhu and Zhu (2023) proposes a family of semiparametric approaches that use the target maximum likelihood estimation to analyze the correlated data from CARA designs. Their proposed method allows to incorporation of the effect of a large number of covariates on the responses without model misspecification.

2. Design-based optimal variance lower bound

In this section, we introduce a design-based optimal variance lower bound for estimating the quantile treatment effect. This optimal variance lower bound allows us to see how covariate information may be used to improve experimental design efficiency in CARA. Here, we measure the efficiency of the design by evaluating the asymptotic variance of the constructed causal effect estimator. We are also in a scenario where the covariate information can take continuous value in a bounded set.

We start by introducing the notations to be used throughout this manuscript. Let $A \in \{0, 1\}$ denote the treatment assignment status. Following the Neyman–Rubin causal model, let $Y(a)$ be the potential outcome if the subject receives the treatment $a \in \{0, 1\}$, and the observed outcome $Y = AY(1) + (1 - A)Y(0)$. Let τ be a real number in $(0, 1)$. For a given quantile level τ, the QTE refers to the difference between the quantiles of the distributions of the potential outcomes in the treated and controlled groups:

$$\alpha_\tau = Q_1(\tau) - Q_0(\tau),$$

where

$$Q_a(\tau) = \inf\left\{y \in \mathbb{R} : F_{Y(a)}(y) \geq \tau\right\},$$

is the quantile function of the potential outcome in the treatment arm a evaluated at a percentile level τ. Next, let X denote the pre-treatment covariates defined on the sample space $\mathcal{X} \subset \mathbb{R}^d$. Here, the covariates X can contain both continuous or discrete variables. As in clinical studies researchers may intentionally discretize continuous covariates to enhance interpretability, we also denote Z as the discretized covariate X.

For all $x \in \mathcal{X}$, suppose the treatment is randomly assigned based on the propensity score $e(x) = \mathbb{P}(A = 1|X = x)$. When researchers use discretized covariate in the experiment, we denote such propensity scores as $e(z) = \mathbb{P}(A = 1|Z = z)$. When the treatment assignment is independent of the potential outcomes conditional on the covariate, that is,

$$A \perp\!\!\!\perp (Y(1), Y(0))|X = x, \quad x \in \mathcal{X},$$

Firpo (2007) provides the semiparametric efficiency bound $\mathbb{V}_{\mathtt{SPEB},\tau}$ for estimating the QTE:

$$\mathbb{V}_{\mathtt{SPEB},\tau} = \mathbb{E}\left[\frac{\sigma_1^2(X)}{f_1^2(\tau)e(X)} + \frac{\sigma_0^2(X)}{f_0^2(\tau)\big(1-e(X)\big)} + \left[\frac{\mathbb{E}\big[\psi_\tau\big(Y-Q_1(\tau)\big)|X,A=1\big]}{f_1(\tau)} - \frac{\mathbb{E}\big[\psi_\tau\big(Y-Q_0(\tau)\big)|X,A=0\big]}{f_0(\tau)}\right]^2\right]. \tag{1}$$

Here, for $a\in\{0,1\}$, $f_a(\tau)\triangleq f_a\big(Q_a(\tau)\big)$ is the marginal density of the potential outcome $Y(a)$ evaluated at $Q_a(\tau)$, $\sigma_a^2(x)=\mathrm{Var}\big[\psi_\tau\big(Y-Q_a(\tau)\big)\big|X=x,A=a\big]$ is the conditional variance of the quantile rank score of the potential outcomes, where $\psi_\tau(u)=\tau-\mathbf{1}(u\le 0)$.

The semiparametric efficiency bound $\mathbb{V}_{\mathtt{SPEB},\tau}$ provided above is dependent on the treatment assignment. If we know $f_a(\tau)$ and $\sigma_a^2(x)$, we can hypothetically choose how the treatments are assigned to achieve the lowest possible variance in the estimated QTE. More concretely, this can be achieved by assigning treatment based on an optimal propensity score $e^*(x)$:

$$e^*(x)=\frac{\sigma_1(x)/f_1(\tau)}{\sigma_1(x)/f_1(\tau)+\sigma_0(x)/f_0(\tau)}.$$

If we assign the treatment using $e^*(x)$, the optimal variance lower bound of estimating QTE is:

$$\mathbb{V}_{\mathtt{Optimal},\tau} = \mathbb{E}\left[\left(\frac{\sigma_1(X)}{f_1(\tau)}+\frac{\sigma_0(X)}{f_0(\tau)}\right)^2 + \left[\frac{\mathbb{E}\big[\psi_\tau\big(Y-Q_1(\tau)\big)|X,A=1\big]}{f_1(\tau)} - \frac{\mathbb{E}\big[\psi_\tau\big(Y-Q_0(\tau)\big)|X,A=0\big]}{f_0(\tau)}\right]^2\right]. \tag{2}$$

The optimal variance lower bound result suggests that if we have perfect knowledge about the marginal density of potential outcomes and how the conditional variances of the quantile rank score of potential outcomes vary across different covariate values x, we can construct a QTE estimator with the lowest possible variance $V_{\mathtt{Optimal},\tau}$. In practice, when we neither have prior information on $\sigma_a^2(x)$ and $f_a(\tau)$ nor are willing to assume them to follow parametric models which can be arbitrarily misspecified, CARA offers an ideal environment to sequentially discover the conditional variability of the outcome variable given the covariates as the experiment proceeds.

In the following section, we propose two CARA design strategies that ultimately provide QTE estimators that attain the optimal variance lower bound $V_{\mathtt{Optimal}}$. It is worth noting that the variance lower bound is influenced by the choice of covariates used in the randomization process. Specifically, the use of continuous covariates versus discretized covariates can result in different efficiency bounds (refer to Theorems 1 and 2 for more detailed results).

3. Data collection mechanism and design strategies

In this section, we provide two CARA experimental design strategies that provide efficient QTE estimators attaining the derived variance lower bound $V_{\mathtt{Optimal},\tau}$. To better synthesize the two proposed design strategies, we start by introducing a general data collection mechanism of two-arm adaptive experiments.

3.1. Data collection mechanism

Suppose experiment participants are sequentially enrolled in different stages $t=1,\dots,T$, and respond to treatments with no delays. We denote the number of participants contained in stage t as n_t and the total number of participants $N_T=\sum_{t=1}^T n_t$, and let $N_t=\sum_{s=1}^t n_s$ denote the cumulative sample size up to stage t. For the ith enrolled participant in the tth stage, we introduce an ancillary stage variable $S_i\in\{1,\dots,T\}$ that records the stage of the ith participant is enrolled in the experiment. Each stage has enrolled a proportion of r_t patients with

$$\mathbb{P}(S_i=t)=r_t,\quad \sum_{t=1}^T r_t=1.$$

The observed outcome in such multi-stage randomized experiments can thus be written as

$$Y_i=\sum_{t=1}^T \mathbf{1}_{(S_i=t)}\Big(A_iY_i(1)+(1-A_i)Y_i(0)\Big).$$

In line with existing literature on response adaptive designs (Hu and Rosenberger, 2006), we are also in a scenario where the joint distributions of the potential outcomes and covariates collected in different stages do not shift over time. As adaptive experiments allow practitioners to revise the treatment assignment probability sequentially in multiple stages, we denote such sequential propensity scores for the subject with either covariate x or discretized covariate z enrolled in stage t as

$$e_t(x)=\mathbb{P}(A_i=1|X_i=x,S_i=t),\quad e_t(z)=\mathbb{P}(A_i=1|Z_i=z,S_i=t).$$

In the following section, we propose two design strategies that can accommodate different types of covariates, either original ones (X) or discretized ones (Z). In fully adaptive experiments, where T may be large and n_t is small, treatment assignment needs

to be updated frequently, making discretization of covariates a common practice in these designs. Note in the fully adaptive design, we do not restrict n_t to be one. From a practical perspective, this means that it is not completely necessary to adjust the treatment assignment mechanism each time a new subject is enrolled. Rather, we can make revisions to the treatment assignment mechanism after enrolling a few subjects. From a theoretical point of view, if n_t is set as a fixed constant, whether it is one or five, as long as it remains a fixed positive integer, it will not affect the theoretical properties of our design (as sample size N is of the same order as the stage number T). In multi-stage designs, where T may be small and n_t is large, each stage has enough data to estimate unknown parameters, making it more feasible to directly incorporate continuous covariates.

3.2. Design strategy I: Fully adaptive design

In this section, we provide our design strategy in the classical fully adaptive setting (that is, small n_t and $T \to \infty$) with discretized covariates Z. Fully adaptive designs can be more preferable in digital experiments, such as online A/B testing or digital clinical trials. For example, in online A/B testing, an internet company aims to compare two promotional strategies, A and B, showcased on its website. This is done by directing users into two groups, with each group exposed to a different promotional strategy. The effectiveness of each strategy is measured based on the averaged user responses (such as clicking or not clicking). Typically, an A/B test sets a total sample size of N users, where each user is presented with Strategy A (with a probability e) or Strategy B (with a probability $1-e$). In a digital environment, due to the ability to quickly process user feedback and adjust parameters in real-time, it is feasible to rapidly adjust the randomization probabilities. Therefore, we believe that a fully adaptive design, which allows for continuous modification of these probabilities based on ongoing results, may be more suitable in such contexts. This is also demonstrated in studies such as Kharitonov et al. (2015) and Robertson et al. (2020). In these digital experiments, the treatment assignment can be updated in real-time based on the accumulated data, allowing for efficient use of resources and faster evaluation of treatment effects. Our design strategy proceeds as follows:

Stage 1. *Assign each enrolled subject $i = 1, \dots, n_1$ to the treatment arm with a pre-specified propensity score e_1. At the end of stage 1, we collect observations $\{(Y_i, A_i, Z_i)\}_{i=1}^{n_1}$ for these enrolled subjects where Z_i denotes the discretized covariate for subject i.*

As we do not have prior information about the enrolled participants, stage 1 serves as an exploration stage in which we hope obtain initial estimates of the unknown quantities.

Stage t, **for** $t = 2, \dots, T-1$. *Given the estimates of the unknown quantities obtained in stage $t-1$ shall be shown in Eqs. (3)–(5), we calculate the estimated optimal propensity score for stage t with*

$$\widehat{e}_t(z) = \frac{\widehat{\sigma}_{1,t-1}(z)/\widehat{f}_{1,t-1}(\tau)}{\widehat{\sigma}_{1,t-1}(z)/\widehat{f}_{1,t-1}(\tau) + \widehat{\sigma}_{0,t-1}(z)/\widehat{f}_{0,t-1}(\tau)}.$$

We then allocate subjects enrolled in stage t to the treatment arm according to the estimated optimal propensity score $\widehat{e}_t(z)$ for stage t, and collect data $\{(Y_i, A_i, Z_i)\}_{i=N_{t-1}+1}^{N_t}$ for these individuals enrolled in stage t.

Using the data $\{(Y_i, A_i, Z_i)\}_{i=1}^{N_t}$ collected up to stage t and the estimated optimal propensity scores $\{\widehat{e}_s\}_{s=1}^{t}$, we estimate the quantile treatment effect $\widehat{Q}_{a,t}(\tau)$ by solving the following optimization problem:

$$\widehat{Q}_{a,t}(\tau) = \arg\min_{u\in\mathbb{R}} \frac{1}{N_t}\sum_{s=1}^{t}\sum_{i=N_{s-1}+1}^{N_s} \frac{\mathbf{1}(A_i = a)}{A_i\widehat{e}_s(Z_i) + (1-A_i)\big(1-\widehat{e}_s(Z_i)\big)}\rho_\tau(Y_i - u), \tag{3}$$

where $\rho_\tau(a) = a\cdot\big(\tau - \mathbf{1}(a \le 0)\big)$ is the check-loss function. Then, we update the estimates of unknown quantities to $\{\widehat{\sigma}^2_{1,t}(z), \widehat{\sigma}^2_{0,t}(z), \widehat{f}_{1,t}(\tau), \widehat{f}_{0,t}(\tau)\}$. In particular, $\widehat{\sigma}^2_{a,t}(z)$ is the sample variance of $(\mathbb{1}(Y_i \le \widehat{Q}_{a,t-1}(\tau)) - \tau)$ for $A_i = a, Z_i = z, i = 1, \dots, N_t$:

$$\widehat{\sigma}^2_{a,t}(z) = \frac{\sum_{i=1}^{N_t}\left\{\mathbf{1}\left(Y_i \le \widehat{Q}_{a,t-1}(\tau)\right) - \tau\right\}^2 \mathbf{1}(A_i = a, Z_i = z)}{\sum_{i=1}^{N_t}\mathbf{1}(A_i = a, Z_i = z)}, \quad a \in \{0,1\}. \tag{4}$$

$\widehat{f}_{a,t}(\tau)$ is an estimator of the density $f_a(\tau)$ at the end of stage t, and can be obtained via the standard kernel estimation method

$$\widehat{f}_{a,t}(\tau) = \frac{1}{N_{a,t}}\sum_{i=1}^{N_t}\mathcal{K}_{h_f}\left(\frac{Y_i - \widehat{Q}_{a,t-1}(\tau)}{h}\right)\mathbf{1}(A_i = a), \quad a \in \{0,1\}, \tag{5}$$

where $N_{a,t} = \sum_{i=1}^{N_t}\mathbf{1}(A_i = a)$, $\mathcal{K}_{h_f}(\cdot)$ is the kernel function and h_f is a bandwidth (see its data-driven choice in Section 5.1).

Statistical inference after stage T. *At the end of stage T, we construct the QTE estimator along with its variance estimator:*

$$\begin{aligned}\widehat{\alpha}_\tau &= \widehat{Q}_{1,T}(\tau) - \widehat{Q}_{0,T}(\tau),\\ \widehat{\mathbb{V}}_{\tau,FA} &= \frac{1}{N_T}\sum_{t=1}^{T}\sum_{i=N_{t-1}+1}^{N_t}(\widehat{U}_{\tau,it} + \widehat{S}_{\tau,it})^2,\end{aligned} \tag{6}$$

where

$$\widehat{U}_{\tau,it} = \frac{A_i}{\widehat{e}_t(Z_i)}\cdot\frac{\psi_\tau(Y_i - \widehat{Q}_1(\tau))}{\widehat{f}_{1,t}(\widehat{Q}_1(\tau))} - \frac{1-A_i}{1-\widehat{e}_t(Z_i)}\frac{\psi_\tau(Y_i - \widehat{Q}_0(\tau))}{\widehat{f}_{0,t}(\widehat{Q}_0(\tau))},$$

$$\widehat{S}_{\tau,it} = -\frac{1}{\sum_{i=1}^{N_t}\mathbf{1}(Z=Z_i)}\sum_{i=1}^{N_t}\left[\frac{A_i}{(\widehat{e}_t(Z))^2}\cdot\frac{\psi_\tau(Y_i-\widehat{Q}_1(\tau))}{\widehat{f}_{1,t}(\widehat{Q}_1(\tau))}+\frac{1-A_i}{(1-\widehat{e}_t(Z))^2}\cdot\frac{\psi_\tau(Y_i-\widehat{Q}_0(\tau))}{\widehat{f}_{0,t}(\widehat{Q}_0(\tau))}\right]\cdot\mathbf{1}(Z=Z_i)$$
$$\cdot(A_i-\widehat{e}_t(Z_i)).$$

Together with our theoretical results in Theorem 2, we construct a 95% level confidence interval for α_τ with $\left[\widehat{\alpha}_\tau \pm 1.96\times\widehat{\mathbb{V}}_{\tau,FA}^{\frac{1}{2}}\right]$.

