

Quantile regression approach to conditional mode estimation

Hirofumi Ota*

*Department of Statistics, Rutgers, the State University of New Jersey
110 Frelinghuysen Road, Piscataway, NJ 08854, USA
e-mail: ho105@rutgers.edu*

Kengo Kato

*Department of Statistics and Data Science, Cornell University
1194 Comstock Hall, Ithaca, NY 14853, USA
e-mail: kk976@cornell.edu*

and

Satoshi Hara

*The Institute of Scientific and Industrial Research, Osaka University
Mihogaoka 8-1, Ibaraki, Osaka 567-0047, Japan
e-mail: satohara@ar.sanken.osaka-u.ac.jp*

Abstract: In this paper, we consider estimation of the conditional mode of an outcome variable given regressors. To this end, we propose and analyze a computationally scalable estimator derived from a linear quantile regression model and develop asymptotic distributional theory for the estimator. Specifically, we find that the pointwise limiting distribution is a scale transformation of Chernoff’s distribution despite the presence of regressors. In addition, we consider analytical and subsampling-based confidence intervals for the proposed estimator. We also conduct Monte Carlo simulations to assess the finite sample performance of the proposed estimator together with the analytical and subsampling confidence intervals. Finally, we apply the proposed estimator to predicting the net hourly electrical energy output using Combined Cycle Power Plant Data.

MSC 2010 subject classifications: Primary 62J02; secondary 62G20.

Keywords and phrases: Chernoff’s distribution, cube root asymptotics, modal regression, quantile regression.

Received November 2018.

Contents

1	Introduction	3121
2	Setup and estimator	3124
3	Limiting distributions	3127
	3.1 Limiting distributions	3127

*Partially supported by the Grant-in-Aid for JSPS Fellows (18J20499) from the Japan Society for the Promotion of Science.

3.2	Inference	3132
3.2.1	Analytical confidence intervals	3132
3.2.2	Subsampling	3134
4	Numerical results	3135
4.1	Bandwidth selection	3135
4.2	Simulation results	3135
4.2.1	Comparison of RMSEs	3135
4.2.2	Coverage probabilities of confidence intervals	3136
4.3	Combined cycle power plant data	3138
5	Discussion	3140
A	Proofs	3141
A.1	Preliminaries	3141
A.2	Proof of Theorem 1	3143
A.3	Proof of Corollary 1	3152
A.4	Proof of Proposition 1	3154
B	Convergence of maximum of Chernoff random variables	3155
	Acknowledgments	3157
	References	3157

1. Introduction

Estimation of the conditional mode of an outcome variable given regressors, called *modal regression*, is an active research area in the recent statistics literature. In particular, if the conditional distribution is highly skewed or has fat tails, then one would be more interested in the conditional mode than the conditional mean or median since in such cases the mean or median may fail to capture a major trend of the conditional distribution. As such, modal regression has a wide variety of applications including the analysis of traffic and forest fire data [14, 53], econometrics [34, 35, 25, 21], and machine learning [46, 16]. For example, [25] argue that the mode is the most intuitive measure of central tendency for positively skewed data found in many econometric applications such as wages, prices, and expenditures ([25], p. 93). See also [7] and [5] for recent reviews on modal regression.

Existing approaches to estimation of the conditional mode includes nonparametric kernel estimation [8] and linear modal regression [34, 35, 25, 53], among others. The nonparametric estimation is able to avoid model misspecification but has slow rates of convergence that deteriorate as the number of regressors increases. Namely, if the number of continuous regressors is p , then the rate of convergence of the kernel density based estimator in [8] is at best $n^{-2/(p+7)}$ under four times differentiability of the joint density. On the other hand, the linear modal regression is able to avoid such “curse of dimensionality” but requires to solve a multi-dimensional non-convex optimization problem.

In this paper, we propose a new estimator for the conditional mode that is able to avoid the curse of dimensionality and at the same time is computationally scalable, thereby complementing the above existing methods. The

proposed method is based on the observation that the derivative of the conditional quantile function with respect to the quantile index is the reciprocal of the conditional density evaluated at the conditional quantile function and hence the conditional mode is obtained by minimizing the derivative of the conditional quantile function. Specifically, we assume a linear quantile regression model to estimate the conditional quantile function as in [29] (see also [28]), and estimate its derivative by a numerical differentiation of the estimated conditional quantile function. The proposed estimator is then obtained by minimizing the estimated derivative. Notably, the proposed method is computationally attractive since the computation of the quantile regression estimate can be formulated as a linear programming problem and so is highly scalable (cf. Chapter 6 in [28]), and the minimization of the estimated derivative is a one-dimensional optimization problem and so can be carried out by a grid search.

We develop asymptotic theory for the proposed estimator, which turns out to be non-standard. Specifically, we find that the proposed estimator has convergence rate $(nh^2)^{-1/3}$ where n is the sample size and $h = h_n \rightarrow 0$ is a sequence of bandwidths, and the limiting distribution is a scale transformation of Chernoff's distribution [9]. Chernoff's distribution is defined as the distribution of a maximizer of a two-sided Brownian motion with a negative quadratic drift, and appears as e.g. limiting distributions of estimators for monotone functions; see [20]. Our result on the limiting distribution would be of interest from theoretical and practical perspectives. First, the proposed estimator provides a new example of estimators having Chernoff's distribution as limiting distributions, which would be of theoretical interest. Second, the fact that the limiting distribution is a scale transformation of Chernoff's distribution makes inference for our estimator relatively simple. This is in contrast to e.g. Manski's maximum score [39] whose limiting distribution is a maximizer of a Gaussian process with its covariance function depending on the distribution of regressors; see [27]. Building upon the limiting distribution, we develop inference methods for our estimator. The one is an analytical confidence interval based on consistently estimating the scaling constant, and the other is based on the subsampling [41, 42]. We also derive a multivariate limit theorem for the proposed estimator, which can be used to construct simultaneous confidence intervals for the modal function over finite design points.

In addition to the theoretical results, we conduct Monte Carlo simulations to assess the finite sample performance of the proposed estimator together with the analytical and subsampling confidence intervals. We suggest a practical method to choose the bandwidth based upon the idea suggested in [30]. We compare the performance of the proposed estimator with the linear modal regression estimator of [25, 53] via the root mean square error for the two data generating processes where the true modal function is linear or nonlinear. Finally, we apply the proposed estimator to predicting the net hourly electrical energy output using Combined Cycle Power Plant Data [24, 49]. These numerical results show evidence that the proposed estimator works well in the finite sample.

The literature related to this paper is broad. Nonparametric estimation of the unconditional mode goes back to Parzen [40] and Chernoff [9] in 1960s;

see also [44]. Modal regression originates from [45] and the literature has flourished since then [34, 35, 14, 25, 54, 53, 8, 55, 46, 21, 32, 26, 16]. However, none of these papers do not consider a quantile regression based estimator for the conditional mode. [34, 35, 25, 53] consider linear modal regression; [34, 35] assume a restrictive condition that the conditional distribution is symmetric around the origin to derive limiting distributions of the estimators. The symmetry of the conditional distribution implies that the conditional mean, median, and mode are all identical. Subsequently, [25, 53] relax the symmetry assumption and propose estimators that enjoy asymptotic normality. In the present paper, instead of linearity of the conditional mode, we assume a linear quantile regression model. Importantly, the linear quantile regression model does not imply linearity of the conditional mode, and so there are no strict inclusion relations between the two assumptions; see Remark 1 ahead. The recent work of [8] studies nonparametric kernel estimation of the conditional mode. To be precise, [8] do not assume the existence of the unique global mode and allow for multiple local modes. Extension of our approach to multiple local modes would be of interest but is beyond the scope of the present paper. [54] propose a local modal regression (LMR) estimator that can be seen as a local linear estimator for the conditional mode, and establish asymptotic results analogous to those of a local linear estimator for the conditional mean. In particular, the rate of convergence of the LMR estimator is faster than that of the kernel density based estimator of [8]. This is, however, due to Condition (A6) in [54] that is essentially the conditional symmetry assumption on the error term (note that h_2 in [54] is fixed) and under which the conditional mode and mean coincide. In the present paper, we assume no symmetry assumptions on the conditional distribution.

From a technical point of view, derivation of the limiting distribution of the proposed estimator is by no means trivial. First of all, it is not a priori straightforward to foresee that the convergence rate is $(nh^2)^{-1/3}$ and the limiting distribution is a scale transformation of Chernoff's distribution. Second, because our objective function depends on the bandwidth tending to zero as the sample size increases, our result does not follow from the general theorem, Theorem 1.1, in [27], which is a pioneering work on cube root asymptotic theory. The recent work of [48] extends [27] to allow the objective function to depend on the bandwidth, but some of their regularity conditions are severely restrictive or difficult to verify in our problem. Hence, we provide a separate and self-contained proof of the main theorem, Theorem 1 ahead, which requires a substantial work. See also the discussion after Theorem 1.

The rest of the paper is organized as follows. In Section 2, we state the formal setup and define the estimator. In Section 3, we derive limiting distributions of the proposed estimator and develop inference methods for it. In Section 4, we conduct Monte Carlo simulations to assess the finite sample performance of the proposed estimator together with the analytical and subsampling confidence intervals. In addition, we apply the proposed estimator to predicting the net hourly electrical energy output using Combined Cycle Power Plant Data. Section 5 concludes. All the proofs are gathered in Appendix.

2. Setup and estimator

In this paper, we are interested in estimating the conditional mode of an outcome variable $Y \in \mathbb{R}$ given a vector of regressors $X = (X_1, \dots, X_d)^T \in \mathbb{R}^d$. In what follows, we assume that there exists a conditional density $f(y | x)$ of Y given X that is (at least) continuous in y , and for each design point x in the support of X , there exists a unique mode $m(x)$, i.e., there exists a unique maximizer of the function $y \mapsto f(y | x)$:

$$f(m(x) | x) = \max_{y \in \mathbb{R}} f(y | x).$$

The function $m(x)$ is called the modal function.

We base our estimation strategy of the modal function $m(x)$ on inverting a quantile regression model. Let $Q(\tau | X)$ denote the conditional τ -quantile of Y given X for $\tau \in (0, 1)$. For the notational convenience, we also write $Q_x(\tau) = Q(\tau | X = x)$. To see the link between the conditional quantile function and the modal function, we begin with observing that

$$s_x(\tau) := Q'_x(\tau) = \frac{\partial Q_x(\tau)}{\partial \tau} = \frac{1}{f(Q_x(\tau) | x)}$$

assuming some regularity conditions that will be clarified below. Hence, defining

$$\tau_x = \arg \min_{\tau \in (0, 1)} s_x(\tau),$$

which exists and is unique (by continuity and strict positivity of the function $y \mapsto f(y | x)$ around the mode $m(x)$), we arrive at the key identity

$$m(x) = Q_x(\tau_x).$$

The function $\tau \mapsto s_x(\tau)$ (called the “sparsity” function) can be estimated by a numerical differentiation of an estimator of the conditional quantile function $\tau \mapsto Q_x(\tau)$, and so the problem boils down to estimating the conditional quantile function. To this end, we assume a linear quantile regression model:

$$Q(\tau | X) = X^T \beta(\tau), \quad \tau \in (0, 1), \quad (1)$$

where $\beta(\tau) \in \mathbb{R}^d$ is an unknown slope vector for each $\tau \in (0, 1)$.

Pick any design point x in the support of X , and consider to estimate $m(x)$. Let $(Y_1, X_1), \dots, (Y_n, X_n)$ be i.i.d. observations of (Y, X) . We estimate the slope vector $\beta(\tau)$ by

$$\hat{\beta}(\tau) = \arg \min_{\beta \in \mathbb{R}^d} \sum_{i=1}^n \rho_\tau(Y_i - X_i^T \beta), \quad (2)$$

where $\rho_\tau(u) = \{\tau - I(u \leq 0)\}u$ is the check function [29]. This leads to an estimator $\hat{Q}_x(\tau) = x^T \hat{\beta}(\tau)$ of $Q_x(\tau)$. To estimate $s_x(\tau) = Q'_x(\tau)$, let $h = h_n \rightarrow 0$

be a sequence of bandwidths such that $nh^2 \rightarrow \infty$; then we estimate $s_x(\tau)$ by a numerical differentiation:

$$\widehat{s}_x(\tau) = \frac{\widehat{Q}_x(\tau + h) - \widehat{Q}_x(\tau - h)}{2h}.$$

Finally, we estimate $m(x)$ by $\widehat{m}(x) = \widehat{Q}_x(\widehat{\tau}_x) = x^T \widehat{\beta}(\widehat{\tau}_x)$, where $\widehat{\tau}_x$ is an approximate minimizer of $\widehat{s}_x(\tau)$ on $[\varepsilon, 1 - \varepsilon]$ with sufficiently small parameter $\varepsilon \in (0, 1/2)$ chosen by users, in the sense that

$$\widehat{s}_x(\widehat{\tau}_x) \leq \inf_{\tau \in [\varepsilon, 1 - \varepsilon]} \widehat{s}_x(\tau) + o((nh^2)^{-2/3}).$$

The objective function $\widehat{s}_x(\tau)$ may not admit strict minimizers, and so we allow $\widehat{\tau}_x$ to be an approximate minimizer in the above sense, which always exists. In practice, our estimator requires to choose the bandwidth h , which will be discussed in Section 4.1.

Importantly, our estimate $\widehat{m}(x)$ is easy to compute even when the sample size n and the dimension d of X are large. The quantile regression problem (2) can be formulated as a linear programming problem and hence can be efficiently solved even when n and d are large (cf. Chapter 6 in [28]). Furthermore, the entire path $\tau \mapsto \widehat{\beta}(\tau)$ can be computed by a parametric linear programming or discretizing the interval $(0, 1)$ into fine grids. The minimization of $\widehat{s}_x(\tau)$ is a one-dimensional optimization problem and can be solved by a grid search. On the other hand, the linear modal regression estimator [34, 35, 25, 53] requires to solve a multi-dimensional non-convex optimization problem. For example, [53] assume that the modal function is linear $m(x) = x^T \gamma$ for some $\gamma \in \mathbb{R}^d$ and propose the following estimator:

$$\widehat{\gamma}_{YL} = \arg \max_{\gamma} \sum_{i=1}^n \phi_h(Y_i - X_i^T \gamma), \quad (3)$$

where $\phi(y) = (2\pi)^{-1/2} e^{-y^2/2}$ is the density of the standard normal distribution and $\phi_h(y) = h^{-1} \phi(y/h)$. The optimization problem (3) is non-convex. [53] propose an EM like algorithm for (3), but “there is no guarantee that the algorithm will converge to the global optimal solution” ([53], p. 659).

Remark 1 (Generality of linear quantile regression model). The linear quantile regression model (1) is flexible enough to cover many data generating processes. In general, if $\tau \mapsto \beta(\tau)$ is a function on $(0, 1)$ such that the map $\tau \mapsto X^T \beta(\tau)$ is strictly increasing almost surely and Y is generated as $Y = X^T \beta(U)$ for $U \sim U(0, 1)$ independent of X , then the pair (Y, X) satisfies the linear quantile regression model (1). In particular, it is worth pointing out that the linear quantile regression model (1) does not imply linearity of the modal function $m(x)$. For example, consider the simple case where $X = (1, X_2)^T$ with $X_2 \in (0, 1)$ and $Y = U^3/3 - X_2(U - 1)^2$ for $U \sim U(0, 1)$ independent of X . In this case, the pair (Y, X) satisfies the model (1) with $\beta(\tau) = (\tau^3/3, -(\tau - 1)^2)^T$ and so $Q_x(\tau) = \tau^3/3 - (\tau - 1)^2 x_2$. Since $Q'_x(\tau) = \tau^2 - 2(\tau - 1)x_2$ is minimized at $\tau = x_2$, the modal function $m(x) = Q_x(x_2) = -2x_2^3/3 + 2x_2^2 - x_2$ is nonlinear.

Remark 2 (Case with no regressors). In the simple case where there are no regressors, i.e., $X = 1$, our estimator of the mode reduces to $\hat{m} = \hat{Q}(\hat{\tau})$, where $\hat{Q}(\tau) = \hat{F}^{-1}(\tau) = \inf\{y : \hat{F}(y) \geq \tau\}$ is the empirical quantile function (with $\hat{F}(y) = n^{-1} \sum_{i=1}^n I(Y_i \leq y)$ being the empirical distribution function) and

$$\hat{\tau} = \arg \min_{\tau} \frac{\hat{Q}(\tau + h) - \hat{Q}(\tau - h)}{2h}.$$

Our estimator \hat{m} can also be described by using order statistics $Y_{(1)} \leq \dots \leq Y_{(n)}$. Since in general $\hat{Q}(\tau) = Y_{(\lceil n\tau \rceil)}$ where $\lceil \cdot \rceil$ is the ceiling function, our estimator \hat{m} coincides with the order statistic $Y_{(\lceil n\hat{\tau} \rceil)}$ where $\hat{\tau}$ minimizes the spacing $Y_{(\lceil n(\tau+h) \rceil)} - Y_{(\lceil n(\tau-h) \rceil)}$.

It is then clear that our estimator is (related to but) markedly different from Chernoff's [9] estimator of the unconditional mode of Y that is defined by

$$\hat{m}_C = \arg \max_y \frac{\hat{F}(y + h) - \hat{F}(y - h)}{2h},$$

namely, \hat{m}_C is the point whose local neighborhood contains the most observations.

Remark 3 (Alternative objective function). The estimator $\hat{s}_x(\tau)$ of $s_x(\tau)$ contains a deterministic bias of order h^2 under the conditions stated in the next section. Alternatively, we may estimate $s_x(\tau)$ by

$$\tilde{s}_x(\tau) = \frac{\frac{2}{3}\{\hat{Q}_x(\tau + h) - \hat{Q}_x(\tau - h)\} - \frac{1}{12}\{\hat{Q}_x(\tau + 2h) - \hat{Q}_x(\tau - 2h)\}}{h}, \quad (4)$$

which has a bias of order h^4 under additional smoothness conditions; cf. [3]. In the present paper, however, we shall use a simpler objective function $\hat{s}_x(\tau)$.