3.3. Design strategy II: Multi-stage design

In this section, we consider multi-stage designs. Unlike fully adaptive designs, where the treatment assignment strategy is modified for an infinite number of times (large T), multi-stage designs allow for finitely many modifications of the treatment assignment strategy. Therefore, the multi-stage design under consideration has a small T but a large n_t, for $t \geq 1$. Although multi-stage designs only revise the treatment assignment for finitely many times, they have clear benefits in successfully handling continuous covariates and being more efficient asymptotically (as shown in Theorem 3). Without loss of generality, we start with showcasing the design strategy when $T = 2$.

In the initial stage of the experiment, we lack prior knowledge about the underlying data distribution. Therefore, we randomly assign all enrolled participants to either the treatment or control group with a pre-specified probability $e_1 = \mathbb{P}(A = 1)$, which is independent of any observed covariate information:

Stage 1. *Assign each enrolled subject $i = 1, \ldots, n_1$ to the treatment arm with probability e_1, irrespective of their covariate values, and collect observations $\{(Y_i, A_i, X_i)\}_{i=1}^{n_1}$ for these subjects.*

Then, in the second stage, we choose the treatment assignment probability $e_2(x)$ based on the individual's covariate information to minimize the variance bound in Eq. (2). This means that we need to choose an optimal $e_2(x)$ that minimizes the following variance bound, that is

$$\min_{e_2(x)}\ \mathbb{V}(e_2(x)) \triangleq \frac{\sigma_1^2(x)/f_1^2(\tau)}{r_1e_1+r_2e_2(x)}+\frac{\sigma_0^2(x)/f_0^2(\tau)}{r_1\left(1-e_1\right)+r_2\left(1-e_2(x)\right)}.$$

The closed-form solution of the above optimization problem yields an optimal propensity score:

$$e_2^\star(x) = \frac{1}{r_2}\left\{\frac{\sigma_1(x)/f_1(\tau)}{\sigma_1(x)/f_1(\tau)+\sigma_0(x)/f_0(\tau)}-r_1e_1\right\}. \tag{7}$$

This optimal $e_2^\star(x)$ yields a design that provides an estimate of the quantile treatment effect with minimum (asymptotic) variance $\mathbb{V}(e_2^\star) = V_{\texttt{Optimal}}$. By replacing the unknown quantities with their estimates, Stage 2 of our design proceeds as follows:

Stage 2. *Assign subjects $i = n_1 + 1, \ldots, n_1 + n_2$ in the second stage to treatment with probabilities $\widehat{e}_2(X_i)$:*

$$\widehat{e}_2(x) = \frac{1}{r_2}\left\{\frac{\widehat{\sigma}_1(x)/\widehat{f}_1(\tau)}{\widehat{\sigma}_1(x)/\widehat{f}_1(\tau)+\widehat{\sigma}_0(x)/\widehat{f}_0(\tau)}-r_1e_1\right\}. \tag{8}$$

Here, we estimate the conditional variance with a polynomial approximation using the stage 1 data, that is.

$$\widehat{\sigma}_a^2(x) = \frac{1}{\sum_{i=1}^{n_1}\mathbf{1}(A_i=a)}\sum_{i=1}^{n_1}\mathcal{K}_{h_\sigma}\left(\frac{(\mathbf{1}(Y_i\leq\widehat{Q}_a(\tau))-\tau)^2}{h_\sigma}\right)\mathbf{1}(A_i=a) \tag{9}$$

where h_σ is bandwidth (see their data driven choice in Section 5.1). If X is a discrete variable, we adopt the sample variance estimator in (4). Similar to the estimator of the quantile treatment effect in the fully adaptive design in Eq. (3), the one-stage quantile treatment effect $Q_a(\tau)$ is estimated with

$$\widehat{Q}_{a,1}(\tau) = \arg\ \min_{u\in\mathbb{R}}\ \frac{1}{n_1}\sum_{i=1}^{n_1}\frac{\mathbf{1}(A_i=a)}{A_ie_1+(1-A_i)\left(1-e_1\right)}\rho_\tau(Y_i-u).$$

The density estimator $\widehat{f}_a(\tau)$ of the potential outcome $Y(a)$ evaluated at $\widehat{Q}_{a,1}(\tau)$ is obtained similarly to that in Eq. (5). Here, we calculate the density estimator $\widehat{f}_a(\tau)$ with the data collected from stage 1. We estimate the optimal propensity score in Eq. (7) with its sample analogue.

Statistical inference after two stages. At the end of the second stage, we estimate $Q_a(\tau)$ with

$$\widehat{Q}_{a,2}(\tau) = \arg\min_{u\in\mathbb{R}}\ \frac{1}{N}\sum_{i=1}^{N}\frac{\mathbf{1}(A_i=a)}{A_i\widehat{e}(X_i)+(1-A_i)\left(1-\widehat{e}(X_i)\right)}\rho_\tau(Y_i-u),$$

where $\widehat{e}(x) = r_1e_1 + r_2\widehat{e}_2(x)$, and $N = n_1 + n_2$. We then construct the QTE estimator along with its variance estimator and confidence interval using:

$$\widehat{\alpha}_\tau = \widehat{Q}_{1,2}(\tau)-\widehat{Q}_{0,2}(\tau),\quad \widehat{\mathbb{V}}_{\tau,TS} = \frac{1}{N}\sum_{i=1}^{N}(\widehat{U}_{\tau,i}+\widehat{S}_{\tau,i})^2, \tag{10}$$

where

$$\widehat{U}_{\tau,i} = \frac{A_i}{\widehat{e}(X_i)} \cdot \frac{\psi_\tau(Y_i - \widehat{Q}_1(\tau))}{\widehat{f}_1(\widehat{Q}_1(\tau))} - \frac{1 - A_i}{1 - \widehat{e}(X_i)} \frac{\psi_\tau(Y_i - \widehat{Q}_0(\tau))}{\widehat{f}_0(\widehat{Q}_0(\tau))},$$

$$\widehat{S}_{\tau,i} = -\widehat{\mathbb{E}}\left[\frac{A}{\widehat{e}(X)^2} \cdot \frac{\psi_\tau(Y - \widehat{Q}_1(\tau))}{\widehat{f}_1(\widehat{Q}_1(\tau))} + \frac{1 - A}{(1 - \widehat{e}(X))^2} \cdot \frac{\psi_\tau(Y - \widehat{Q}_0(\tau))}{\widehat{f}_0(\widehat{Q}_0(\tau))} \mid X = X_i\right] \cdot (A_i - \widehat{e}(X_i)).$$

Similar to Section 3.2, the conditional expectation in $\widehat{S}_{\tau,it}$ can be estimated by a polynomial approximation. A 95% level confidence interval for α_τ can be constructed with $\left[\widehat{\alpha}_\tau \pm 1.96 \times \widehat{\mathbb{V}}^{\frac{1}{2}}_{\tau,TS}\right]$.

In a more general multi-stage design with $T > 2$, for the final estimator of the QTE to attain the optimal variance lower bound in Eq. (1), the optimal propensity scores $(e_1^\star(x), \ldots, e_T^\star(x))$ should satisfy the following equation:

$$\sum_{t=1}^{T} r_t e_t^*(x) = e^*(x) = \frac{\sigma_1(x)/f_1(\tau)}{\sigma_1(x)/f_1(\tau) + \sigma_0(x)/f_0(\tau)}$$

where $n_t/N \to r_t$ as $N \to \infty$, $\sum_{t=1}^{T} r_t = 1$ and r_t is a positive constant bounded away from 0 and 1. Similar to the two-stage design, in the first stage, we can randomly assign the treatment with $e_1(x) = e_1 \in (0, 1)$ independent of any observed covariate information. For subsequent stages $s = 2, \ldots, T$, we update the estimates of $\sigma_a(x)$ and $f_a(\tau)$, for $a \in \{0, 1\}$, using the cumulative data up to stage $s - 1$ following the derivation provided in Eqs. (5) and (9). Treatment assignments in these stages are based on the estimated propensity score, determined by the equation:

$$\sum_{t=1}^{s} r_t \widehat{e}_t(x) = \frac{\widehat{\sigma}_1(x)/\widehat{f}_1(\tau)}{\widehat{\sigma}_1(x)/\widehat{f}_1(\tau) + \widehat{\sigma}_0(x)/\widehat{f}_0(\tau)},$$

where $\widehat{\sigma}_a(x)$ and $\widehat{f}_a(\tau)$ denote the estimates of $\sigma_a(x)$ and $f_a(\tau)$ using the accumulated data up to stage $s - 1$. Finally, we estimate $Q_a(\tau)$ with

$$\widehat{Q}_{a,T}(\tau) = \arg\min_{u \in \mathbb{R}} \frac{1}{N} \sum_{i=1}^{N} \frac{\mathbf{1}(A_i = a)}{A_i \widehat{e}(X_i) + (1 - A_i)\left(1 - \widehat{e}(X_i)\right)} \rho_\tau(Y_i - u),$$

where $\widehat{e}(x) = \sum_{t=1}^{T} r_t \widehat{e}_t(x)$. The corresponding QTE estimator, along with its variance estimator and confidence interval, can be estimated according to Eq. (10).

4. Theoretical property

In this section, we establish the theoretical guarantees for the QTE estimators obtained from both the two-stage and fully adaptive design procedures. Specifically, we prove that both designs provide root-N consistent estimators for α_τ, and these estimators are asymptotically normal with asymptotic variances that attain the semiparametric efficiency bound. Additionally, we demonstrate the consistency of the estimated optimal propensity score in both designs. Lastly, we provide an efficiency comparison between these two proposed designs.

As suggested by the data collection mechanism in the previous sections, we work under the following assumptions:

Assumption 1 (*Unconfoundedness Assumption Reached by the Design*)**.** (i) For the multi-stage design, given the covariate X, the potential outcomes are jointly independent from the treatment A: $\big(Y(1), Y(0)\big) \perp\!\!\!\perp A \mid X$; (ii) For the fully adaptive design, given the discretized covariate Z, the potential outcomes are jointly independent from the treatment A: $\big(Y(1), Y(0)\big) \perp\!\!\!\perp A \mid Z$.

Assumption 2 (*Positivity*)**.** For $t = 1, \ldots, T$, the propensity scores are bounded away from zero and one.

Assumption 1 is satisfied by the adaptive design procedure where the assignment to the treatment is completely at random within each strata of X. Assumption 2 states that for almost all values of X, both treatment assignment levels have a positive probability of occurrence.

Assumption 3 (*Uniqueness of Quantiles*)**.** For $a = 0, 1$, $Y(a)$ is a continuous random variable with support in $\mathbb{R}$ and where the following statements apply:

(i) There are nonempty sets $\mathcal{Y}_1$ and $\mathcal{Y}_0$, such that $\mathcal{Y}_a = \{\tau \in (0, 1); \mathbb{P}\left[Y(a) \le Q_a(\tau) - c\right] < \mathbb{P}\left[Y(a) \le Q_a(\tau) + c\right], \forall c \in \mathbb{R}, c > 0\}$.

(ii) The density function of $Y(a)$, $a \in \{0, 1\}$ are uniform continuous and bounded away form zero. The bandwidths h_f, h_σ used in $\widehat{f}_a(\tau)$, $\widehat{\sigma}_a^2(x)$ satisfy $h_f \propto N^{-1/5}$ and $h_\sigma \propto N^{-1/5}$, respectively.

Assumption 4 (*Kernel Function*)**.** The kernel function $\mathcal{K}_h(x) \doteq K(x/h)/h$ with $\int K(x)dx = 1, \int xK(x)dx = 0, \int x^2K(x)dx < \infty, \int |K(x)|dx < \infty, \int |x \log |x||^{1/2} dK(x) < \infty$. And $K(x) \to 0$ as $x \to \infty$.

Assumption 5 (*Continuous*). If X is a continuous random variable, for $a = 0, 1$, $\mathbb{E}(\psi_\tau(Y(a)-Q_a(\tau))|X = x)$ is continuously differentiable with respect to x for every x in $\mathcal{X}$ and the variance $\sigma_1^2(x)$ and $\sigma_0^2(x)$ are twice continuously differentiable. The bandwidths h_σ used in $\widehat{\sigma}_a^2(x)$ satisfy $h_\sigma \propto N^{-1/(d+4)}$. The bandwidths h_S used in $\widehat{S}_{\tau,i}$ also satisfy $h_S \propto N^{-1/(d+4)}$.

Assumption 3-(i) is the same as Assumption 2-(i) in Firpo (2007), which ensures that QTE become estimable from the data on (Y, A, X). are commonly assumed in nonparametric regression (Fan and Gijbels, 1996). Here we directly assume the bandwidth has the optimal order to ensure the consistency and efficiency of the kernel smooth estimators.

Theorem 1 (*Two-stage Design*). *Let* $e_2^*(x) = \lim_{n_2\to\infty} \widehat{e}_2(x)$ *and* $e^*(x) = r_1 e_1 + r_2 e_2^*(x)$. *Under Assumptions 1(i) and 2–5 we have*

$$\sqrt{N}(\widehat{\alpha}_\tau - \alpha_\tau) \xrightarrow{d} \mathcal{N}(0, \mathbb{V}_{\tau,TS}), \tag{11}$$

where

$$\mathbb{V}_{\tau,TS} = \mathbb{E}\left[\left(\frac{\sigma_1(X)}{f_1(\tau)} + \frac{\sigma_0(X)}{f_0(\tau)}\right)^2 + \left[\frac{\mathbb{E}\left[\psi_\tau\left(Y - Q_1(\tau)\right)|X, A = 1\right]}{f_1(\tau)} - \frac{\mathbb{E}\left[\psi_\tau\left(Y - Q_0(\tau)\right)|X, A = 0\right]}{f_0(\tau)}\right]^2\right]. \tag{12}$$

And the estimator of $\widehat{\mathbb{V}}_\tau$ *in (10) are consistent, i.e.* $\widehat{\mathbb{V}}_\tau = \mathbb{V}_{\tau,TS} + o_p(1)$.

This asymptotic variance attains the optimal variance lower bound of QTE estimators in Eq. (2). A consistent estimator of $\mathbb{V}_\tau$ can be obtained via Eq. (10). Next, we state the following theoretical results of the fully adaptive design.

Theorem 2 (*Fully Adaptive Design*). *The discretized covariate vector* Z *follows a multinomial distribution with finite support. Under Assumptions 1(ii) and 2–4, we have*

$$\sup_{z\in\mathcal{Z}} |\widehat{e}_T(z) - e_T(z)| = o_p(1), \quad \sqrt{N_T}(\widehat{\alpha}_\tau - \alpha_\tau) \xrightarrow{d} \mathcal{N}(0, \mathbb{V}_{\tau,FA}), \text{ as } N_T \to \infty,$$

where

$$\mathbb{V}_{\tau,FA} = \mathbb{E}\left[\left(\frac{\sigma_1(Z)}{f_1(\tau)} + \frac{\sigma_0(Z)}{f_0(\tau)}\right)^2 + \left[\frac{\mathbb{E}\left[\psi_\tau\left(Y - Q_1(\tau)\right)|Z, A = 1\right]}{f_1(\tau)} - \frac{\mathbb{E}\left[\psi_\tau\left(Y - Q_0(\tau)\right)|Z, A = 0\right]}{f_0(\tau)}\right]^2\right].$$

Under a fully adaptive design setting, the proposed QTE estimator also attains the optimal variance lower bound $\mathbb{V}_{\texttt{optimal},\tau}$. A consistent estimator $\widehat{\mathbb{V}}_\tau$ can be calculated via Eq. (6) in Section 3.2.