Remark 4 (Implementation detail). In the finite sample, $[\tau - h, \tau + h]$ may not be included in $(0, 1)$ for some $\tau \in [\varepsilon, 1 - \varepsilon]$. To fix this, we suggest the following simple modification. Suppose that we compute $\hat{Q}_x(\tau)$ on $[\tau_{\min}, \tau_{\max}] \supset [\varepsilon, 1 - \varepsilon]$; then in practice we suggest to replace $\hat{s}_x(\tau)$ by

$$\hat{s}_x(\tau) = \frac{\hat{Q}_x(\tau + \min\{h, \tau_{\max} - \tau\}) - \hat{Q}_x(\tau - \min\{h, \tau - \tau_{\min}\})}{\min\{h, \tau_{\max} - \tau\} + \min\{h, \tau - \tau_{\min}\}},$$

which asymptotically coincides with the original definition of $\hat{s}_x(\tau)$ uniformly in $\tau \in [\varepsilon, 1 - \varepsilon]$ (as long as $(\tau_{\min}, \tau_{\max}) \supset [\varepsilon, 1 - \varepsilon]$).

Remark 5 (Alternative specifications to the conditional quantile function). In the present paper, we assume that the conditional quantile function is linear in X . The linear quantile regression model is the most fundamental modeling in conditional quantile estimation, and is computationally attractive since the computation of the Koenker-Bassett [29] estimate can be formulated as a linear

programming problem. Indeed, the computational attractiveness is one of the main motivations to study the proposed estimator of the conditional mode.

Having said that, we could use alternative specifications to the conditional quantile function to estimate the conditional mode. One possible alternative is a nonlinear quantile regression model $Q_x(\tau) = g(x, \beta(\tau))$ where g is some known smooth function (the dimensions of x and $\beta(\tau)$ need not be matched); see e.g. Section 4.4 of [28]. In this case, we can estimate $\beta(\tau)$ by

$$\widehat{\beta}(\tau) = \arg \min_{\beta} \sum_{i=1}^n \rho_{\tau}(Y_i - g(X_i, \beta)),$$

and thus can estimate $s_x(\tau) = Q'_x(\tau)$ by $\widehat{s}_x(\tau) = \{\widehat{Q}_x(\tau + h) - \widehat{Q}_x(\tau - h)\}/(2h)$ with $\widehat{Q}_x(\tau) = g(x, \widehat{\beta}(\tau))$. Alternatively, we can use the expression

$$s_x(\tau) = \left[\underbrace{\frac{\partial g(x, \beta)}{\partial \beta} \Big|_{\beta = \beta(\tau)}}_{=: g_{\beta}(x, \beta(\tau))} \right]^T \frac{d\beta(\tau)}{d\tau},$$

and estimate $s_x(\tau)$ by $\widehat{s}_x(\tau) = g_{\beta}(x, \widehat{\beta}(\tau))^T \{\widehat{\beta}(\tau + h) - \widehat{\beta}(\tau - h)\}/(2h)$. It is known that under regularity conditions, similar asymptotic properties to those of the linear quantile regression estimator hold for the nonlinear case (cf. Section 4.4 of [28]), and hence it is natural to expect that asymptotic results analogous to those developed in the next section can be extended to the resulting conditional mode estimator under the nonlinear quantile regression model.

A yet alternative specification would be a semiparametric single index model $Q_x(\tau) = \psi(x^T \beta(\tau))$ where ψ is some unknown function. For given estimators $\widehat{\psi}$ and $\widehat{\beta}(\tau)$ of ψ and $\beta(\tau)$, we can estimate $Q_x(\tau)$ and $s_x(\tau)$ by $\widehat{Q}_x(\tau) = \widehat{\psi}(x^T \widehat{\beta}(\tau))$ and $\widehat{s}_x(\tau) = \{\widehat{Q}_x(\tau + h) - \widehat{Q}_x(\tau - h)\}/(2h)$, respectively. Alternatively, we can use the expression $s_x(\tau) = \psi'(x^T \beta(\tau)) d(x^T \beta(\tau))/d\tau$ and estimate $d(x^T \beta(\tau))/d\tau$ by a difference quotient. Methods to estimate the parametric and nonparametric components in the single index quantile regression model can be found in e.g. [6, 52], and [38]. In the single index case, the nonparametric estimation of the link function ψ is involved, whose effect has to be taken into account when considering asymptotic properties of the resulting conditional mode estimator, which would be a nontrivial challenge.

3. Limiting distributions

3.1. Limiting distributions

In this section, we derive limiting distributions of $\widehat{\tau}_x$ and $\widehat{m}(x)$. To this end, we make the following assumption. Let \mathcal{X} denote the support of X .

Assumption 1. *In addition to the baseline assumption stated in the previous section, we assume the following conditions.*

- (i) $\mathbb{E}[X_j^4] < \infty$ for all $j = 1, \dots, d$.
- (ii) The matrix $\mathbb{E}[XX^T]$ is positive definite.
- (iii) The conditional density $f(y | x)$ is three times continuously differentiable with respect to y for each $x \in \mathcal{X}$. Let $f^{(j)}(y | x) = \partial^j f(y | x) / \partial y^j$ for $j = 0, 1, 2, 3$, where $f^{(0)}(y | x) = f(y | x)$. There exists a constant C such that $|f^{(j)}(y | x)| \leq C$ for all $(y, x) \in \mathbb{R} \times \mathcal{X}$ and $j = 0, 1, 2, 3$.
- (iv) There exists a positive constant c (that may depend on ε) such that $f(y | x) \geq c$ for all $y \in [Q_x(\varepsilon/2), Q_x(1 - \varepsilon/2)]$ and $x \in \mathcal{X}$.
- (v) As $n \rightarrow \infty$, $nh^8 \rightarrow 0$ and $nh^5 \rightarrow \infty$.

Conditions (i)–(iv) are more or less standard in the quantile regression literature; cf. [28]. In particular, they require no moment conditions on Y . For instance, they allow $\mathbb{E}[|Y|] = \infty$. Conditions (iii) and (iv) allow $Q_x(\tau)$ to be four times continuously differentiable on $(\varepsilon/2, 1 - \varepsilon/2)$ with

$$s_x(\tau) := Q'_x(\tau) = \frac{1}{f(Q_x(\tau) | x)}, \quad s'_x(\tau) = \frac{-f^{(1)}(Q_x(\tau) | x)}{f(Q_x(\tau) | x)^3},$$

$$s''_x(\tau) = \frac{3f^{(1)}(Q_x(\tau) | x)^2 - f(Q_x(\tau) | x)f^{(2)}(Q_x(\tau) | x)}{f(Q_x(\tau) | x)^5}.$$

Condition (v) is concerned with the bandwidth. The condition $nh^8 \rightarrow 0$ is an “undersmoothing” condition. The proof of Theorem 1 shows that the estimator $\hat{m}(x)$ contains a deterministic bias of order h^2 , while the stochastic error decreases at rate $(nh^2)^{-1/3}$. To guarantee that $h^2 = o((nh^2)^{-1/3})$, we need $nh^8 \rightarrow 0$.

Let $\{B(t) : t \in \mathbb{R}\}$ be a two-sided standard Brownian motion, i.e., a centered Gaussian process with continuous sample paths and covariance function

$$\mathbb{E}[B(t_1)B(t_2)] = \begin{cases} t_1 & \text{if } 0 \leq t_1 \leq t_2 \\ -t_2 & \text{if } t_1 \leq t_2 \leq 0 \\ 0 & \text{if } t_1 \leq 0 \leq t_2 \end{cases}.$$

Such a two-sided standard Brownian motion can be constructed by generating independent standard Brownian motions $\{W_1(t) : t \geq 0\}$ and $\{W_2(t) : t \geq 0\}$, and then defining $B(t) = W_1(t)$ for $t \geq 0$ and $B(t) = W_2(-t)$ for $t < 0$. In addition, let

$$Z = \arg \max_{t \in \mathbb{R}} \{B(t) - t^2\},$$

which exists and is unique almost surely by Lemmas 2.5 and 2.6 in [27]. The distribution of Z is called *Chernoff's distribution* [9], and numerical values of quantiles of Chernoff's distribution can be found in [20].

Finally, define the matrix

$$J(\tau) = \mathbb{E}[f(X^T \beta(\tau) | X)XX^T],$$

which is positive definite for every $\tau \in [\varepsilon/2, 1 - \varepsilon/2]$ under our assumption. We are now in position to state the main theorem of this paper.