Based on the theoretical results presented above, it can be concluded that when the original covariates are discrete and no discretization is required for implementing the proposed design, both the fully adaptive design and the two-stage design achieve the same efficiency bound. However, in finite samples, as will be demonstrated later, the QTE estimator obtained from the fully adaptive design has a smaller variance than the one obtained from the two-stage design. We conjecture that this difference in finite sample performance is due to the fully adaptive design being able to revise the randomization probability an infinite number of times, thereby providing more robustness.

Finally, when the original covariates are continuous and discretization is used in the fully adaptive design, the following theorem demonstrates that the two-stage design is more asymptotically efficient than the fully adaptive design. This implies that the two-stage design has a smaller asymptotic variance than that obtained from the fully adaptive design. We shall verify the following results in simulation studies in the following section:

Theorem 3 (*Efficiency Gain of the Two-Stage Design by Handling Continuous Covariates*). *Assuming both Assumptions 1 and 2 hold, and with the fully adaptive design using discretized original covariates to adaptively assign treatment while the two-stage design uses the original covariates for treatment assignment, we obtain the following:* $\mathbb{V}_{\tau,FA} \geq \mathbb{V}_{\tau,TS}$.

5. Simulation studies and a synthetic case study

This section presents an investigation into the performance of our proposed designs for estimating quantile treatment effects at different levels. We summarize our main observations as follows: Firstly, our proposed adaptive experimental designs have the capability to achieve a smaller variance in estimating the QTE, when compared to the QTE obtained from completely randomized experiments. This confirms our theoretical results provided in Theorems 1 and 2 in Section 4. Secondly, upon comparing the performance of the two-stage design and the fully adaptive design, we observe that their performances are similar when the covariate is discrete and the sample size is large. However, the fully adaptive design tends to provide QTE estimators with smaller variance when the sample size is small. Lastly, the two-stage design is more efficient in estimating QTE when dealing with continuous covariates. This observation is in line with our theoretical result presented in Theorem 3.

5.1. Simulation design

In this section, we evaluate the performance of our proposed design strategies with two simulation setups: one with a discrete covariate and the other with a continuous covariate. For the continuous covariate in a fully adaptive design, we discretize it into multiple strata and adaptively assign treatments within each stratum.

When the covariate X is a discrete random variable, we generate it from a Bernoulli distribution with probability $\mathbb{P}(X = 1) = 0.4$. The potential outcomes are generated from two linear models $Y(0) = X - 6 + \epsilon_0$ and $Y(1) = -0.5X + \epsilon_1$, respectively, where $\epsilon_0 \sim \chi^2_{v_X}, v_X = 5 - X$ and $\epsilon_1 \sim N(0, 2 - 0.5X)$. When X is a continuous random variable, we generate X from the uniform distribution $U(0, 1)$. The potential outcomes are generated from $Y(0) = X + e^{2X} - 9 + \epsilon_0$ and $Y(1) = 0.5X + \cos(2\pi x) + \epsilon_1$ where $\epsilon_0 \sim \chi^2_5$ and $\epsilon_1 \sim N(0, 2 + 0.5X)$. As fully adaptive design cannot directly work with a continuous covariate, we discretize X into two strata: $[0, 0.5)$ and $[0.5, 1]$. In both discrete and continuous cases, we set two different sample sizes $N = 1000$ and 5000, and the results are summarized over 1000 replications.

In the above simulation setup, we simulate potential outcomes with two characteristics: first, they follow skewed distributions that deviate from the normal distribution, and second, their conditional distributions vary with the covariate X. This approach allows us to, on the one hand, comprehensively assess the robustness and adaptability of our designs under diverse possible realistic scenarios and, on the other hand, showcase the benefit of analyzing the QTE. This is because the current data-generating process enables us to, on the one hand, thoroughly evaluate the robustness and adaptability of our designs across a variety of plausible real-world scenarios and, on the other hand, highlight the advantages of analyzing the QTE. Because QTE offers a more nuanced understanding of treatment impacts across different quantile levels in the outcome distribution, it provides critical insights in cases where the outcome distribution is skewed (so that QTE has values different from ATE) or has distribution depending on the covariates (so that the efficiency bound of the QTE differs across different quantile levels).

In the two-stage design, we set $e_1 = 0.5$ and $n_1 \in \{0.2N, 0.5N\}$ for the first stage. In contrast, for the fully adaptive design, we set $e_1 = 0.5$ and $n_1 = 0.2N$ for the first stage. For subsequent stages, we set $n_t = 5$, indicating that we revise the treatment allocation strategy for every 5 individuals. As later we set $T \in \{200, 1000\}$, which is a relatively large number compared to $n_t = 5$, it might be reasonable to consider that we are operating within the fully adaptive setting.

We compare the performances of QTE estimators obtained in (i) a completely randomized experiment where the treatment is assigned randomly to each participant with a 1/2 probability throughout the experiment, (ii) a naive method that ignores the treatment allocation based on covariate and takes the quantiles of a simple difference between the observed outcomes in treated and control groups, (iii) a two-stage design as shown in Eq. (10), and (iv) a fully adaptive design as shown in Eq. (6). We measure the performance of various estimators in terms of their biases (computed as the differences between the mean of the Monte Carlo estimates of α_τ and the true QTE), variances (computed as the variances of the Monte Carlo estimates of α_τ), mean squared errors (MSE) and coverage probabilities of the confidence intervals with the nominal coverage probability of 95%.

Throughout our simulation study, we choose the Gaussian Kernel function for all nonparametric estimators. The bandwidth is chosen as $h_x = 1.5 \times sd(x) \times n^{-1/5}$ where n denotes the sample size, and $sd(x)$ is the sample standard deviation of $x_1, \ldots, x_n$.

5.2. Simulation results

The simulation results of the above estimators, including those from both two-stage and fully adaptive designs, are reported in Tables 1 and 2, respectively. These results pertain to the case of a discrete X, with $N = 1000$, $N = 5000$ and $\tau \in \{0.3, 0.5, 0.8\}$, respectively. As expected, the estimators from both the two-stage and fully adaptive designs outperform the one obtained from a completely randomized design. The variance of the latter is greater than those from our proposed designs, demonstrating the efficiency gain from our proposed designs.

Additionally, when X is discrete, the QTE estimator obtained from a fully adaptive design has a smaller variance than the one from a two-stage design when the sample size $N = 1000$. The naive estimator, which does not take into account the nature of the covariate-based design, has the worst performance in all metrics. Cross-comparison between Tables 1 and 2 reveals that almost all evaluation metrics diminish rapidly as the sample size increases, except for the bias of the naive estimator.

Results for a continuous covariate are summarized in Tables 3 and 4. It can be clearly seen from these two tables that the standard error in QTE from a fully adaptive design is larger than that from a two-stage design. This observation suggests that discretization of the covariate in a fully adaptive design may lead to a loss of statistical efficiency, which corroborates the results in Theorem 3 in Section 4.

5.3. Synthetic case study

In this section, we illustrate the performance of our adaptive design algorithm using data from a direct mail fundraising experiment reported in Karlan and List (2007). The original design of this study is one-stage with an aim to test the effectiveness of a matching grant on charitable giving. In our synthetic experiment, we are interested in how we should allocate treatments, conditional on covariates, in the second wave of a two-stage design or in the subsequent stages of a fully adaptive design. Besides, we would like to investigate the effect of a matching grant effect in the upper tail of the donation amount distribution. According to our preliminary analysis, the donation amount distribution is right-skewed and QTE is clearly more relevant than the conventionally used ATE.

Table 1
Performances of different estimators in the case where X is discrete. The total sample size is $N = 1000$. For the two-stage design and the fully adaptive design, we set $n_1 = 0.2N$.

τ	Design	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
0.3	Completely random	129.47	3.00	16.60	–
	Naive	134.06	49.65	20.25	–
	Two stage	128.48	10.93	16.46	0.97
	Fully adaptive	116.02	23.58	13.88	0.97
0.5	Completely random	152.98	3.05	23.18	–
	Naive	159.50	72.27	30.41	–
	Two stage	138.61	19.51	19.40	0.96
	Fully adaptive	123.27	34.98	16.42	0.96
0.8	Completely random	237.53	111.02	68.18	–
	Naive	248.60	13.10	62.00	–
	Two stage	229.82	18.68	52.63	
	Fully adaptive	214.71	35.54	46.90	0.94

Table 2
Performance of each estimator with discrete covariate with $N = 5000$. For the two-stage design, we set $n_1 = 0.5N$.

τ	Design	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
0.3	Completely random	65.14	−4.67	4.26	–
	Naive	65.27	6.24	4.30	–
	Two stage	63.99	−3.22	4.10	0.96
	Fully adaptive	64.30	−5.26	4.16	0.96
0.5	Completely random	77.99	4.15	6.09	–
	Naive	75.06	18.58	5.97	–
	Two stage	74.4	3.27	5.54	0.95
	Fully adaptive	75.77	2.37	5.74	0.95
0.8	Completely random	119.60	−5.89	14.32	–
	Naive	106.28	17.90	11.60	–
	Two stage	104.05	−3.54	10.83	0.96
	Fully adaptive	109.08	0.55	11.89	0.96

Table 3
Performance of each estimator with continuous case with $N = 1000$.

τ	Design	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
0.3	Completely random	216.4	0.2	46.8	–
	Naive	222.8	−43.3	51.5	–
	Two stage	206.1	4.9	42.5	0.96
	Fully adaptive	207.4	0.6	43.0	0.96
0.5	Completely random	230.3	6.7	53.1	–
	Naive	230.9	−33.9	54.5	–
	Two stage	216.8	−0.8	47.0	0.96
	Fully adaptive	218.4	6.2	47.7	0.96
0.8	Completely random	301.7	6.8	91.1	–
	Naive	287.3	−37.2	83.9	–
	Two stage	276.8	−4.5	76.6	0.94
	Fully adaptive	278.4	−0.2	77.5	0.95

The original dataset contains 50,083 observations among which two-thirds (33,396) received a matching grant offer ($A = 1$), and one-third (16,687) received the same solicitation but without mention of a matching grant ($A = 0$). In this case study, we applied our algorithm without the constraint that the overall treatment probability be two-thirds. We choose to use the highest previous amount (HPA) donated as the covariate (X) for the design, which is featured prominently in the original analysis. In a fully adaptive design, we discretize HPA into two categories based on its quartiles: (i) $X \leq \$45$; (ii) $X > \$45$. In a two-stage design, we can either use its original continuous form or the discretized version. To compare the efficiency, we adopt both ways in a two-stage design. The outcome of interest (Y) is the individual's donation amount in dollars after receiving the direct mail solicitation. In particular, we estimate the 30%, 50%, and 80% quantiles of a matching grant effect.

We assume that the observations of this large-scale field experiment are observed sequentially over time, and we can adaptively revise the probability of receiving a matching grant offer. To be consistent with our simulation setup, we use a subsample of the original full dataset with sample sizes $N = 1000$ and 5000. In a fully adaptive design, we initialize the experiment with $n_1 = 0.2N$ allocated to two arms with equal probability in stage 1, and $n_t = 5$ for the following stages. In a two-stage design, we enroll $n_1 = 0.5N$

Table 4
Performance of each estimator with continuous case with $N = 5000$.

τ	Design	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE($\times 10^{-3}$)	Coverage probability
0.3	Completely random	100.43	−6.79	10.13	–
	Naive	102.59	−38.49	12	–
	Two stage	96.4	−5.62	9.32	0.96
	Fully adaptive	96.46	−5.96	9.34	0.95
0.5	Completely random	102.55	−6.71	10.56	–
	Naive	103.82	−35	12	–
	Two stage	96.35	−5.06	9.31	0.96
	Fully adaptive	99.12	−4.49	9.84	0.95
0.8	Completely random	135.6	5.8	18.42	–
	Naive	128.56	−24.57	17.13	–
	Two stage	123.65	1.54	15.29	0.94
	Fully adaptive	127.43	3.11	16.24	0.94

Table 5
Performances of different estimators in the synthetic case study. The total sample size is $N = 1000$. "MSE" denotes the mean squared error.

τ	Design	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
0.3	Completely random	96.05	51.71	11.9	–
	Discrete two-stage	74.44	35.64	6.81	0.93
	Continuous two-stage	71.10	33.06	6.15	0.94
	Discrete fully adaptive	77.03	36.81	7.29	0.91
0.5	Completely random	160.44	3.79	25.75	–
	Discrete two-stage	130.61	−2.40	17.06	0.88
	Continuous two-stage	123.76	−1.49	15.32	0.93
	Discrete fully adaptive	171.94	−20.06	29.96	0.82
0.8	Completely random	250.94	−34.61	64.15	–
	Discrete two-stage	221.07	−36.29	50.18	0.91
	Continuous two-stage	214.73	−33.14	47.20	0.92
	Discrete fully adaptive	222.69	−41.19	51.28	0.91

to the first stage and $n_2 = 0.5N$ to the second one. We compare four different design strategies: (i) completely randomized design where the treatment is assigned randomly to each participant with a $1/2$ probability throughout the experiment, (ii) two-stage design with the discretized covariate, (iii) two-stage design with the original continuous covariate, and (iv) fully adaptive design with the discretized covariate.

In both designs, we use the observed data $\{(Y_i, A_i, X_i)\}_{i=1}^{n_1}$ to initialize the process. For later stages, given the covariate information $\{X_i\}_{i=n_1+1}^{N}$, we assign a certain subject i to either the treatment or control arm according to the optimal propensity score in that stage. The potential outcomes are simulated via $Y(1) = \beta_{01} + \beta_{11}X + \epsilon_1$ and $Y(0) = \beta_{00} + \beta_{10}X + \epsilon_0$ where the regression coefficients $(\beta_{01}, \beta_{11}, \beta_{00}, \beta_{10})$ are estimated with the original full dataset. To better mimic the right-skewed outcome distributions in the real dataset, the error terms ϵ_1 and ϵ_0 are generated from a mixture of chi-squared distributions. Results from this synthetic case study are summarized in Tables 5 and 6.

Similar to our observations in the simulation studies, bias, standard error, and MSE all decrease when we increase the total sample size from $N = 1000$ to 5000. Furthermore, as we expect, our proposed adaptive rule will lead to a smaller asymptotic variance than the completely randomized design. In two-stage designs, there is clear evidence of loss of information when the covariates are discretized: the bias, the standard error, and the MSE of the QTE estimates in the continuous case are smaller than those in the discrete one. More specifically, the efficiency gain varies from 2.87% to 5.24% when $N = 1000$ while it ranges from 2.27% to 17.07% when $N = 5000$. The QTE estimates obtained in a fully adaptive design are generally less efficient potentially due to the discretization of the covariates.

6. Discussion

In this manuscript, we proposed two new CARA designs for estimating quantile treatment effects that are easy to implement and enjoy strong theoretical guarantees. Some points merit further research in the future.

First, when the covariates are discrete, both designs attain the same semiparametric efficiency bound asymptotically, but they exhibit different finite sample performances. While we believe that the observed difference is due to the frequent update nature of the fully adaptive experiments, a rigorous justification of this observation may require a careful investigation of multi-stage and fully adaptive designs as the number of stages slowly increases with the sample size.

Second, in the existing literature, CARA is often proposed to have both efficiency and ethical advantages. Efficiency generally refers to the power of detecting treatment differences in clinical trials, and ethics often concerns the number of treatment successes

Table 6
Performances of different estimators in the synthetic case study. The total sample size is $N = 5000$. "MSE" denotes the mean squared error.

τ	Design	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
0.3	Completely random	30.1	8.55	0.98	–
	Discrete two stage	21.56	5.05	0.49	0.96
	Continuous two stage	17.88	3.61	0.33	0.96
	Discrete fully adaptive	23.58	6.05	0.59	0.95
0.5	Completely random	67.61	1.92	4.57	–
	Discrete two-stage	54.14	4.38	2.95	0.90
	Continuous two-stage	50.88	1.91	2.59	0.96
	Discrete fully adaptive	64.25	−1.37	4.13	0.89
0.8	Completely random	111.73	−9.17	12.55	–
	Discrete two-stage	94.08	−14.26	9.05	0.93
	Continuous two-stage	91.94	−23.3	8.99	0.93
	Discrete fully adaptive	95.45	−22.19	9.59	0.92

(for binary outcomes) or the total response (for continuous outcomes) in the trial. As the proposed design aims to maximize the statistical power of detecting the quantile treatment effect, it does not necessarily improve the ethical aspect of the design. Guided by Hu et al. (2015), we believe that it is possible to integrate measurements of both efficiency and ethics into the current design by revising the treatment allocation probability as

$$\widehat{e}_t(z) = \frac{\widehat{\sigma}_{1,t-1}(z)/\widehat{f}_{1,t-1}(\tau)\cdot\widehat{\theta}_{1,t-1}(z)}{\widehat{\sigma}_{1,t-1}(z)/\widehat{f}_{1,t-1}(\tau)\cdot\widehat{\theta}_{1,t-1}(z)+\widehat{\sigma}_{0,t-1}(z)/\widehat{f}_{0,t-1}(\tau)\cdot\widehat{\theta}_{0,t-1}(z)}.$$

The choice of the ethics measurements $\widehat{\theta}_{1,t-1}(z)$ and $\widehat{\theta}_{0,t-1}(z)$ greatly depends on the specific target of the experiment. We hope to incorporate some ethical considerations of our design in our future work.

Third, as indicated in Tables B.3–B.4, the fully adaptive design may not perform well when the sample sizes, N, are relatively small. This might be because there are not sufficient samples to accurately estimate the variance, $\sigma^2_{a,t}(z)$, even when assuming discrete covariates. The situation becomes even more challenging when the covariates are continuous, as the number of effective samples for variance estimation decreases further. Consequently, a fully adaptive design for continuous covariates warrants further studies.

Lastly, we hope to comment that the existing literature documents an infrequent application of the QTE in designing randomized experiments and clinical trials; we conjecture that this can be attributed to the fact that the QTE presents complexities in interpretation. Unlike the ATE, QTE does not provide a single summary measure but rather a distribution of effects. Additionally, establishing theoretical guarantees of QTE requires assumptions about the distribution of potential outcomes, which can be difficult to verify in practice. We hope that this manuscript serves as a tentative effort to promote the use of QTE in designing randomized experiments.

Acknowledgments

The authors are grateful to the editor, the associate editor, and two anonymous referees for their comments that have greatly improved this paper. The authors would also like to thank helpful discussions with Xinwei Ma on the theoretical properties of the proposed design. Li was partially supported by the National Key R &D Program of China (Grant 2022ZD0114801) and the National Natural Science Foundation of China (Grant 12071233).

Appendix A. Technical proofs

Proof of Theorem 1. First, if X_i has finite support, we will show that $\sup_{x\in\mathcal{X}}|\widehat{e}_2(x)-e_2(x)| = O_p(N^{-2/5})$. Obviously, $\widehat{\sigma}^2_a(x)$ is a $\sqrt{N}$-consistent estimator of $\sigma^2_a(x)$ and $\widehat{f}_a(\tau)$ is a $N^{-2/5}$ consistent estimator of $f_a(\tau)$ by setting the bandwidth with $O(N^{-1/5})$. Then, $\widehat{e}_2(x)$ is also a $N^{-2/5}$ consistent estimator of $e_2(x)$. Additionally, $\mathcal{X}$ is a finite set. So we obtain that $\sup_{x\in\mathcal{X}}|\widehat{e}_2(x)-e_2(x)| = O_p(N^{-2/5})$.

Second, if X_i is a continuous random variable, we adopt the nonparametric estimator of $\widehat{\sigma}^2_a(x)$. By Theorem 2.3 in Li et al. (2016), we have $\sup_{x\in\mathcal{X}}|\widehat{\sigma}^2_a(x)-\sigma^2_a(x)| = O_p\left(\frac{\sqrt{\log N}}{N^{-\frac{2}{d+4}}}\right)$ where d is the dimension of X_i. Similarly, we have $\widehat{f}_a(\tau)$ is a $N^{-2/5}$ consistent estimator of $f_a(\tau)$. So we obtain that $\sup_{x\in\mathcal{X}}|\widehat{e}_2(x)-e_2(x)| = O_p\left(\frac{\sqrt{\log N}}{N^{-\frac{2}{d+4}}}\right)$.

Consequently, similar to the proof of Theorem 1 in Firpo (2007), we can show that

$$\sqrt{N}(\widehat{\alpha}_\tau-\alpha_\tau) = \frac{1}{\sqrt{N}}\sum_{i=1}^{N}\phi_\tau(Y_i, A_i^*, X_i)+o_p(1)$$

where

$$\begin{aligned}
\phi_\tau(y,a,x) =& \phi_{1,\tau}(y,a,x) - \phi_{0,\tau}(y,a,x),\\
\phi_{1,\tau}(y,a,x) =& \frac{a}{e^*(x)}\frac{\psi_\tau(y-Q_1(\tau))}{f_1(\tau)} - \frac{a-e^*(x)}{e^*(x)f_1(\tau)}E(\psi_\tau(y-Q_1(\tau))|x,A=1),\\
\phi_{0,\tau}(y,a,x) =& \frac{1-a}{1-e^*(x)}\frac{\psi_\tau(y-Q_0(\tau))}{f_0(\tau)} + \frac{a-e^*(x)}{1-e^*(x)f_0(\tau)}E(\psi_\tau(y-Q_0(\tau))|x,A=1),\\
A_i^* =& \mathbf{1}(U_i \le e_1),\ i=1,\dots,n_1,\ A_i^* = \mathbf{1}(U_i \le e_2^*(X_i)),\ i=n_1+1,\dots,N.
\end{aligned}$$

U_i's are iid random variables that follow Uniform(0,1), and are independent of X and Y. Therefore, we can write

$$\sqrt{N}(\hat{\alpha}_\tau - \alpha_\tau) = \frac{\sqrt{n_1}}{\sqrt{N}}L_1 + \frac{\sqrt{n_2}}{\sqrt{N}}L_2 + o_p(1)$$

where

$$L_1 = \frac{1}{\sqrt{n_1}}\sum_{i=1}^{n_1}\phi(Y_i, A_i^*, X_i),\ L_2 = \frac{1}{\sqrt{n_2}}\sum_{i=n_1+1}^{N}\phi(Y_i, A_i^*, X_i).$$

By the Central Limit Theorem, we obtain that

$$L_1 \xrightarrow{d} \mathcal{N}(0, V_{J1}),\quad L_2 \xrightarrow{d} \mathcal{N}(0, V_{J2}),$$

where

$$\begin{aligned}
V_{L1} =& \mathbb{E}\left[\frac{e_1\sigma_1^2(X)}{f_1^2(\tau)e^*(X)^2} + \frac{(1-e_1)\sigma_0^2(X)}{f_0^2(\tau)\big(1-e^*(X)\big)^2}\right.\\
&\left. + \left[\frac{\mathbb{E}\big[\psi_\tau\big(Y-Q_1(\tau)\big)|X,A=1\big]}{f_1(\tau)} - \frac{\mathbb{E}\big[\psi_\tau\big(Y-Q_0(\tau)\big)|X,A=0\big]}{f_0(\tau)}\right]^2\right],\\
V_{L2} =& \mathbb{E}\left[\frac{e_2^*(X)\sigma_1^2(X)}{f_1^2(\tau)e^*(X)^2} + \frac{(1-e_2^*(X))\sigma_0^2(X)}{f_0^2(\tau)\big(1-e^*(X)\big)^2}\right.\\
&\left. + \left[\frac{\mathbb{E}\big[\psi_\tau\big(Y-Q_1(\tau)\big)|X,A=1\big]}{f_1(\tau)} - \frac{\mathbb{E}\big[\psi_\tau\big(Y-Q_0(\tau)\big)|X,A=0\big]}{f_0(\tau)}\right]^2\right].
\end{aligned}$$

Noting that L_1 and L_2 are independent of each other and by the definition of $e^*(x)$, we obtain the result.

Next, because $\hat{e}(X)$, $\widehat{Q}_a(\tau)$, $\widehat{f}_a$ are all consistent estimators, we will show the consistency of $\widehat{\mathbb{V}}_{\tau,TS}$ by taking the same procedure as the proof of Theorem 2 in Firpo (2007).

Define

$$\begin{aligned}
\widetilde{U}_{\tau,i} =& \frac{A_i}{\hat{e}(X_i)}\cdot\frac{\psi_\tau(Y_i-Q_1(\tau))}{f_1(Q_1(\tau))} - \frac{1-A_i}{1-\hat{e}(X_i)}\frac{\psi_\tau(Y_i-Q_0(\tau))}{f_0(Q_0(\tau))},\\
\widetilde{S}_{\tau,i} =& -\widehat{\mathbb{E}}\left[\frac{A}{\hat{e}(X)^2}\cdot\frac{\psi_\tau(Y-Q_1(\tau))}{f_1(Q_1(\tau))} + \frac{1-A}{(1-\hat{e}(X))^2}\cdot\frac{\psi_\tau(Y-Q_0(\tau))}{f_0(Q_0(\tau))}\mid X=X_i\right]\cdot(A_i-\hat{e}(X_i)).
\end{aligned}$$

Here $\widetilde{U}_{\tau,i}$ and $\widetilde{S}_{\tau,i}$ use the true quantiles and the true density of potential outcomes. Then, we have

$$\begin{aligned}
|\widehat{\mathbb{V}}_{\tau,TS} - \mathbb{V}_{\tau,TS}| \le& \left|\frac{1}{N}\sum_{i=1}^{N}\left\{(\widehat{U}_{\tau,i}+\widehat{S}_{\tau,i})^2 - (\widetilde{U}_{\tau,i}+\widetilde{S}_{\tau,i})^2\right\}\right| + \left|\frac{1}{N}\sum_{i=1}^{N}(\widetilde{U}_{\tau,i}+\widetilde{S}_{\tau,i})^2 - \mathbb{V}_\tau\right|\\
\doteq& J_1 + J_2.
\end{aligned}$$

Taking the same procedure as the proof of Theorem 2 in Hirano et al. (2003), using $\frac{\psi_\tau(Y_i-Q_1(\tau))}{f_1(Q_1(\tau))}$ and $\frac{\psi_\tau(Y_i-Q_0(\tau))}{f_0(Q_0(\tau))}$, which are bounded functions of Y, instead of Y itself as in their case, we can easily obtain that $J_2 = o_p(1)$. So we only need to show that $J_1 = o_p(1)$.

Then, by triangle inequality, we have

$$\begin{aligned}
&\left|\frac{1}{N}\sum_{i=1}^{N}\left\{(\widehat{U}_{\tau,i}+\widehat{S}_{\tau,i})^2 - (\widetilde{U}_{\tau,i}+\widetilde{S}_{\tau,i})^2\right\}\right|\\
\le& \left|\frac{1}{N}\sum_{i=1}^{N}\left\{\widehat{U}_{\tau,i}^2 - \widetilde{U}_{\tau,i}^2\right\}\right| + \left|\frac{1}{N}\sum_{i=1}^{N}\left\{\widehat{S}_{\tau,i}^2 - \widetilde{S}_{\tau,i}^2\right\}\right| + \left|\frac{2}{N}\sum_{i=1}^{N}\left\{\widehat{U}_{\tau,i}\widehat{S}_{\tau,i} - \widetilde{U}_{\tau,i}\widetilde{S}_{\tau,i}\right\}\right|\\
\doteq& J_{11} + J_{12} + J_{13}.
\end{aligned}$$

Let $g_{a,\tau}(Y_i) = \frac{\psi_\tau(Y_i - Q_a(\tau))}{f_a(Q_a(\tau))}$ and $\hat{g}_{a,\tau}(Y_i) = \frac{\psi_\tau(Y_i - \hat{Q}_a(\tau))}{\hat{f}_a(\hat{Q}_a(\tau))}$ for $a = 0, 1$. Then,

$$J_{11} \le \left| \frac{1}{N} \sum_{i=1}^{N} \frac{e^\star(X_i)}{\hat{e}^2(X_i)} \cdot \frac{A_i}{e^\star(X_i)} \cdot \left(\hat{g}^2_{1,\tau}(Y_i) - g^2_{1,\tau}(Y_i) \right) \right| + \left| \frac{1}{N} \sum_{i=1}^{N} \frac{1 - e^\star(X_i)}{(1 - \hat{e}(X_i))^2} \cdot \frac{1 - A_i}{1 - e^\star(X_i)} \cdot \left(\hat{g}^2_{0,\tau}(Y_i) - g^2_{0,\tau}(Y_i) \right) \right|.$$

Here the cross-product term will be zero because of $A_1(1 - A_i) = 0$ for all i. Note that the first term of the forgoing sum is the same as equation (A-12) in Firpo (2007), so we can also obtain that $J_{11} = o_p(1)$ by taking the same procedure as Firpo (2007).