Theorem 1 (Limiting distributions). *Pick any $x \in \mathcal{X}$. Suppose that Assumption 1 holds, and in addition that $f^{(2)}(m(x) | x) < 0$ and $m(x) \in (Q_x(\varepsilon), Q_x(1 - \varepsilon))$ (or equivalently $\tau_x \in (\varepsilon, 1 - \varepsilon)$). Then we have*

$$(nh^2)^{1/3}(\widehat{\tau}_x - \tau_x) \xrightarrow{d} (\sigma_x/v_x)^{2/3}Z \quad (5)$$

as $n \rightarrow \infty$, where $\sigma_x^2 = \mathbb{E}[(x^T J(\tau_x)^{-1} X)^2]/2$, $v_x = s_x''(\tau_x)/2 = -f^{(2)}(m(x) | x)/\{2f(m(x) | x)^4\}$ (> 0), and $\sigma_x = \sqrt{\sigma_x^2}$. In addition, we have

$$(nh^2)^{1/3}(\widehat{m}(x) - m(x)) \xrightarrow{d} s_x(\tau_x)(\sigma_x/v_x)^{2/3}Z. \quad (6)$$

Remark 6 (Rates of convergence). The rate of convergence of $\widehat{m}(x)$ toward $m(x)$ is $(nh^2)^{-1/3}$ and can be arbitrarily close to $n^{-1/4}$ under Condition (v), which is independent of the dimension d of the regressor vector. The $n^{-1/4}$ rate is likely to be suboptimal from a minimax point of view since it is known that when $X = 1$, the minimax rate of estimating the mode under three time differentiability of the underlying density is $n^{-2/7}$; see Theorem 3.1 in [44]. It is worth noting that if we use the alternative objective function $\widetilde{s}_x(\tau)$ in (4), the bias of the resulting estimator $\widehat{m}(x)$ is reduced to $O(h^4)$ (this however requires additional smoothness conditions on the conditional density), and therefore the rate of convergence can be arbitrarily close to $n^{-2/7}$.

Remark 7 (Case with no regressors). In the simple case where there are no regressors, i.e., $X = 1$, the limiting distribution of our estimator \widehat{m} is as follows. Let f denote the density of Y with mode m ; then $(nh^2)^{1/3}(\widehat{m} - m) \xrightarrow{d} \{2f(m)^3/f''(m)^2\}^{1/3}Z$. In contrast, the limiting distribution of Chernoff's mode estimator (see Remark 2) is $(nh^2)^{1/3}(\widehat{m}_C - m) \xrightarrow{d} \{2f(m)/f''(m)^2\}^{1/3}Z$, which is slightly different from our limiting distribution. It is worth mentioning that Chernoff's derivation of the preceding limiting distribution in [9] is only heuristic, but can be made rigorous (under regularity conditions) by mimicking the proof of Theorem 1.

Interestingly, despite the presence of regressors, the limiting distribution of our estimator $\widehat{m}(x)$ is a scale transformation of Chernoff's distribution, which is in contrast to e.g. Manski's maximum score [39] whose limiting distribution is given by a maximizer of a Gaussian process whose covariance function depends on the distribution of regressors; see Example 6.4 in [27]. The fact that the limiting distribution is a scale transformation of Chernoff's distribution makes inference for our estimator $\widehat{m}(x)$ relatively simple. Namely, an asymptotic confidence interval can be constructed by consistently estimating the constant $s_x(\tau_x)(\sigma_x/v_x)^{2/3}$, which will be discussed in the next section.

The main part of the proof is the proof of the first result (5). The second result (6) follows from the \sqrt{n} -uniform consistency of the quantile regression estimator and the delta method. To prove the first result (5), we begin with expanding the objective function $\widehat{s}_x(\tau)$ and showing that $\widehat{\tau}_x$ is an approximate minimizer of the sample average of kernel functions with a uniform kernel; see (13) in the proof. Since those kernel functions depend on the sample size n via the bandwidth

$h = h_n$, the result (5) does not follow from the general theorem, Theorem 1.1, in [27], which is a pioneering work on cube root asymptotic theory. Theorem 1.1 in [27] covers the case where the objective function is the sample average of functions that do not depend on n and the estimator is $n^{1/3}$ -consistent, but its proof does not carry over to our case (cf. the second paragraph in page 192 of [27]). The recent work of [48] extends the results of [27] to allow the objective function to depend on the bandwidth (and the data to be dependent), but some of their assumptions are severely restrictive or difficult to verify in our problem. Specifically, Assumption M (i) in [48] requires $h_n f_{n,\theta}$ (in their notation) to be uniformly bounded, which in our problem requires the regressor vector X to be bounded (recall that we only assume each coordinate of X to have finite fourth moment); and we (the authors) found that Assumption M (ii) is difficult to verify in our problem. Hence, instead of checking the assumptions of [48], we provide a separate and self-contained proof of the result (5), which requires a substantial work. Specifically, we show that the “rescaled” objective function for which the rescaled estimator $\hat{t} = (nh^2)^{1/3}(\hat{\tau}_x - \tau_x)$ is an approximate maximizer converges weakly to the process $\{\sigma_x B(t) - v_x t^2 : t \in \mathbb{R}\}$ in the space of locally bounded functions on \mathbb{R} , and apply Theorem 2.7 in [27] to conclude that the approximate maximizer $\hat{t} = (nh^2)^{1/3}(\hat{\tau}_x - \tau_x)$ converges weakly to $\arg \max_{t \in \mathbb{R}} \{\sigma_x B(t) - v_x t^2\}$, which is shown to be equal in distribution to $(\sigma_x/v_x)^{2/3}Z$; see Step 5 of the proof.

Next, we consider a multivariate limit theorem for the proposed estimator. Let $x^1, \dots, x^L \in \mathcal{X}$ be a finite number of design points with L independent of n , and let

$$\tau_{(1)} > \tau_{(2)} > \dots > \tau_{(M)}$$

denote the distinct values of $\tau_{x^1}, \dots, \tau_{x^L}$. Set $S_k = \{j \in \{1, \dots, L\} : \tau_{x^j} = \tau_{(k)}\}$ with $s_k = \text{Card}(S_k)$ for $k = 1, \dots, M$. For each $k = 1, \dots, M$, let $\{\mathbb{B}_k((t_j)_{j \in S_k}) : (t_j)_{j \in S_k} \in \mathbb{R}^{s_k}\}$ denote a centered Gaussian process with covariance function

$$\begin{aligned} & \text{Cov}(\mathbb{B}_k((t_i)_{i \in S_k}), \mathbb{B}_k((t'_j)_{j \in S_k})) \\ &= \frac{1}{2} \sum_{i,j \in S_k} (x^i)^T J(\tau_{(k)})^{-1} \mathbb{E}[X X^T] J(\tau_{(k)})^{-1} x^j \mathbb{E}[B(t_i)B(t'_j)]. \end{aligned}$$

We note that the construction of the \mathbb{B}_k -process depends on the design points x^1, \dots, x^L . Recall that a version of a stochastic process is another process with the same finite dimensional distributions.

Corollary 1. *Suppose that Assumption 1 holds, and in addition that $f^{(2)}(m(x) | x) < 0$ and $m(x) \in (Q_x(\varepsilon), Q_x(1 - \varepsilon))$ for all $x \in \{x^1, \dots, x^L\}$. Then, for each $k = 1, \dots, M$, there exists a version of the \mathbb{B}_k -process with continuous paths, and denoting the continuous version by the same symbol \mathbb{B}_k , we have*

$$(nh^2)^{1/3}(\hat{\tau}_{x^1}, \dots, \hat{\tau}_{x^L})^T \xrightarrow{d} (W_1, \dots, W_L)^T$$

as $n \rightarrow \infty$, where $(W_j)_{j \in S_k}, k = 1, \dots, M$ are independent, and for each $k =$

$1, \dots, M,$

$$(W_j)_{j \in S_k} \stackrel{d}{=} \arg \max_{(t_j)_{j \in S_k} \in \mathbb{R}^{s_k}} \left\{ \mathbb{B}_k((t_j)_{j \in S_k}) - \sum_{j \in S_k} v_{x^j} t_j^2 \right\}.$$

In addition, we have

$$(nh^2)^{1/3} (\widehat{m}(x^1) - m(x^1), \dots, \widehat{m}(x^L) - m(x^L))^T \\ \xrightarrow{d} (s_{x^1}(\tau_{x^1})W_1, \dots, s_{x^L}(\tau_{x^L})W_L)^T.$$

In the special case when $\tau_{x^1}, \dots, \tau_{x^L}$ are all distinct, we have

$$(nh^2)^{1/3} (\widehat{m}(x^1) - m(x^1), \dots, \widehat{m}(x^L) - m(x^L))^T \\ \xrightarrow{d} \left(s_{x^1}(\tau_{x^1})(\sigma_{x^1}/v_{x^1})^{2/3} Z_1, \dots, s_{x^L}(\tau_{x^L})(\sigma_{x^L}/v_{x^L})^{2/3} Z_L \right)^T,$$

where Z_1, \dots, Z_L are independent Chernoff random variables.

Corollary 1 implies that

$$(nh^2)^{1/3} \max_{1 \leq j \leq L} |\widehat{m}(x^j) - m(x^j)| \xrightarrow{d} \max_{1 \leq j \leq L} |s_{x^j}(\tau_{x^j})W_j|,$$

which can be used to construct simultaneous confidence intervals for $m(x)$ over the design points x^1, \dots, x^L ; see Remark 11 ahead.