Next, we consider the term J_{12}. Define

$$h(X_i) = \mathbb{E}\left[\frac{A}{e(X)^2} \cdot \frac{\psi_\tau(Y - Q_1(\tau))}{f_1(Q_1(\tau))} + \frac{1 - A}{(1 - e(X))^2} \cdot \frac{\psi_\tau(Y - Q_0(\tau))}{f_0(Q_0(\tau))} \mid X = X_i \right].$$

If X is discrete, the estimators of $h(X_i)$ in $\hat{S}_{\tau,i}$ and $\widetilde{S}_{\tau,i}$ are

$$\hat{h}(X_i) = \frac{1}{N} \sum_{j=1}^{N} \widehat{W}_j I(X_j = X_i),$$

$$\tilde{h}(X_i) = \frac{1}{N} \sum_{j=1}^{N} \widetilde{W}_j I(X_j = X_i),$$

respectively, where

$$\widehat{W}_i = \frac{A_i}{\hat{e}(X_i)^2} \cdot \frac{\psi_\tau(Y_i - \hat{Q}_1(\tau))}{\hat{f}_1(\hat{Q}_1(\tau))} + \frac{1 - A_i}{(1 - \hat{e}(X_i))^2} \cdot \frac{\psi_\tau(Y_i - \hat{Q}_0(\tau))}{\hat{f}_0(\hat{Q}_0(\tau))},$$

$$\widetilde{W}_i = \frac{A_i}{\hat{e}(X_i)^2} \cdot \frac{\psi_\tau(Y_i - Q_1(\tau))}{f_1(Q_1(\tau))} + \frac{1 - A_i}{(1 - \hat{e}(X_i))^2} \cdot \frac{\psi_\tau(Y_i - Q_0(\tau))}{f_0(Q_0(\tau))}.$$

Thus, by the law of large numbers, we can easily have $J_{12} = o_p(1)$. If X is continuous, we have the expression of estimator of $h(X_i)$ in $\hat{S}_{\tau,i}, \widetilde{S}_{\tau,i}$, i.e.

$$\hat{h}(X_i) = \frac{\sum_{j=1}^{N} \widehat{W}_j \mathcal{K}_h(X_i - X_j)}{\sum_{j=1}^{N} \mathcal{K}_h(X_i - X_j)},$$

$$\tilde{h}(X_i) = \frac{\sum_{j=1}^{N} \widetilde{W}_j \mathcal{K}_h(X_i - X_j)}{\sum_{j=1}^{N} \mathcal{K}_h(X_i - X_j)},$$

respectively. So

$$\begin{aligned} J_{12} &= \left| \frac{1}{N} \sum_{i=1}^{N} (A_i - \hat{e}(X_i))^2 \left(\hat{h}^2(X_i) - \tilde{h}^2(X_i) \right) \right| \\ &= \left| \frac{1}{N} \sum_{i=1}^{N} (A_i - \hat{e}(X_i))^2 \left(\hat{h}(X_i) + \tilde{h}(X_i) \right) \left(\hat{h}(X_i) - \tilde{h}(X_i) \right) \right| \\ &\le \max_{1 \le i \le N} |A_i - \hat{e}(X_i)|^2 \max_{1 \le i \le N} |\hat{h}(X_i) + \tilde{h}(X_i)| \left| \frac{1}{N} \sum_{i=1}^{N} \left| \hat{h}(X_i) - \tilde{h}(X_i) \right| \right|. \end{aligned}$$

Because $|A_i - \hat{e}(X_i)|$ and $|\hat{h}(X_i) + \tilde{h}(X_i)|$ are all bounded variables, so we only need to show that $\left| \hat{h}(X_i) - \tilde{h}(X_i) \right| = o_p(1)$. Next,

$$\begin{aligned} &\hat{h}(X_i) - \tilde{h}(X_i) \\ &= \frac{\sum_{j=1}^{N} (\widehat{W}_j - \widetilde{W}_j) \mathcal{K}_h(X_i - X_j)}{\sum_{j=1}^{N} \mathcal{K}_h(X_i - X_j)} \\ &= \frac{1}{\sum_{j=1}^{N} \mathcal{K}_h(X_i - X_j)} \sum_{j=1}^{N} \frac{A_i}{\hat{e}(X_j)^2} \mathcal{K}_h(X_i - X_j) \left(\frac{\psi_\tau(Y_j - \hat{Q}_1(\tau))}{\hat{f}_1(\hat{Q}_1(\tau))} - \frac{\psi_\tau(Y_j - Q_1(\tau))}{f_1(Q_1(\tau))} \right) \\ &\quad + \frac{1}{\sum_{j=1}^{N} \mathcal{K}_h(X_i - X_j)} \sum_{j=1}^{N} \frac{1 - A_i}{(1 - \hat{e}(X_j))^2} \mathcal{K}_h(X_i - X_j) \left(\frac{\psi_\tau(Y_j - \hat{Q}_0(\tau))}{\hat{f}_0(\hat{Q}_0(\tau))} - \frac{\psi_\tau(Y_j - Q_0(\tau))}{f_0(Q_0(\tau))} \right). \end{aligned}$$

Then, by the consistency of $\hat{Q}_{a,\tau}$ and $\hat{f}_a(\hat{Q}_{a,\tau})$, we have $\hat{h}(X_i) - \tilde{h}(X_i) = o_p(1)$. Here we complete the proof of $J_{12} = o_p(1)$ for continuous case.

For the last term J_{13}, we have

$$\begin{aligned} J_{13} &= \left| \frac{2}{N} \sum_{i=1}^{N} \left\{ \widehat{U}_{\tau,i} \widehat{S}_{\tau,i} - \widetilde{U}_{\tau,i} \widetilde{S}_{\tau,i} \right\} \right| \\ &\leq 2 \left| \frac{1}{N} \sum_{i=1}^{N} \widetilde{U}_{\tau,i} (\widehat{S}_{\tau,i} - \widetilde{S}_{\tau,i}) \right| + 2 \left| \frac{1}{N} \sum_{i=1}^{N} \widehat{S}_{\tau,i} (\widehat{U}_{\tau,i} - \widetilde{U}_{\tau,i}) \right| \\ &\doteq J_{131} + J_{132}. \end{aligned}$$

By taking the same procedure as J_{12}, we also have $J_{131} = o_p(1)$. Similarly, $J_{132} = o_p(1)$ by similar proof of J_{11}. Here we complete the proof. □

Proof of Theorem 2. First, we show the consistency of $\widehat{Q}_{a,1}(\tau), \widehat{f}_{a,1}(\tau)$ and $\widehat{\sigma}^2_{a,1}(z)$. Obviously, $\widehat{Q}_{a,1}(\tau)$ is the minimizer of

$$W(u) = \frac{1}{n_1} \sum_{i=1}^{n_1} \rho_\tau(Y_i - u) I(A_i = a).$$

Because we use the random design in the first step, so A_i is independent of Y_i. Thus, by the law of large numbers, we have $W(u) \to E(\rho_\tau(Y_i) - u), a.s.$ Consequently, we have $\widehat{Q}_{a,1}(\tau) \to Q_a(\tau), a.s.$ and $\widehat{Q}_{a,1}(\tau) - Q_a(\tau) = O\left(\sqrt{\frac{\log N}{N}}\right), a.s.$ as $n_1/N \to c \in (0,1)$ (Bssett and Koenker, 1986). So,

$$\widehat{f}_{a,1}(\tau) = \frac{1}{n_1} \sum_{i=1}^{n_1} \mathcal{K}_h \left(\frac{Y_i - Q_a(\tau)}{h} \right) \mathbf{1}(A_i = a) + O\left(\sqrt{\frac{\log N}{N}} \right), a.s.$$

Similarly, according to Theorem B in Silverman (1978), we also have

$$\frac{1}{n_1} \sum_{i=1}^{n_1} \mathcal{K}_h \left(\frac{Y_i - Q_a(\tau)}{h} \right) \mathbf{1}(A_i = a) = f_a(\tau) + O(h^2) + O\left(\sqrt{\frac{\log(1/h)}{n_1 h}} \right), a.s.$$

by the independence between Y_i and A_i. Thus, by the condition $h \propto N^{-1/5}$ and $n_1/N \to c \in (0,1)$, we have $\widehat{f}_{a,1}(\tau) = f_a(\tau) + O\left(\frac{\log^{1/2} N}{N^{2/5}}\right), a.s.$ Similarly, we also have

$$\widehat{\sigma}^2_{a,1}(z) = \sigma^2_a(x) + O\left(\frac{\log^{1/2} N}{N^{1/2}} \right), a.s.$$

Next, we will prove the consistency of $\widehat{Q}_{a,1}(\tau), \widehat{f}_{a,1}(\tau)$ and $\widehat{\sigma}^2_{a,1}(z)$ by mathematical induction. Suppose

$$\widehat{Q}_{a,t-1}(\tau) = Q_a(\tau) + O\left(\frac{\log^{1/2} N}{N^{1/2}} \right), a.s. \tag{A.1}$$

$$\widehat{f}_{a,t-1}(\tau) = f_a(\tau) + O\left(\frac{\log^{1/2} N}{N^{2/5}} \right), a.s. \tag{A.2}$$

$$\widehat{\sigma}^2_{a,t-1}(z) = \sigma^2_a(x) + O\left(\frac{\log^{1/2} N}{N^{1/2}} \right), a.s. \tag{A.3}$$

By (A.1), we know that

$$\widehat{f}_{a,t}(\tau) = \frac{1}{N_{a,t}} \sum_{i=1}^{N_t} \mathcal{K}_h \left(\frac{Y_i - Q_a(\tau)}{h} \right) \mathbf{1}(A_i = a) + O\left(\frac{\log^{1/2} N}{N^{1/2}} \right), a.s.$$

Let $\{W_i\}$ be an independent copy of $\{Y_i\}$, which is also independent of Z_i. Define $\nu_i^a = \min\{j \mid N_{a,j} = i\}$, $\widetilde{Y}_{i,a} = Y_{\nu_i^a} \mathbf{1}(\nu_i^a < +\infty) + W_i \mathbf{1}(\nu_i^a = +\infty)$, $i = 1, 2, \ldots$. Using the same argument of Doob (1936), we can show that $\{\widetilde{Y}_{i,a}, i = 1, 2, \ldots\}$ is a sequence of i.i.d random variable with a common distribution the same as that of Y_i. Thus, we have

$$\frac{1}{N_{a,t}} \sum_{i=1}^{N_{a,t}} \mathcal{K}_h \left(\frac{\widetilde{Y}_{i,a} - Q_a(\tau)}{h} \right) = f_a(\tau) + O\left(\frac{\log^{1/2} N}{N^{2/5}} \right), a.s.$$

if $N_{a,t} \to \infty$. Thus, on the event $\{\lim_{t\to\infty} N_{a,t} = \infty\}$, we have

$$\widehat{f}_{a,t}(\tau) = f_a(\tau) + O(h^2) + O\left(\frac{\log^{1/2} N}{N^{2/5}} \right), a.s.$$

Similarly, we have

$$\widehat{\sigma}^2_{a,t}(z) = \sigma^2_a(z) + O\left(\frac{\log^{1/2} N}{N^{1/2}} \right), a.s.$$

on the event $\{\lim_{t\to\infty} N_{a,t} = \infty\}$. As a result, by the definition of $\hat{e}_t(z)$, we have

$$\hat{e}_t(z) = e^\star(z) + O(h^2) + O\left(\frac{\log^{1/2} N}{N^{2/5}}\right), a.s.$$

as $N_{a,t} \to \infty$. Because $\hat{e}_t(z)$ is a consistent estimator of $e^\star(z)$ as $t \to \infty$. So taking the same procedure as the proof of Theorem 1 in Firpo (2007), we can show that $\hat{Q}_{a,t}(\tau) = \check{Q}_{a,t}(\tau) + o(\log^{1/2} N / N^{1/2}), a.s.$, where

$$\check{Q}_{a,t}(\tau) = \arg\min_u \frac{1}{N_t} \sum_{s=1}^{t} \sum_{i=N_{s-1}+1}^{N_s} \frac{\mathbf{1}(A_i = a)}{A_i e_s^\star(Z_i) + (1 - A_i)\left(1 - e_s^\star(Z_i)\right)} \rho_\tau(Y_i - u).$$

Therefore,

$$\check{Q}_{1,t}(\tau) = \arg\min_u \frac{1}{N_{1,t}} \sum_{s=1}^{t} \sum_{i=N_{s-1}+1}^{N_s} \frac{\mathbf{1}(A_i = 1)}{e_s^\star(Z_i)} \rho_\tau(\widetilde{Y}_i - u), \tag{A.4}$$

$$\check{Q}_{0,t}(\tau) = \arg\min_u \frac{1}{N_{0,t}} \sum_{s=1}^{t} \sum_{i=N_{s-1}+1}^{N_s} \frac{\mathbf{1}(A_i = 0)}{1 - e_s^\star(Z_i)} \rho_\tau(\widetilde{Y}_i - u). \tag{A.5}$$

on the event $\{N_{a,t} \to \infty\}$, respectively. By the strong consistency of the quantile regression, we also have

$$\check{Q}_{1,t}(\tau) = Q_{1,t} + O\left(\frac{\log^{1/2} N}{N^{1/2}}\right), a.s., \quad \check{Q}_{0,t}(\tau) = Q_{0,t} + O\left(\frac{\log^{1/2} N}{N^{1/2}}\right), a.s.$$

if $N_{a,t} \to \infty, a = 0, 1$.

So we need to show that $N_{a,t} \to \infty, a.s.$ for $a = 0, 1$. On the event $\{\lim_{t\to\infty} N_{a,t} < \infty\}$, the sequence $\hat{Q}_{a,t}(\tau), \hat{f}_{a,t}(\tau), \hat{\sigma}^2_{a,t}(z)$ takes a finite number of values. So $\hat{e}_t(z)$ takes a finite number of values. Also, on the event $\{\lim_{t\to\infty} N_{a,t} = \infty\}$, $\hat{e}_t(z) \to e^\star(z)$. This shows that $\{\hat{e}_t(z)\}$ is almost surely a relatively compact set. So there exist a $0 < \delta < 1$ such that $\hat{e}_t(z) \geq \delta$. Define $N_{a,t,z} = \sum_{i=1}^{t} \mathbf{1}(A_i = a, Z_i = z)$. So $N_{a,t} = \sum_{z\in\mathcal{Z}} N_{a,t,z}$. Note that

$$\mathbb{P}\left(A_t = 1 \mid \mathcal{F}_{t-1}, Z_t = z\right) = \hat{e}_t(z) \geq \delta.$$

Then,

$$\sum_{t=2}^{\infty} \mathbb{P}\left(A_t = 1 \mid \mathcal{F}_{t-1}, Z_t = z\right) = +\infty, a.s.$$

which implies $\{A_t = 1, i.o.\} = \{N_{a,t,z} \to \infty\}$ almost surely by the generalized Borel–Cantelli lemma. Consequently, $N_{a,t} \to \infty$ almost surely.

Next, we show the asymptotic normality of $\hat{Q}_{a,T}(\tau)$. By (A.4) and (A.5), we can obtain that

$$\sqrt{N_T}(\check{Q}_{1,T}(\tau) - Q_1(\tau)) \xrightarrow{d} \mathcal{N}(0, V_1),$$
$$\sqrt{N_T}(\check{Q}_{0,T}(\tau) - Q_0(\tau)) \xrightarrow{d} \mathcal{N}(0, V_0),$$

where

$$V_1 = \mathbb{E}\left(\frac{\sigma_1^2(Z)}{f_1^2(Z)e^\star(Z)} + \left[\frac{\mathbb{E}\left[\psi_\tau\left(Y - Q_1(\tau)\right)|Z, A = 1\right]}{f_1(\tau)}\right]^2\right),$$
$$V_0 = \mathbb{E}\left(\frac{\sigma_0^2(Z)}{f_0^2(Z)(1 - e^\star(Z))} + \left[\frac{\mathbb{E}\left[\psi_\tau\left(Y - Q_0(\tau)\right)|Z, A = 0\right]}{f_0(\tau)}\right]^2\right).$$

Similar to Lemma A.5 in Firpo (2007), the difference between $\check{Q}_{1,T} - \check{Q}_{0,T}$ still have an asymptotically linear influence function and will be asymptotically normal, i.e.

$$\sqrt{N_T}(\check{Q}_{1,T} - \check{Q}_{0,T}) \xrightarrow{d} \mathcal{N}(0, \mathbb{V}_{\tau,FA})$$

where

$$\mathbb{V}_{\tau,FA} = \mathbb{E}\left[\left(\frac{\sigma_1(Z)}{f_1(\tau)} + \frac{\sigma_0(Z)}{f_0(\tau)}\right)^2 + \left[\frac{\mathbb{E}\left[\psi_\tau\left(Y - Q_1(\tau)\right)|Z, A = 1\right]}{f_1(\tau)} - \frac{\mathbb{E}\left[\psi_\tau\left(Y - Q_0(\tau)\right)|Z, A = 0\right]}{f_0(\tau)}\right]^2\right].$$

Next, we prove the consistency of the estimator $\widehat{\mathbb{V}}_{\tau,FA}$. Because of the consistency of $\hat{Q}_a(\tau), \hat{e}(z), \hat{f}_a$, by taking the same procedure as the proof of Theorem 1, we can show that

$$\widehat{\mathbb{V}}_{\tau,FA} = \frac{1}{N_T} \sum_{t=1}^{T} \sum_{i=N_{t-1}+1}^{N_t} (U_{\tau,it} + S_{\tau,it})^2 + o_p(1),$$

where

$$U_{\tau,it} = \frac{A_i}{e_t^\star(Z_i)} \cdot \frac{\psi_\tau(Y_i - Q_1(\tau))}{f_{1,t}(Q_1(\tau))} - \frac{1-A_i}{1-e_t^\star(Z_i)} \frac{\psi_\tau(Y_i - Q_0(\tau))}{f_{0,t}(Q_0(\tau))},$$

$$S_{\tau,it} = -\mathbb{E}\left[\frac{A}{(e_t^\star(Z))^2} \cdot \frac{\psi_\tau(Y - Q_1(\tau))}{f_{1,t}(Q_1(\tau))} + \frac{1-A}{(1-e_t^\star(Z))^2} \cdot \frac{\psi_\tau(Y - Q_0(\tau))}{f_{0,t}(Q_0(\tau))} \mid Z = Z_i \right] \cdot (A_i - e_t^\star(Z_i)).$$

On the event $\{N_{a,t} \to \infty\}$, we have

$$\widehat{\mathbb{V}}_{\tau,FA} = \frac{1}{N_T} \sum_{t=1}^{T} \sum_{i=N_{t-1}+1}^{N_t} (\widetilde{U}_{\tau,it} + S_{\tau,it})^2 + o_p(1),$$

where

$$\widetilde{U}_{\tau,it} = \frac{A_i}{e_t^\star(Z_i)} \cdot \frac{\psi_\tau(\widetilde{Y}_i - Q_1(\tau))}{f_{1,t}(Q_1(\tau))} - \frac{1-A_i}{1-e_t^\star(Z_i)} \frac{\psi_\tau(\widetilde{Y}_i - Q_0(\tau))}{f_{0,t}(Q_0(\tau))}.$$

By the law of large numbers, we obtain $\widehat{\mathbb{V}}_{\tau,FA} = \mathbb{V}_{\tau,FA} + o_p(1)$. □

Proof for Theorem 3. Recall that X_i are the original covariates. The discretized covariates Z_i is coarsened X_i, and it contains less information. By design, the conditional independence holds for both X_i and Z_i, that is,

$$(Y_i(1), Y_i(0)) \perp\!\!\!\perp A_i \mid X_i, \quad \text{and} \quad (Y_i(1), Y_i(0)) \perp\!\!\!\perp A_i \mid Z_i.$$

Our goal is to show that when the treatment is assigned using the discretized covariate Z_i, meaning that $\mathbb{P}[A_i = 1|X_i] = \mathbb{P}[A_i = 1|Z_i] = e(Z_i)$, the following inequality holds:

$$\mathbb{V}[\psi_{(\mathrm{ii})}(Y_i, X_i, X_i)] \geq \mathbb{V}[\psi_{(\mathrm{i})}(Y_i, X_i, X_i)] \geq \mathbb{V}_{\tau,\mathrm{TS}},$$

where

$$\psi_{(\mathrm{i})}(Y_i, A_i, X_i) = \frac{A_i(q_1(\tau) - g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau) - g_0(X_i))}{1-e(Z_i)}$$

$$\psi_{(\mathrm{ii})}(Y_i, A_i, X_i) = \frac{A_i(q_1(\tau) - g_1(Z_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau) - g_0(Z_i))}{1-e(Z_i)},$$

with $q_a(\tau) = \tau / f_a(\tau)$ and

$$g_1(X_i) = \mathbb{E}\left[\frac{\mathbf{1}_{Y(1) \leq Q_1(\tau)}}{f_1(\tau)} \middle| X_i \right], \quad g_0(X_i) = \mathbb{E}\left[\frac{\mathbf{1}_{Y(0) \leq Q_0(\tau)}}{f_0(\tau)} \middle| X_i \right],$$

$$g_1(Z_i) = \mathbb{E}\left[\frac{\mathbf{1}_{Y(1) \leq Q_1(\tau)}}{f_1(\tau)} \middle| Z_i \right], \quad g_0(Z_i) = \mathbb{E}\left[\frac{\mathbf{1}_{Y(0) \leq Q_0(\tau)}}{f_0(\tau)} \middle| Z_i \right].$$

As the above inequality holds for all possible propensity scores, we establish $\mathbb{V}_{\tau,\mathrm{TS}} \leq \mathbb{V}_{\tau,\mathrm{FA}}$.

The second inequality holds by the definition of the optimal variance lower bound defined in Section 2. To prove the first inequality, we need to compare two efficiency bounds. We notice that

$$\begin{aligned}\mathbb{V}[\psi_{(\mathrm{ii})}(Y_i, X_i, X_i)] &= \mathbb{V}\Big[\mathbb{E}[\psi_{(\mathrm{ii})}(Y_i, X_i, X_i)|Z_i]\Big] + \mathbb{E}\Big[\mathbb{V}[\psi_{(\mathrm{ii})}(Y_i, X_i, X_i)|Z_i]\Big] \\ &= \mathbb{V}[g_1(Z_i) - g_0(Z_i)] + \mathbb{E}\left[\mathbb{V}\left[\frac{A_i(q_1(\tau) - g_1(Z_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau) - g_0(Z_i))}{1-e(Z_i)} \middle| Z_i \right]\right].\end{aligned}$$

We may adopt a similar decomposition. To be precise, we write

$$\begin{aligned}\psi_{(\mathrm{i})}(Y_i, X_i, X_i) = \frac{A_i(q_1(\tau) - g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau) - g_0(X_i))}{1-e(Z_i)} + (g_1(X_i) - g_0(X_i) - g_1(Z_i) + g_0(Z_i)) \\ + g_1(Z_i) - g_0(Z_i) - (q_1(\tau) - q_0(\tau)),\end{aligned}$$

which gives

$$\begin{aligned}\mathbb{V}[\psi_{(\mathrm{ii})}(Y_i, X_i, X_i)] &= \mathbb{V}[g_1(Z_i) - g_0(Z_i)] \\ &+ \mathbb{E}\left[\mathbb{V}\left[\frac{A_i(q_1(\tau) - g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(Q_0(\tau) - g_0(X_i))}{1-e(Z_i)} + (g_1(X_i) - g_0(X_i) - g_1(Z_i) + g_0(Z_i)) \middle| Z_i \right]\right].\end{aligned}$$

We thus compare the two conditional variances in the expressions of $\psi_{(\mathrm{i})}(Y_i, X_i, X_i)$ and $\psi_{(\mathrm{ii})}(Y_i, X_i, X_i)$. Notice that

$$\begin{aligned}&\mathbb{V}\left[\frac{A_i(q_1(\tau) - g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau) - g_0(X_i))}{1-e(Z_i)} + (g_1(X_i) - g_0(X_i) - g_1(Z_i) + g_0(Z_i)) \middle| Z_i \right] \\ &= \mathbb{V}\left[\mathbb{E}\left[\frac{A_i(q_1(\tau) - g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau) - g_0(X_i))}{1-e(Z_i)} + (g_1(X_i) - g_0(X_i) - g_1(Z_i) + g_0(Z_i)) \middle| X_i \right] \middle| Z_i \right]\end{aligned}$$

$$
\begin{aligned}
&+ \mathbb{E}\left[\mathbb{V}\left[\frac{A_i(q_1(\tau)-g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(X_i))}{1-e(Z_i)} + (g_1(X_i)-g_0(X_i)-g_1(Z_i)+g_0(Z_i))\middle| X_i\right]\middle| Z_i\right] \\
&= \mathbb{V}\left[g_1(X_i)-g_0(X_i)-g_1(Z_i)+g_0(Z_i)\middle| Z_i\right] \\
&\quad + \mathbb{E}\left[\mathbb{V}\left[\frac{A_i(q_1(\tau)-g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(X_i))}{1-e(Z_i)}\middle| X_i\right]\middle| Z_i\right].
\end{aligned}
$$

Similarly, the conditional variance in the expression of $\psi_{(\mathrm{ii})}(Y_i, X_i, X_i)$ is

$$
\begin{aligned}
&\mathbb{V}\left[\frac{A_i(Y_i-g_1(Z_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(Z_i))}{1-e(Z_i)}\middle| Z_i\right] \\
&= \mathbb{V}\left[\frac{A_i(q_1(\tau)-g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(X_i))}{1-e(Z_i)} + \frac{A_i(g_1(X_i)-g_1(Z_i))}{e(Z_i)} - \frac{(1-A_i)(g_0(X_i)-g_0(Z_i))}{1-e(Z_i)}\middle| Z_i\right] \\
&= \mathbb{V}\left[\mathbb{E}\left[\frac{A_i(q_1(\tau)-g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(X_i))}{1-e(Z_i)} + \frac{A_i(g_1(X_i)-g_1(Z_i))}{e(Z_i)} - \frac{(1-X_i)(g_0(X_i)-g_0(Z_i))}{1-e(Z_i)}\middle| X_i\right]\middle| Z_i\right] \\
&\quad + \mathbb{E}\left[\mathbb{V}\left[\frac{A_i(q_1(\tau)-g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(X_i))}{1-e(Z_i)}\middle| X_i\right]\middle| Z_i\right] \\
&\quad + \mathbb{E}\left[\mathbb{V}\left[\frac{A_i(g_1(X_i)-g_1(Z_i))}{e(Z_i)} - \frac{(1-A_i)(g_0(X_i)-g_0(Z_i))}{1-e(Z_i)}\middle| X_i\right]\middle| Z_i\right] \\
&= \mathbb{V}\left[g_1(X_i)-g_0(X_i)-g_1(Z_i)+g_0(Z_i)\middle| Z_i\right] \\
&\quad + \mathbb{E}\left[\mathbb{V}\left[\frac{A_i(q_1(\tau)-g_1(Z_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(Z_i))}{1-e(Z_i)}\middle| X_i\right]\middle| Z_i\right].
\end{aligned}
$$

To summarize,

$$
\begin{aligned}
&\mathbb{V}[\psi_{(\mathrm{ii})}(Y_i, X_i, X_i)] - \mathbb{V}[\psi_{(\mathrm{i})}(Y_i, X_i, X_i)] \\
&= \mathbb{E}\left[\mathbb{V}\left[\frac{A_i(q_1(\tau)-g_1(Z_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(Z_i))}{1-e(Z_i)}\middle| X_i\right]\right. \\
&\quad \left. - \mathbb{V}\left[\frac{A_i(q_1(\tau)-g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(X_i))}{1-e(Z_i)}\middle| X_i\right]\right].
\end{aligned}
$$

As a last step, we compute the above conditional variances explicitly:

$$
\mathbb{V}\left[\frac{A_i(q_1(\tau)-g_1(X_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(X_i))}{1-e(Z_i)}\middle| X_i\right] = \frac{\sigma_1(X_i)^2}{f_1^2(\tau)e(X_i)} + \frac{\sigma_0(X_i)^2}{f_0^2(\tau)(1-e(X_i))},
$$

where $\sigma_1(X_i)^2$ and $\sigma_0(X_i)^2$ are the conditional variances of the quantile rank scores of the potential outcomes, respectively. Similarly (variance equals the difference of the second moment and the squared first moment),

$$
\begin{aligned}
&\mathbb{V}\left[\frac{A_i(q_1(\tau)-g_1(Z_i))}{e(Z_i)} - \frac{(1-A_i)(q_0(\tau)-g_0(Z_i))}{1-e(Z_i)}\middle| X_i\right] \\
&= \frac{\mathbb{E}[(q_1(\tau)-g_1(Z_i))^2|X_i]}{e(Z_i)} + \frac{\mathbb{E}[(q_0(\tau)-g_0(Z_i))^2|X_i]}{1-e(Z_i)} - \Big(g_1(X_i)-g_0(X_i)-g_1(Z_i)+g_0(Z_i)\Big)^2 \\
&= \frac{\sigma_1(X_i)^2}{f_1^2(\tau)e(X_i)} + \frac{\sigma_0(X_i)^2}{f_0^2(\tau)(1-e(X_i))} + \frac{(g_1(X_i)-g_1(Z_i))^2}{e(Z_i)} + \frac{(g_0(X_i)-g_0(Z_i))^2}{1-e(Z_i)} \\
&\quad - \Big(g_1(X_i)-g_0(X_i)-g_1(Z_i)+g_0(Z_i)\Big)^2.
\end{aligned}
$$

We therefore conclude that

$$
\begin{aligned}
&\mathbb{V}[\psi_{(\mathrm{ii})}(Y_i, X_i, X_i)] - \mathbb{V}[\psi_{(\mathrm{i})}(Y_i, X_i, X_i)] \\
&= \mathbb{E}\left[\frac{(g_1(X_i)-g_1(Z_i))^2}{e(Z_i)} + \frac{(g_0(X_i)-g_0(Z_i))^2}{1-e(Z_i)} - \Big(g_1(X_i)-g_0(X_i)-g_1(Z_i)+g_0(Z_i)\Big)^2\right] \geq 0,
\end{aligned}
$$

with the inequality following from

$$
\begin{aligned}
&\frac{(g_1(X_i)-g_1(Z_i))^2}{e(Z_i)} + \frac{(g_0(X_i)-g_0(Z_i))^2}{1-e(Z_i)} - \Big(g_1(X_i)-g_0(X_i)-g_1(Z_i)+g_0(Z_i)\Big)^2 \\
&= \left(\sqrt{\frac{1}{e(Z_i)}-1}\Big(g_1(X_i)-g_1(Z_i)\Big) + \sqrt{\frac{1}{1-e(Z_i)}-1}\Big(g_0(X_i)-g_0(Z_i)\Big)\right)^2 \geq 0.
\end{aligned}
$$

Table B.1
Performances of different estimators in the case where X is discrete. The total sample size is $N = 500$. For the two-stage design, we set $n_1 = 0.5N$ and the fully adaptive design, we set $n_1 = 0.2N$.

τ	Design	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
0.3	Completely random	188.51	−7.13	35.58	–
	Naive	173.9	1.41	30.24	–
	Two stage	173.71	0.74	30.17	0.97
	Fully adaptive	178.2	−0.6	31.75	0.97
0.5	Completely random	231.39	3.09	53.54	–
	Naive	205.51	9.65	42.32	–
	Two stage	206.33	4.72	42.59	0.96
	Fully adaptive	209.48	4.02	43.89	0.96
0.8	Completely random	368.51	12.25	135.93	–
	Naive	317.58	10.25	100.95	–
	Two stage	314.39	0.69	98.83	0.94
	Fully adaptive	332.46	0.29	110.52	0.94

Table B.2
Performance of each estimator with continuous case with $N = 500$.