Remark 8 (Uniform rate over expanding sets of design points). It is of interest to study the rate of convergence and limiting distribution of the L^∞ -distance between the proposed estimator and the true modal function on a continuum set of design points or expanding sets of design points, since e.g. such limiting distribution enables us to construct simultaneous confidence bands. To the best of our knowledge, however, much less is known about the rate of convergence and (especially) limiting distribution for the L^∞ -distance in nonstandard nonparametric estimation problems than standard nonparametric estimation problems with Gaussian limits, and we believe that the problem is challenging. One exception is the work of [12], which derives the uniform rate of convergence and the limiting distribution of the L^∞ -distance for the Grenander [18] estimator (precisely speaking [12] cover more general Grenander-*type* estimators); see also the recent review article by [13]. Their argument depends substantially on the specific construction of the Grenander estimator and can not be directly extended to our estimator. It is thus beyond the scope of this paper to formally study the uniform rate and the limiting distribution of the L^∞ -distance to our estimator, but we will give some heuristic discussion on this question, which we believe is of some interest to the reader.

To simplify the question, we confine ourselves to the maximum distance on expanding sets of design points x^1, \dots, x^L with $L = L_n \rightarrow \infty$. Suppose in addition that $\tau_{x^1}, \dots, \tau_{x^{L_n}}$ are all distinct. Then by Corollary 1 it is expected that

$\max_{1 \leq j \leq L_n} (nh^2)^{1/3} |\widehat{m}(x^j) - m(x^j)| / \{s_{x^j}(\tau_{x^j})(\sigma_{x^j}/v_{x^j})^{2/3}\}$ could be approximated by $\max_{1 \leq j \leq L_n} |Z_j| =: |Z|_{(L_n)}$ as long as $L_n \rightarrow \infty$ sufficiently slowly. In Appendix B, we will show that, for the norming constants

$$a_{L_n} = 3 \left(\frac{2}{3} \right)^{1/3} (\log L_n)^{2/3},$$

$$b'_{L_n} = \left(\frac{3}{2} \log L_n \right)^{1/3} - \frac{1}{a_{L_n}} \left[\kappa \left(\frac{3}{2} \log L_n \right)^{1/3} + \frac{1}{3} \log \log L_n + \frac{1}{3} \log \frac{3}{2} - \log(2\lambda) \right],$$

where λ and κ are positive constants (see Appendix B), we have

$$a_{L_n} (|Z|_{(L_n)} - b'_{L_n}) \xrightarrow{d} \underbrace{\Lambda}_{\text{Gumbel distribution}}.$$

In particular, $|Z|_{(L_n)} = b'_{L_n} + O_{\mathbb{P}}(1/a_{L_n}) = O_{\mathbb{P}}((\log L_n)^{1/3})$, and as long as L_n grows at most polynomially fast in n , $|Z|_{(L_n)} = O_{\mathbb{P}}((\log n)^{1/3})$. This suggests that the uniform rate of the proposed estimator would be $O_{\mathbb{P}}((nh^2/\log L_n)^{-1/3})$ and the maximum distance would converge in distribution to the Gumbel distribution after normalization. The preceding argument is heuristic since Corollary 1 only holds with fixed L (and extending the corollary to the case where $L = L_n \rightarrow \infty$ is a substantial technical challenge), and the rigorous result is left to future research.

3.2. Inference

3.2.1. Analytical confidence intervals

Theorem 1 allows us to construct pointwise confidence intervals for $m(x)$ by consistently estimating the nuisance parameters σ_x^2 , v_x , and $s_x(\tau_x)$.

The parameter $s_x(\tau_x)$ can be estimated by $\widehat{s}_x(\widehat{\tau}_x)$. Next, consider to estimate σ_x^2 . For the notational convenience, let $\Sigma = \mathbb{E}[XX^T]$ and so $\sigma_x^2 = x^T J(\tau_x)^{-1} \Sigma J(\tau_x)^{-1} x / 2$. The matrices Σ and $J(\tau)$ can be estimated by

$$\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n X_i X_i^T \quad \text{and} \quad \widehat{J}(\tau) = \frac{1}{2nh} \sum_{i=1}^n I(|Y_i - X_i^T \widehat{\beta}(\tau)| \leq h) X_i X_i^T,$$

respectively, so that we can estimate σ_x^2 by

$$\widehat{\sigma}_x^2 = \frac{1}{2} x^T \widehat{J}(\widehat{\tau}_x)^{-1} \widehat{\Sigma} \widehat{J}(\widehat{\tau}_x)^{-1} x,$$

where $\widehat{J}(\tau)$ is Powell's kernel estimator [43]. Finally, consider to estimate $v_x = s_x''(\tau_x)/2$. To this end, we estimate $s_x''(\tau) = Q_x'''(\tau)$ by a numerical differentiation

of $\widehat{Q}_x(\tau)$. Namely, define the operator Δ_h by $\Delta_h g(\tau) = (g(\tau+h) - g(\tau-h))/(2h)$, and $\Delta_h^j g = \Delta_h(\Delta_h^{j-1} g)$ recursively for $j = 2, 3, \dots$. Then we can estimate v_x by

$$\widehat{v}_x = \frac{1}{2} \Delta_h^3 \widehat{Q}_x(\widehat{\tau}_x). \tag{7}$$

The bandwidths used in $\widehat{J}(\tau)$ and \widehat{v}_x can be different from that for $\widehat{\tau}_x$. See Remark 9 ahead for alternative estimators for v_x . The following proposition shows that these estimators are indeed consistent under almost the same conditions as in Theorem 1.

Proposition 1 (Consistency of estimators for nuisance parameters). *Suppose that the conditions of Theorem 1 hold and in addition that $nh^5/\log n \rightarrow \infty$. Then we have $\widehat{\sigma}_x^2 \xrightarrow{\mathbb{P}} \sigma_x^2$, $\widehat{v}_x \xrightarrow{\mathbb{P}} v_x$, and $\widehat{s}_x(\widehat{\tau}_x) \xrightarrow{\mathbb{P}} s_x(\tau_x)$ as $n \rightarrow \infty$.*

Now, since Chernoff’s distribution is symmetric about the origin, an asymptotic $(1 - \alpha)$ -confidence interval for $m(x)$ is given by

$$\left[\widehat{m}(x) \pm \frac{\widehat{s}_x(\widehat{\tau}_x)(\widehat{\sigma}_x/\widehat{v}_x)^{2/3}}{(nh^2)^{1/3}} q_{1-\alpha/2} \right],$$

where $q_{1-\alpha/2}$ is the $(1 - \alpha/2)$ -quantile of Chernoff’s distribution. For example, Table 2 in [20] yields that $q_{0.975} \approx 0.998181$.

Remark 9 (Alternative estimators for v_x). Alternatively to the estimator \widehat{v}_x , we may use

$$\widetilde{v}_x = \frac{1}{2h^3} \left(\widehat{Q}_x(\widehat{\tau}_x + 2h) - \widehat{Q}_x(\widehat{\tau}_x - 2h) - 2\{\widehat{Q}_x(\widehat{\tau}_x + h) - \widehat{Q}_x(\widehat{\tau}_x - h)\} \right),$$

which is consistent under additional smoothness conditions on the conditional density.

Still, higher order numerical differentials tend to be unstable in the finite sample. Instead, we may use the expression $v_x = -f^{(2)}(m(x) | x)s(\tau_x)^4/2$, and estimate $f^{(2)}(m(x) | x)$ by a kernel method. Suppose that X is decomposed as $X = (X^C, X^D)$ where $X^C \in \mathbb{R}^{d_C}$ is continuous and $X^D \in \mathbb{R}^{d-D}$ is discrete. Let $\mathbb{K}_1 : \mathbb{R} \rightarrow \mathbb{R}$ and $\mathbb{K}_2 : \mathbb{R}^{d_C} \rightarrow \mathbb{R}$ be kernel functions (i.e., functions that integrate to 1) where \mathbb{K}_1 is twice differentiable. For given bandwidths $b_X = b_{X,n} \rightarrow 0$ and $b_Y = b_{Y,n} \rightarrow 0$, we may estimate $f^{(2)}(m(x) | x)$ with $x = (x^C, x^D)$ by

$$\begin{aligned} & \widehat{f}^{(2)}(\widehat{m}(x) | x) \\ &= \frac{(nb_Y^3 b_X^{d_C})^{-1} \sum_{i=1}^n \mathbb{K}_1''((\widehat{m}(x) - Y_i)/b_Y) \mathbb{K}_2((x^C - X_i^C)/b_X) I(X_i^D = x^D)}{(nb_X^{d_C})^{-1} \sum_{i=1}^n \mathbb{K}_2((x^C - X_i^C)/b_X) I(X_i^D = x^D)}, \end{aligned}$$

which is consistent under appropriate conditions. This leads to an alternative estimator for v_x :

$$\check{v}_x = -\widehat{f}^{(2)}(\widehat{m}(x) | x) \widehat{s}(\widehat{\tau}_x)^4/2. \tag{8}$$

In the simulation study, we use the kernel-based estimator \check{v}_x for v_x .