τ	Design	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
0.3	Completely random	305.34	−9.95	93.32	–
	Naive	316.2	−48.04	102.28	–
	Two stage	288.54	9.01	83.33	0.97
	Fully adaptive	321.13	−2.87	103.12	0.94
0.5	Completely random	325.54	−2.04	105.97	–
	Naive	328.46	−47.23	110.11	–
	Two stage	306.59	−12.8	94.15	0.96
	Fully adaptive	311.43	0.96	96.98	0.96
0.8	Completely random	422.66	6.49	178.67	–
	Naive	405.43	−43.93	166.29	–
	Two stage	391.56	−8.99	153.38	0.93
	Fully adaptive	398.38	−5.65	158.73	0.95

Table B.3
Performances of different estimators in the case where X is discrete. The total sample size is $N = 200$. For the two-stage design, we set $n_1 = 0.5N$ and the fully adaptive design, we set $n_1 = 0.2N$.

τ	Design	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
0.3	Completely random	302.592	−20.209	91.961	–
	Naive	278.438	−1.89	77.523	–
	Two stage	279.869	−2.793	78.327	0.99
	Fully adaptive	336.108	2.145	112.962	0.96
0.5	Completely random	364.521	−10.361	132.97	–
	Naive	326.699	0.804	106.722	–
	Two stage	326.461	−9.064	106.649	0.97
	Fully adaptive	339.486	−4.844	115.263	0.97
0.8	Completely random	578.978	32.002	336.207	–
	Naive	509.37	19.505	259.813	–
	Two stage	505.242	4.23	255.262	0.91
	Fully adaptive	544.458	3.714	296.419	0.93

Appendix B. Additional simulations

In order to investigate the effectiveness of the proposed design at different sample sizes, we have considered $N = 200, 500$ in our study. The simulation results for the four methods are reported in Tables B.1–B.4. Our findings suggest that our proposed designs continue to perform exceptionally well at $N = 500$. However, at smaller potential sample sizes, specifically $N = 200$, our proposed designs do not outperform the naive design. This is to be expected as smaller sample sizes may not provide a consistent estimator of the variance and density function.

In order to demonstrate the sensitivity of our proposed designs to the initial stage sample sizes, n_1, we examined $n_1 = 0.3n, 0.4n$ for the two-stage design and $n_1 = 0.1n, 0.16n$ for the fully adaptive design. The performance of the corresponding estimators is reported in Tables B.5–B.6. Our findings indicate that the two estimators yield results comparable to those obtained with $n_1 = 0.5n$ for the two-stage estimator and $n_1 = 0.2n$ for the fully adaptive estimator. Thus, our methods exhibit low sensitivity to the initial stage sample sizes.

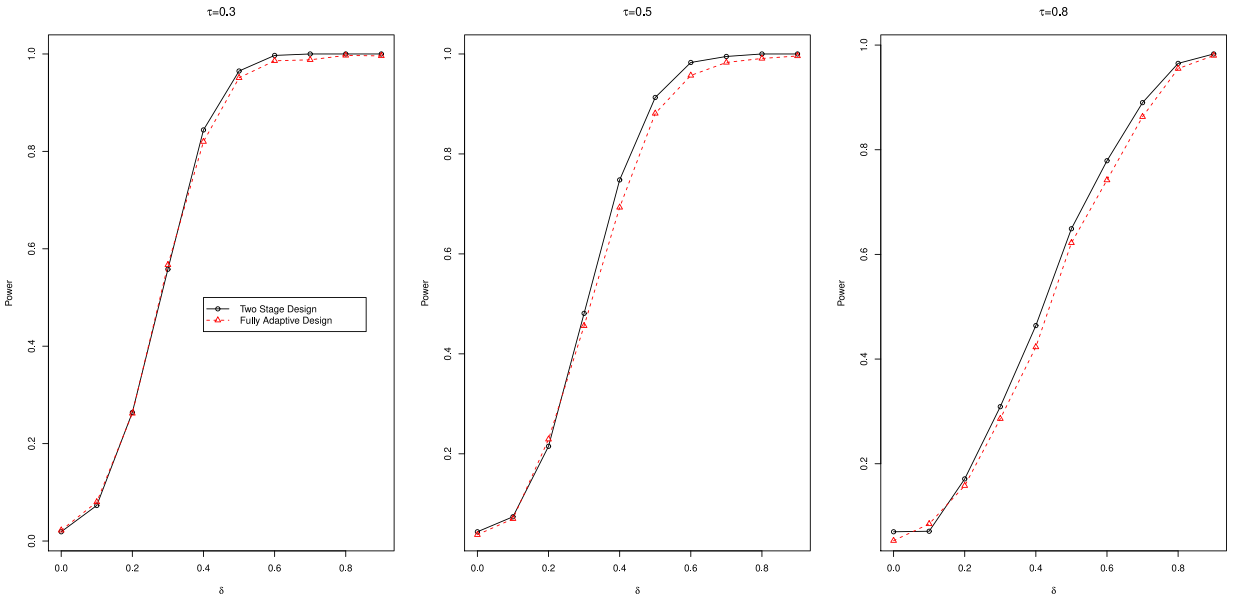

Fig. B.1. Empirical power curves of different design under discrete case with $N = 1000$.

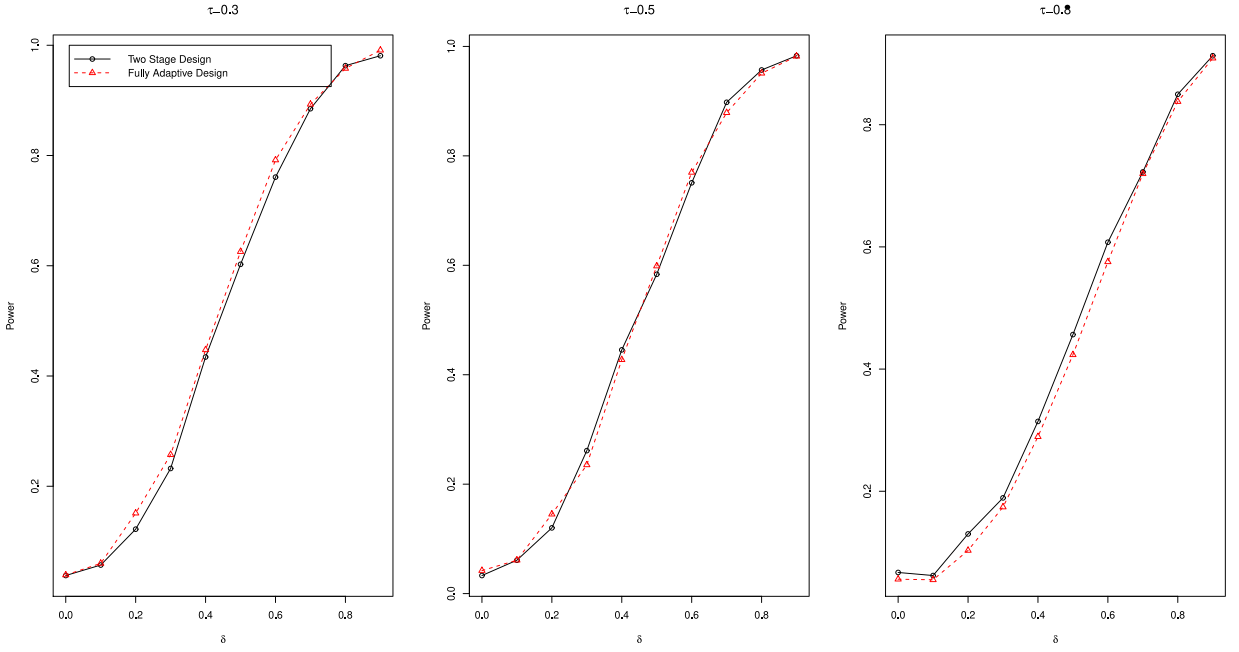

Fig. B.2. Empirical power curves of different design under continuous case with $N = 1000$.

Table B.4
Performance of each estimator with continuous case with $N = 200$.

τ	Design	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
0.3	Completely random	480.89	−28.47	232.04	–
	Naive	505.21	−65.51	259.5	–
	Two stage	462.98	3.57	214.34	0.96
	Fully adaptive	661.9	−24.54	438.66	0.88
0.5	Completely random	517.1	−4.59	267.39	–
	Naive	524.43	−57.19	278.27	–
	Two stage	484.6	−27.75	235.58	0.96
	Fully adaptive	538.79	−4.27	290.28	0.95
0.8	Completely random	676.79	9.31	458.09	–
	Naive	650.31	−75.27	428.53	–
	Two stage	633.06	−50.23	403.25	0.91
	Fully adaptive	666.75	−27.62	445.27	0.95

Table B.5
Performance of two stage estimator with different $n_1 = cn$ for with $N = 1000$.

τ	c	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
Discrete case					
0.3	0.3	123.41	−8.35	15.3	0.97
	0.4	122.46	0.09	15	0.97
0.5	0.3	144.95	−16.69	21.29	0.96
	0.4	145.84	5.82	21.3	0.96
0.8	0.3	223.37	−16.12	50.15	0.93
	0.4	223.26	3.33	49.85	0.93
Continuous case					
0.3	0.3	203.83	−3.59	41.55	0.96
	0.4	204.9	1.01	41.98	0.96
0.5	0.3	215	−17.94	46.54	0.96
	0.4	213.8	−7.15	45.76	0.96
0.8	0.3	278.07	−15.02	77.54	0.94
	0.4	276.88	−2.89	76.66	0.94

Table B.6
Performance of fully adaptive estimator with different $n_1 = cn$ for with $N = 1000$.

τ	c	Standard error ($\times 10^{-3}$)	Bias ($\times 10^{-3}$)	MSE ($\times 10^{-3}$)	Coverage probability
Discrete case					
0.3	0.1	125.91	−9.76	15.95	0.97
	0.16	125.13	−2.29	15.66	0.97
0.5	0.1	146.69	−15.69	21.76	0.96
	0.16	147.2	5.68	21.7	0.96
0.8	0.1	230.14	−11.87	53.1	0.94
	0.16	232.44	6.87	54.07	0.94
Continuous case					
0.3	0.1	219.51	−9	48.26	0.95
	0.16	212.86	−4.02	45.32	0.95
0.5	0.1	217.09	−9.63	47.21	0.96
	0.16	217.77	−1.95	47.42	0.96
0.8	0.1	277.05	−6.25	76.79	0.95
	0.16	281.96	7.57	79.55	0.95

We consider the performance of power of each estimators. For the discrete case, we set $Y(0) = X - 6 + \epsilon_0 + \delta$ and the other settings are the same as above. So if $\delta = 0$, we can calculate the QTEs $\alpha_\tau^{(0)}$ with $\tau = 0.3, 0.5, 0.8$ are $2.21, 1.40, -0.60$, respectively. So we consider the null hypothesis

$$H_0 : \alpha_\tau = \alpha_\tau^{(0)}, \quad vs. \quad H_1 : \alpha_\tau \neq \alpha_\tau^{(0)} \tag{B.1}$$

for each $\tau = 0.3, 0.5, 0.8$. Under the alternative, we set $\delta = 0.1, 0.2, \ldots, 0.9$ and Fig. B.1 shows the power curves of different design under discrete case with $N = 1000$. Similarly, for the continuous case, we also set $Y(0) = X + e^{2X} - 9 + \delta + \epsilon_0$ with $\delta = 0, 0.1, \ldots, 0.9$ and the other settings are the same as above. Fig. B.2 shows the power curves of different design under continuous case with $N = 1000$.

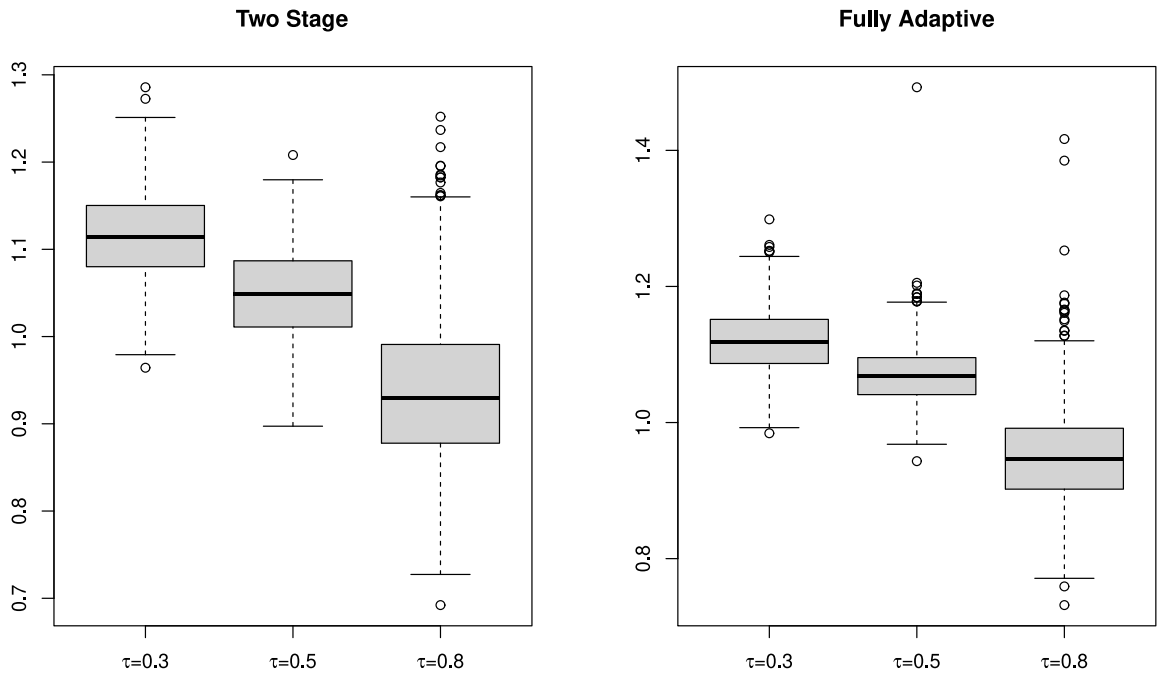

Fig. B.3. Boxplot of $\widehat{\mathbb{V}}_\tau^{1/2}/\mathbb{V}_\tau^{1/2}$ of different design under discrete case with $N = 1000$.