3.2.2. Subsampling

It is known that the nonparametric bootstrap in general fails to be consistent for $n^{1/3}$ -consistent estimators (cf. [1, 36, 31, 47]) and so it is unlikely that the bootstrap would be consistent for our estimator $\widehat{m}(x)$. Instead, since the limiting distribution is a scale transformation of Chernoff's distribution that is absolutely continuous, the subsampling provides a valid inference method for our estimator $\widehat{m}(x)$; see [41, 42]. Let $\widehat{m}(x) = \widehat{m}_n(x) = \widehat{m}_n(x; (Y_1, X_1), \dots, (Y_n, X_n))$ and $h = h_n$, and let W_1, \dots, W_N be the $N = \binom{n}{\ell}$ subsets of $\{(Y_1, X_1), \dots, (Y_n, X_n)\}$ of size $\ell (< n)$. Consider the subsampling distribution

$$U_{n,\ell}(x; t) = \frac{1}{N} \sum_{i=1}^N I \left\{ (\ell h_\ell^2)^{1/3} (\widehat{m}_\ell(x; W_i) - \widehat{m}_n(x)) \leq t \right\}. \quad (9)$$

Then, under the same conditions as in Theorem 1, we have

$$\sup_{t \in \mathbb{R}} \left| U_{n,\ell}(x; t) - \mathbb{P} \left(s_x(\tau_x) (\sigma_x/v_x)^{2/3} Z \leq t \right) \right| \xrightarrow{\mathbb{P}} 0,$$

provided that $\ell = \ell_n \rightarrow \infty$ and $\ell = o(n)$. Hence, denoting by $\widehat{q}_{n,\ell}(x; 1 - \alpha)$ the $(1 - \alpha)$ -quantile of $U_{n,\ell}(x; \cdot)$, i.e.,

$$\widehat{q}_{n,\ell}(x; 1 - \alpha) = \inf \{ t : U_{n,\ell}(x; t) \geq 1 - \alpha \},$$

an asymptotic $(1 - \alpha)$ -confidence interval for $m(x)$ is given by

$$\left[\widehat{m}_n(x) - \frac{\widehat{q}_{n,\ell}(x; 1 - \alpha/2)}{(nh_n^2)^{1/3}}, \widehat{m}_n(x) - \frac{\widehat{q}_{n,\ell}(x; \alpha/2)}{(nh_n^2)^{1/3}} \right].$$

Some comments on the subsampling confidence interval are in order.

Remark 10 (Comments on subsampling confidence interval). (i) In practice, $N = \binom{n}{\ell}$ is too large and so the computation of the complete average over $i = 1, \dots, N$ in (9) is too demanding. Instead, we can take the average of a randomly selected subset of $\{1, \dots, N\}$; see Corollary 2.4.1 in [42].

(ii) The bandwidth h_ℓ used in each subsample may be taken as h_n as long as $\ell_n h_n^8 \rightarrow 0$ and $\ell_n h_n^5 \rightarrow \infty$.

Remark 11 (Simultaneous confidence intervals over finite design points). Consider the setting of Corollary 1, and let $\nu_{1-\alpha}$ denote the $(1 - \alpha)$ quantile of $\max_{1 \leq j \leq L} |s_{x^j}(\tau^j) W_j|$. Then a simultaneous confidence interval for $m(x)$ over the design points x^1, \dots, x^L is given by

$$\left[\widehat{m}(x^j) \pm \frac{\nu_{1-\alpha}}{(nh^2)^{1/3}} \right], \quad j = 1, \dots, L.$$

In general the distribution of $(W_1, \dots, W_L)^T$ is complicated as it depends on whether there are ties in $\tau_{x^1}, \dots, \tau_{x^L}$, so analytical estimation of $\nu_{1-\alpha}$ is difficult. Instead, we can use the subsampling to estimate $\nu_{1-\alpha}$. The procedure is analogous to the pointwise case and hence omitted.

4. Numerical results

4.1. Bandwidth selection

The proposed estimator requires to choose the bandwidth h . We suggest here a simple method to choose the bandwidth, which is based on a modification to the bandwidth selection rule suggested in [30]. The baseline idea of our approach is to select the bandwidth in such a way that the sparsity function $s_x(\tau)$ is well estimated. A similar approach is used in [14] who adapt the smoothing bandwidth to kernel estimation of multi-modal regression by optimizing the conditional density estimation rate. The performance of the sparsity function estimate $\hat{s}_x(\tau)$ depends on the quantile τ of interest, and so the constant involved in the bandwidth should adapt to τ . Since we are interested in $s_x(\tau)$ around $\tau = \tau_x$, we aim at choosing h in such a way that $\hat{s}_x(\tau)$ around $\tau = \tau_x$ tends to be accurate but modify the rate of h so that it satisfies Condition (v) in Assumption 1.

For estimation of $s_x(\tau)$ based on quantile regression, [30] suggest to use the τ -dependent bandwidth

$$h^{KM}(\tau) = n^{-1/3} z_\alpha^{2/3} \left\{ 1.5 \frac{\phi(\Phi^{-1}(\tau))}{2\Phi^{-1}(\tau)^2 + 1} \right\}^{1/3},$$

where ϕ and Φ are the density and distribution functions of $N(0, 1)$, and $z_\alpha = \Phi^{-1}(1 - \alpha/2)$. We set $\alpha = 0.05$. The bandwidth $h^{KM}(\tau)$ does not satisfy Condition (v) in Assumption 1 and is τ -dependent, and so we shall modify $h^{KM}(\tau)$ as follows: (i) pick any design point x in the support of X ; (ii) use the pilot bandwidth $h^{\text{pilot}} = n^{1/6} h^{KM}(0.5) \propto n^{-1/6}$ to construct a preliminary estimator $\hat{\tau}_x^{\text{prelim}}$ of τ_x ; (iii) and use $h_n = h_{n,x} = n^{1/6} h^{KM}(\hat{\tau}_x^{\text{prelim}})$ to construct a final estimator $\hat{m}(x)$. The simulation results suggest that, although it would not be optimal, this bandwidth selection rule works reasonably well.

4.2. Simulation results

4.2.1. Comparison of RMSEs

We compare the performance of our estimator with that of the linear modal regression estimator of [25, 53] via the root mean square error (RMSE)

$$\sqrt{\mathbb{E}_{X^*} [\{\hat{m}(X^*) - m(X^*)\}^2]},$$

where $X^* \stackrel{d}{=} X$ is independent of the data and \mathbb{E}_{X^*} is the expectation with respect to X^* . We consider two settings: the first one is the case where the modal function is linear while the second one is the case where the modal function is nonlinear.

Case (i). Consider a linear location-scale model

$$Y = 1 + X_2 - 3X_3 + X_4 + X_2\nu,$$

where $X = (1, X_2, X_3, X_4)^T$, $X_2, X_3 \sim U(0, 1)$, $X_4 \sim N(0, 1)$, and $\nu \sim Ga(3, 0.5)$ (the Gamma distribution with shape parameter 3 and scale parameter 0.5). In this case, both the conditional quantile and modal functions are linear in X . In fact, $Q_\tau(X) = 1 + (1 + F^{-1}(\tau))X_2 - 3X_3 + X_4$, where F denotes the distribution function of ν . In addition, since the mode of $Ga(3, 0.5)$ is 1, the modal function is $m(X) = 1 + 2X_2 - 3X_3 + X_4$.

Case (ii). Consider the following data generating process

$$Y = U^3/3 - X_2(U - 1)^2,$$

where $X = (1, X_2)^T$, $X_2 \sim U(0, 1)$, and $U \sim U(0, 1)$ independent of X . In this case, the conditional quantile function is linear, $Q_\tau(X) = \tau^3/3 - X_2(\tau - 1)^2$, but the modal function is nonlinear, $m(X) = -2X_2^3/3 + 2X_2^2 - X_2$; see Remark 1.

In this simulation study, we choose $\varepsilon = 0.1$ and compute $\widehat{Q}_x(\tau)$ for 100 equally spaced grids on $[\tau_{\min}, \tau_{\max}] = [0.05, 0.95]$. To implement the linear modal regression estimator, we follow the EM algorithm and the bandwidth selection rule suggested in [53]. The number of Monte Carlo repetitions is 1000 for each case.

Figures 1 and 2 present the box plots of RMSEs of the linear modal regression and proposed estimators for Cases (i) and (ii), respectively, with $n = 500, 1000$, and 2000. These figures lead to the following observations. First, in both cases, the RMSE of the proposed estimator overall decreases as the sample size increases. Second, the proposed estimator tends to be more variable than the linear modal regression estimator, so that the interquartile range of the RMSE is wider for the proposed estimator than the linear modal regression estimator. Third, in both cases, the proposed estimator outperforms the linear modal regression estimator. The superior performance of the proposed estimator in Case (ii) is not surprising since the true modal function is nonlinear in that case and so the linear modal regression estimator is not consistent. Interestingly, even when the true modal function is linear (Case (i)), the proposed estimator performs substantially better than the linear modal regression estimator. This may be partly because the EM algorithm used to compute linear modal regression estimates failed to find global optimal solutions. Overall, the figures confirm that the proposed estimator works well in practice.