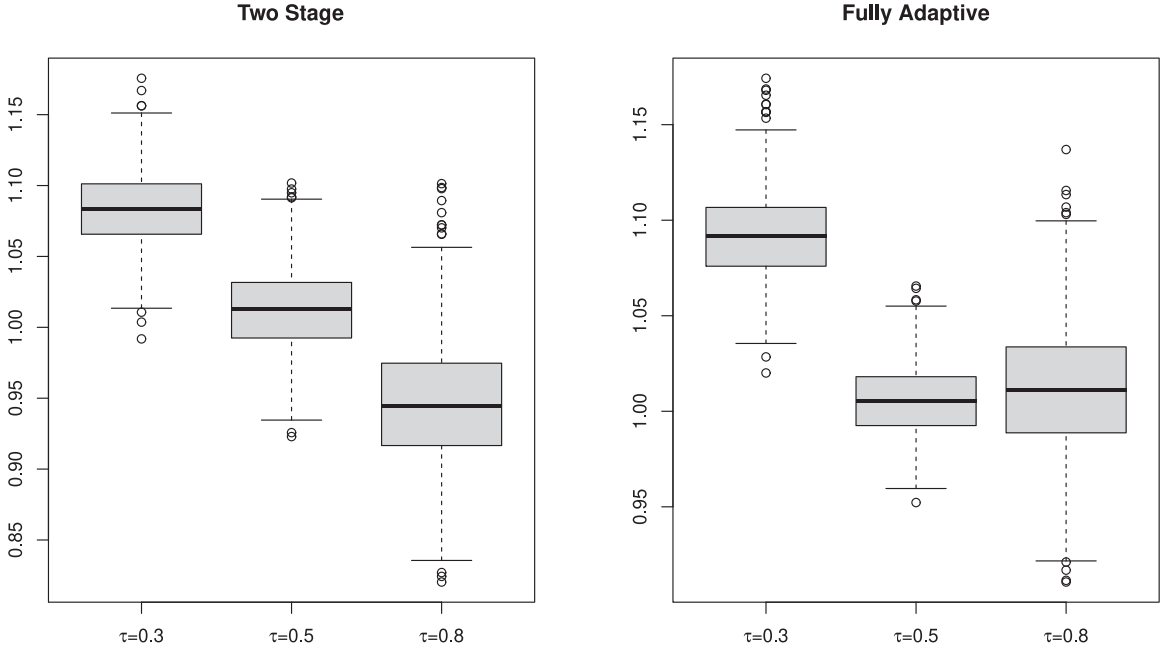

Fig. B.4. Boxplot of $\widehat{\mathbb{V}}_\tau^{1/2}/\mathbb{V}_\tau^{1/2}$ of different design under discrete case with $N = 5000$.

Our findings indicate that our testing procedure is highly effective in managing empirical sizes across all scenarios. Additionally, the power of our proposed testing method increases proportionally with the signal strength. Interestingly, the performance of both the two-stage design and the fully adaptive design are strikingly similar.

In addition, we present the ratio of the estimated standard deviation to the actual standard deviation, denoted as $\widehat{\mathbb{V}}_\tau^{1/2}/\mathbb{V}_\tau^{1/2}$, for both the two-stage design and the fully adaptive design in Figs. B.3–B.6. Our findings indicate that the majority of these ratios are close to one, suggesting that our variance estimator is consistent. However, in some instances, the estimated variance is slightly higher than the actual variance, leading to a coverage probability that is marginally above the targeted 95%.

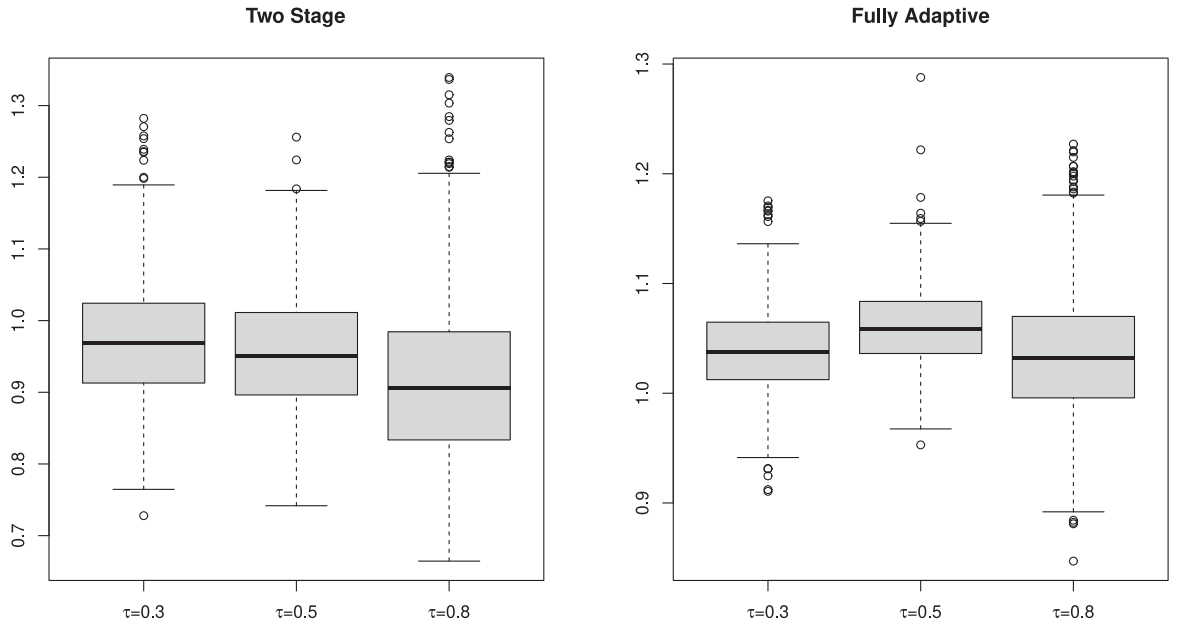

Fig. B.5. Boxplot of $\widehat{\mathbb{V}}_\tau^{1/2}/\mathbb{V}_\tau^{1/2}$ of different design under continuous case with $N = 1000$.

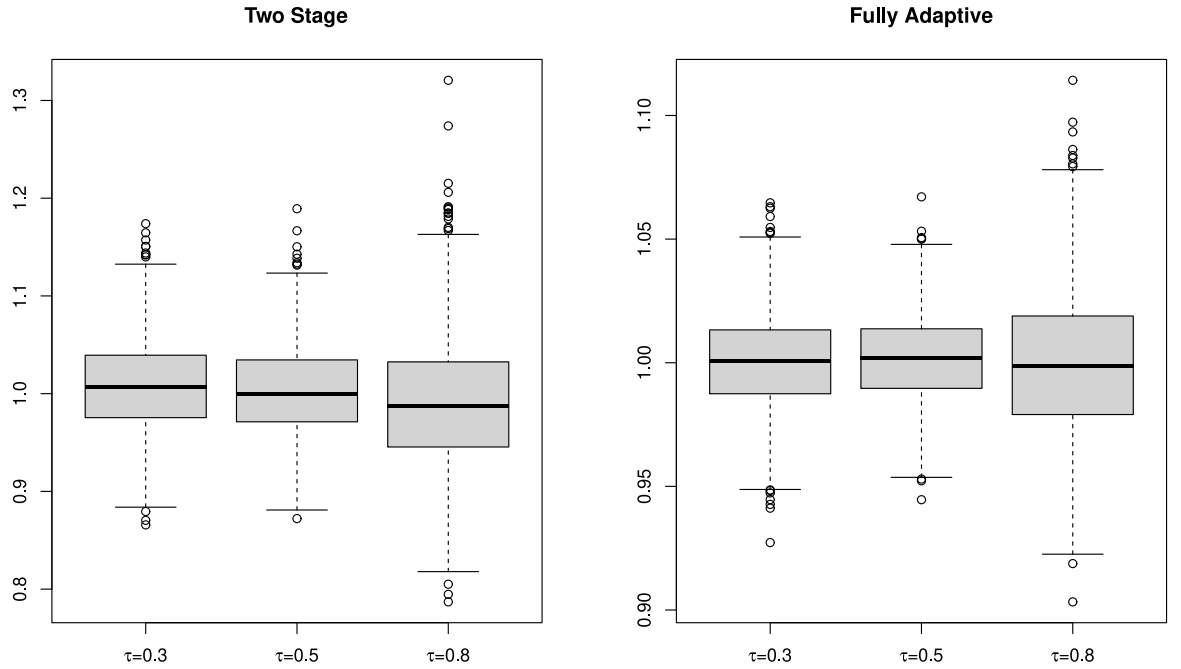

Fig. B.6. Boxplot of $\widehat{\mathbb{V}}_\tau^{1/2}/\mathbb{V}_\tau^{1/2}$ of different design under continuous case with $N = 5000$.

References

Abadie, Alberto, Angrist, Joshua, Imbens, Guido, 1998. Instrumental Variables Estimation of Quantile Treatment Effects. Technical Report, National Bureau of Economic Research, Cambridge, Mass., USA, http://dx.doi.org/10.3386/t0229.

Aletti, Giacomo, Ghiglietti, Andrea, Rosenberger, William F., 2018. Nonparametric covariate-adjusted response-adaptive design based on a functional urn model. Ann. Statist. 46 (6B), 3838–3866.

Anderson, S.J., Mckenzie, D., 2021. Improving business practices and the boundary of the entrepreneur: A randomized experiment comparing training, consulting, insourcing, and outsourcing. J. Polit. Econ. 130 (1), 157–209.
Antognini, Alessandro Baldi, Zagoraiou, Maroussa, 2012. Multi-objective optimal designs in comparative clinical trials with covariates: The reinforced doubly adaptive biased coin design. Ann. Statist. 40 (3), 1315–1345.
Antognini, Alessandro Baldi, Zagoraiou, Maroussa, 2015. On the almost sure convergence of adaptive allocation procedures. Bernoulli 21 (2), 881–908.
Atkinson, A.C., Biswas, A., Pronzato, L., 2011. Covariate-Balanced Response-Adaptive Designs for Clinical Trials with Continuous Responses That Target Allocation Probabilities. Technical Report Technical Report NI11042-DAE, Isaac Newton Institute for Mathematical Sciences.
Baldi Antognini, A., Zagoraiou, Maroussa, 2011. The covariate-adaptive biased coin design for balancing clinical trials in the presence of prognostic factors. Biometrika 98 (3), 519–535.
Bandyopadhyay, Uttam, Biswas, Atanu, 2001. Adaptive designs for normal responses with prognostic factors. Biometrika 88 (2), 409–419.
Bssett, Gilbert W., Koenker, Roger W., 1986. Strong consistency of regression quantiles and related empirical processes. Econometric Theory 2 (2), 191–201.
Cattaneo, Matias D., 2010. Efficient semiparametric estimation of multi-valued treatment effects under ignorability. J. Econometrics 155 (2), 138–154.
Chen, Xiaohong, Hong, Han, Tarozzi, Alessandro, 2008. Semiparametric efficiency in GMM models with auxiliary data. Ann. Stat. 36 (2), 808–843.
Chen, Xiaohong, Linton, Oliver, Van Keilegom, Ingrid, 2003. Estimation of semiparametric models when the criterion function is not smooth. Econometrica 71 (5), 1591–1608.
Chen, Xiaohong, Pouzo, Demian, Powell, James L., 2019. Penalized sieve GEL for weighted average derivatives of nonparametric quantile IV regressions. J. Econometrics 213 (1), 30–53.
Chernozhukov, Victor, Hansen, Christian, 2005. An IV model of quantile treatment effects. Econometrica 73 (1), 245–262.
Chernozhukov, Victor, Hansen, Christian, 2006. Instrumental quantile regression inference for structural and treatment effect models. J. Econometrics 132 (2), 491–525.
Doob, Joseph L., 1936. Note on probability. Ann. of Math. 37 (2), 363–367.
Fan, Jianqing, Gijbels, Irene, 1996. Local Polynomial Modelling and its Applications. Chapman and Hall, London.
Firpo, Sergio, 2007. Efficient semiparametric estimation of quantile treatment effects. Econometrica 75 (1), 259–276.
Gertler, Paul J., Martinez, Sebastian W., Rubio-Codina, Marta, 2012. Investing cash transfers to raise long-term living standards. Am. Econ. J.: Appl. Econ. 4 (1), 164–192.
Hirano, Keisuke, Imbens, Guido W., Ridder, Geert, 2003. Efficient estimation of average treatment effects using the estimated propensity score. Econometrica 71 (4), 1161–1189.
Hu, Feifang, Rosenberger, William F., 2006. The Theory of Response-Adaptive Randomization in Clinical Trials, vol. 525, John Wiley & Sons.
Hu, Jianhua, Zhu, Hongjian, Hu, Feifang, 2015. A unified family of covariate-adjusted response-adaptive designs based on efficiency and ethics. J. Amer. Statist. Assoc. 110 (509), 357–367.
Jiang, Liang, Phillips, Peter C.B., Tao, Yubo, Zhang, Yichong, 2023. Regression-adjusted estimation of quantile treatment effects under covariate-adaptive randomizations. J. Econometrics 234 (2), 758–776.
Karlan, Dean, List, John, 2007. Does price matter in charitable giving? Evidence from a large-scale natural field experiment. Amer. Econ. Rev. 97 (5), 1774–1793.
Karlan, Dean S., Zinman, Jonathan, 2008. Credit elasticities in less-developed economies: Implications for microfinance. Amer. Econ. Rev. 98 (3), 1040–1068.
Kharitonov, Eugene, Vorobev, Aleksandr, Macdonald, Craig, Serdyukov, Pavel, Ounis, Iadh, 2015. Sequential testing for early stopping of online experiments. In: Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 473–482.
Li, Xiaoqin, Yang, Wenzhi, Hu, Shuhe, 2016. Uniform convergence of estimator for nonparametric regression with dependent data. J. Inequal. Appl. 2016, 1–12.
Morgan, Kari Lock, Rubin, Donald B., 2012. Rerandomization to improve covariate balance in experiments. Ann. Statist. 40 (2), 1263–1282.
Morgan, Kari Lock, Rubin, Donald B., 2015. Rerandomization to balance tiers of covariates. J. Amer. Statist. Assoc. 110 (512), 1412–1421.
Pocock, Stuart J., Simon, Richard, 1975. Sequential treatment assignment with balancing for prognostic factors in the controlled clinical trial. Biometrics 31 (1), 103–115.
Robertson, David S, Lee, Kim May, Lopez-Kolkovska, Boryana C, Villar, Sofia S, 2020. Response-adaptive randomization in clinical trials: from myths to practical considerations. arXiv preprint arXiv:2005.00564.
Rosenberger, William F., Sverdlov, Oleksandr, 2008. Handling covariates in the design of clinical trials. Statist. Sci. 23 (3), 404–419.
Rosenberger, William F., Vidyashankar, A.N., Agarwal, Deepak K., 2001. Covariate-adjusted response-adaptive designs for binary response. J. Biopharm. Statist. 11 (4), 227–236.
Silverman, Bernard W., 1978. Weak and strong uniform consistency of the kernel estimate of a density and its derivaives.. Ann. Statist. 6 (1), 177–184.
Zelen, M., 1974. The randomization and stratification of patients of clinical rials. J. Chronic Dis. 27 (7), 365–375.
Zhang, Li-Xin, Hu, Feifang, Cheung, Siu Hung, Chan, Wai Sum, 2007. Asymptotic properties of covariate-adjusted response-adaptive designs. Ann. Statist. 35 (3), 1166–1182.
Zhang, Yichong, Zheng, Xin, 2020. Quantile treatment effects and bootstrap inference under covariate-adaptive randomization. Quant. Econ. 11 (3), 957–982.
Zhao, Wanying, Ma, Wei, Wang, Fan, Hu, Feifang, 2022. Incorporating covariates information in adaptive clinical trials for precision medicine. Pharm. Stat. 21 (1), 176–195.
Zhu, Hongjian, 2015. Covariate-adjusted response adaptive designs incorporating covariates with and without treatment interactions. Canad. J. Statist. 43 (4), 534–553.
Zhu, Hai, Zhu, Hongjian, 2023. Covariate-adjusted response-adaptive designs based on semiparametric approaches. Biometrics 79 (4), 2895–2906.