4.2.2. Coverage probabilities of confidence intervals

Next, we assess the performance of analytical and subsampling confidence intervals considered in Section 3.2. We follow the data generating process of Case (ii) and evaluate Monte Carlo average and median lengths, and coverage probabilities of confidence intervals at three design points $x_2 = 0.25, 0.5$, and 0.75.

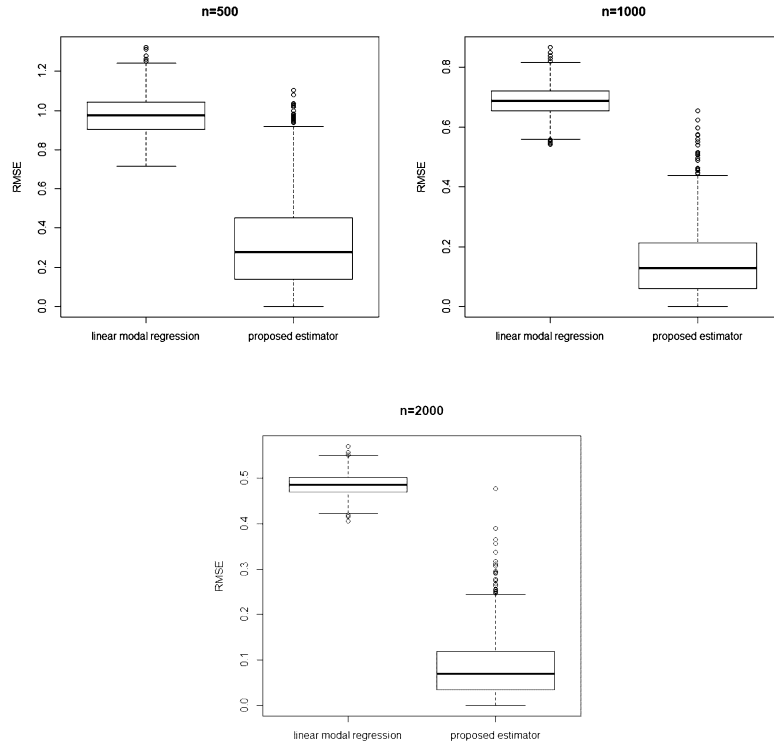


FIG 1. Box plots of RMSEs of the linear modal regression and proposed estimators for Case (i) with $n = 500$ (left), $n = 1000$ (middle), and $n = 2000$ (right).

We consider two nominal coverage probabilities of 99% and 95%. To implement the analytical confidence interval, we use the kernel-based estimator \hat{v}_x given in (8) for v_x . To construct \hat{v}_x , we use the Gaussian kernel for \mathbb{K}_1 and the Epanechnikov kernel for \mathbb{K}_2 together with bandwidths $b_Y = n^{-1/9}\hat{\sigma}_Y$ and $b_X = n^{-1/5}\hat{\sigma}_X$ where $\hat{\sigma}_Y$ and $\hat{\sigma}_X$ are the sample standard deviations of Y and X , respectively. To implement the subsampling confidence interval, we examine two subsample sizes: $\ell = 0.1n$ and $0.2n$. In this simulation study, instead of taking the average of whole subsamples in (9), we take the average of 250 randomly chosen subsamples. When applying the bandwidth selection rule to the subsample, we use the pilot bandwidth computed using the full sample.

Tables 1–4 present the simulation results on the confidence intervals. The tables show that both confidence intervals work reasonable well, given that the convergence rate of the estimator is relatively slow. It is worth noting that the estimators for the nuisance parameters $s_x(\tau_x)$ and v_x tend to be unstable, which results in the discrepancy between the average and median lengths of the analytical confidence interval. The subsample confidence interval is able to avoid estimation of those nuisance parameters, and so the length of the subsampling confidence interval tends to be shorter than that of the analytical confidence in-

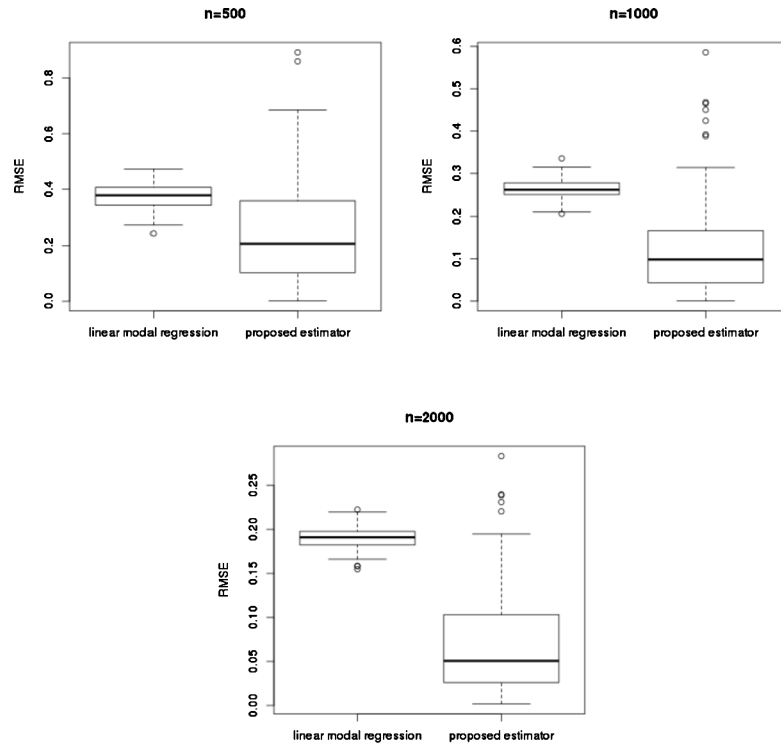


FIG 2. Box plots of RMSEs of the linear modal regression and proposed estimators for Case (ii) with $n = 500$ (left), $n = 1000$ (middle), and $n = 2000$ (right).

terval. In terms of the coverage probability, the subsampling confidence interval with subsample size $0.2n$ works the best.

4.3. Combined cycle power plant data

The electric energy output provided by a power plant fluctuates through the year because of several environmental conditions, and prediction of the electricity output given such environmental conditions is of interest. We apply the proposed estimator to predicting the net hourly electrical energy output using Combined Cycle Power Plant Data [24, 49]. The data set is taken from <https://archive.ics.uci.edu/ml/datasets/Combined+Cycle+Power+Plant> and consists of 9568 data points collected from a Combined Cycle Power Plant over 6 years (2006-2011). It contains hourly average ambient variables Temperature, Ambient Pressure, Relative Humidity, Exhaust Vacuum, and the net hourly electrical energy output, where the first four variables are regressors and the last variable is a response. For this data, the conditional distribution tends to be skewed, and therefore it would be natural to estimate the conditional mode. Figure 3 shows the estimate of the conditional density given one of the regres-

TABLE 1

Monte Carlo average and median lengths, and coverage probabilities of the 99% analytical confidence interval.

Design point	Sample size	Ave. length	Med. length	Cov. probability
$x_2 = 0.25$	$n = 500$	0.494	0.419	0.981
	$n = 1000$	0.359	0.315	0.986
	$n = 2000$	0.247	0.220	0.985
$x_2 = 0.50$	$n = 500$	0.715	0.599	1.000
	$n = 1000$	0.506	0.475	0.997
	$n = 2000$	0.392	0.380	0.992
$x_2 = 0.75$	$n = 500$	1.045	0.878	0.978
	$n = 1000$	0.724	0.653	0.977
	$n = 2000$	0.524	0.488	0.956

TABLE 2

Monte Carlo average and median lengths, and coverage probabilities of the 95% analytical confidence interval.

Design point	Sample size	Ave. length	Med. length	Cov. probability
$x_2 = 0.25$	$n = 500$	0.309	0.242	0.948
	$n = 1000$	0.207	0.175	0.941
	$n = 2000$	0.139	0.128	0.952
$x_2 = 0.50$	$n = 500$	0.459	0.343	0.987
	$n = 1000$	0.302	0.269	0.933
	$n = 2000$	0.226	0.221	0.894
$x_2 = 0.75$	$n = 500$	0.660	0.534	0.873
	$n = 1000$	0.429	0.371	0.869
	$n = 2000$	0.302	0.278	0.845

sors (Exhaust Vacuum). It is seen that the conditional density estimate is highly skewed and the pattern of the skewness depends on the value of the regressor.

To construct prediction intervals, we combine the proposed estimator with the split conformal prediction of [37]. Specifically:

1. Randomly split the index set $\{1, \dots, 9568\}$ into three parts $\mathcal{I}_1, \mathcal{I}_2$, and \mathcal{I}_3 .
2. Use the data $\{(Y_i, X_i) : i \in \mathcal{I}_1\}$ to construct the estimator $\widehat{m}(\cdot)$ for the modal function $m(\cdot)$.
3. Compute the $\alpha/2$ - and $(1 - \alpha/2)$ -quantiles of $\{Y_i - \widehat{m}(X_i) : i \in \mathcal{I}_2\}$ and they are denoted by $\widehat{\xi}_{\alpha/2}$ and $\widehat{\xi}_{1-\alpha/2}$, respectively. In this experiment, $\alpha = 0.05$ is used.
4. Construct $C_{\text{split}}(x) = [\widehat{m}(x) + \widehat{\xi}_{\alpha/2}, \widehat{m}(x) + \widehat{\xi}_{1-\alpha/2}]$.
5. Compute the empirical coverage probability:

$$\frac{1}{|\mathcal{I}_3|} \sum_{i \in \mathcal{I}_3} I\{Y_i \in C_{\text{split}}(X_i)\}.$$

In this experiment, we take $\mathcal{I}_1, \mathcal{I}_2$, and \mathcal{I}_3 in such a way that $|\mathcal{I}_1 \cup \mathcal{I}_2| : |\mathcal{I}_3| \approx 0.95 : 0.05$ and $|\mathcal{I}_1| : |\mathcal{I}_2| \approx 8 : 2$. We repeated this procedure 250 times and report the average of the empirical coverage probabilities together with the average and median lengths. In addition, we compare the proposed estimator with the linear modal regression estimator. Table 5 shows the results. For both methods, the empirical coverage probabilities are surprisingly close to the nominal

TABLE 3
Monte Carlo average and median lengths, and coverage probabilities of the 99% subsampling confidence interval.

Design point	Sample size	Subsample size	Ave. length	Med. length	Cov. probability
$x_2 = 0.25$	$n = 500$	$0.1n$	0.232	0.234	0.959
		$0.2n$	0.250	0.262	0.991
	$n = 1000$	$0.1n$	0.208	0.214	0.966
		$0.2n$	0.191	0.184	0.997
	$n = 2000$	$0.1n$	0.148	0.146	1.000
		$0.2n$	0.146	0.143	1.000
$x_2 = 0.50$	$n = 500$	$0.1n$	0.336	0.337	0.946
		$0.2n$	0.405	0.407	0.999
	$n = 1000$	$0.1n$	0.326	0.327	0.973
		$0.2n$	0.391	0.395	0.998
	$n = 2000$	$0.1n$	0.371	0.382	1.000
		$0.2n$	0.371	0.382	0.999
$x_2 = 0.75$	$n = 500$	$0.1n$	0.447	0.450	0.822
		$0.2n$	0.529	0.538	0.917
	$n = 1000$	$0.1n$	0.430	0.433	0.847
		$0.2n$	0.488	0.508	0.961
	$n = 2000$	$0.1n$	0.416	0.415	0.971
		$0.2n$	0.423	0.416	0.971

coverage probability of 95%, which is consistent with the theory developed in [37]. On the other hand, the average and median lengths of the conformal prediction band with the proposed estimator are substantially smaller than those with the linear modal regression estimator, which is an encouraging sign for the proposed estimator.

5. Discussion

In the present paper we have proposed a new estimator for the conditional mode based on quantile regression. The proposed estimate is computationally scalable since the quantile regression problem can be formulated as a linear programming problem. We have developed asymptotic distributional theory for the proposed estimator, which turns out to be nonstandard. Specifically, we have shown that the rate of convergence of the proposed estimator is $(nh^2)^{1/3}$ where $h = h_n \rightarrow 0$ is a sequence of bandwidths, and that the limiting distribution is a scale transformation of Chernoff's distribution. For inference, we have discussed analytical and subsampling confidence intervals. Finally we have verified the practical usefulness of the proposed method through numerical experiments.

In the present paper, we use the naive quantile regression estimator $\hat{\beta}(\tau)$ that is not smooth in τ to estimate the conditional quantile function, while the true slope vector $\beta(\tau)$ is smooth in τ under our assumption. An interesting alternative approach is to impose smoothness to $\hat{\beta}(\tau)$ so that the estimated conditional quantile function is differentiable in τ . We expect that the resulting conditional mode estimator would have a Gaussian limit (under regularity conditions), which is reminiscent of the smoothed maximum score estimator of [23]. Developing this alternative approach requires a whole new theory and is left as future research.

TABLE 4
 Monte Carlo average and median lengths, and coverage probabilities of the 95% subsampling confidence interval.

Design point	Sample size	Subsample size	Ave. length	Med. length	Cov. probability	
$x_2 = 0.25$	$n = 500$	$0.1n$	0.203	0.208	0.926	
		$0.2n$	0.198	0.195	0.982	
	$n = 1000$	$0.1n$	0.166	0.166	0.947	
		$0.2n$	0.148	0.145	0.993	
		$n = 2000$	$0.1n$	0.120	0.119	0.997
			$0.2n$	0.118	0.116	0.998
$x_2 = 0.50$	$n = 500$	$0.1n$	0.313	0.314	0.899	
		$0.2n$	0.374	0.380	0.989	
	$n = 1000$	$0.1n$	0.304	0.306	0.968	
		$0.2n$	0.353	0.366	0.997	
		$n = 2000$	$0.1n$	0.316	0.326	0.994
			$0.2n$	0.318	0.328	0.996
$x_2 = 0.75$	$n = 500$	$0.1n$	0.413	0.416	0.779	
		$0.2n$	0.473	0.490	0.887	
	$n = 1000$	$0.1n$	0.388	0.396	0.808	
		$0.2n$	0.412	0.415	0.937	
		$n = 2000$	$0.1n$	0.335	0.328	0.958
			$0.2n$	0.342	0.336	0.959

TABLE 5
 Monte Carlo average and median lengths, and empirical coverage probabilities of the 95% conformal prediction intervals.

Method	Average length	Median length	Coverage probability
Proposed method	19.01	19.02	0.950
Modal linear regression	23.71	23.32	0.950

Appendix A: Proofs

A.1. Preliminaries

In what follows, we will obey the following notation. For a given probability space (S, \mathcal{S}, Q) and a measurable function $f : S \rightarrow \mathbb{R}$, we use the notation $Qf = \int fdQ$ whenever the latter integral exists. For a class of measurable real-valued functions \mathcal{F} on S , let $N(\mathcal{F}, \|\cdot\|_{Q,2}, \delta)$ denote the δ -covering number for \mathcal{F} with respect to the $L^2(Q)$ -seminorm $\|\cdot\|_{Q,2}$; see Section 2.1 in [51] for details. In addition, for a (vector-valued) function g on a set T , we use the notation $\|g\|_T = \sup_{x \in T} \|g(x)\|$, where $\|\cdot\|$ denotes the Euclidean norm. We denote by $\stackrel{d}{=}$ the equality in distribution.

The following maximal inequality will be repeatedly used in the proof of Theorem 1.

Lemma 1 (A useful maximal inequality). *Let X_1, \dots, X_n be i.i.d. random variables taking values in a measurable space (S, \mathcal{S}) with common distribution P , and let \mathcal{F} be a pointwise measurable class of (measurable) real-valued functions on S with measurable envelope F .¹ Suppose that there exist constants $A \geq e$*

¹The class \mathcal{F} is said to be pointwise measurable if there exists a countable subclass $\mathcal{G} \subset \mathcal{F}$ such that for every $f \in \mathcal{F}$ there exists a sequence $g_m \in \mathcal{G}$ with $g_m \rightarrow f$ pointwise; see Section 2.3 in [51].

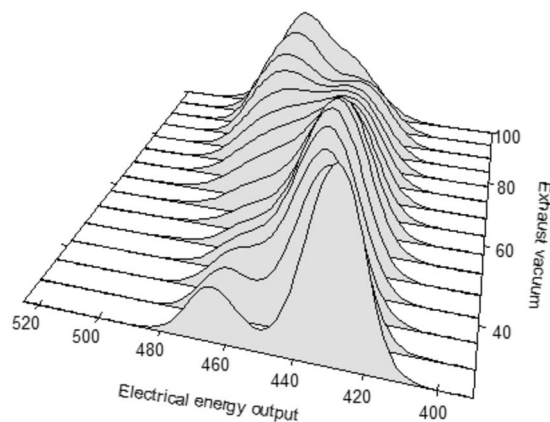


FIG 3. The conditional density estimate of the electronic energy output given Exhaust Vacuum.

and $V \geq 1$ such that $\sup_Q N(\mathcal{F}, \|\cdot\|_{Q,2}, \eta \|F\|_{Q,2}) \leq (A/\eta)^V$ for all $0 < \eta \leq 1$, where \sup_Q is taken over all finitely discrete distributions on S . Furthermore, suppose that $0 < PF^2 < \infty$, and let σ^2 be any positive constant such that $\sup_{f \in \mathcal{F}} Pf^2 \leq \sigma^2 \leq PF^2$. Finally, let $B = \sqrt{\mathbb{E}[\max_{1 \leq i \leq n} F^2(X_i)]}$. Then

$$\mathbb{E} \left[\left\| \sum_{i=1}^n \{f(X_i) - Pf\} \right\|_{\mathcal{F}} \right] \leq C \left[\sqrt{nV\sigma^2 \log(A\|F\|_{P,2}/\sigma)} + VB \log(A\|F\|_{P,2}/\sigma) \right],$$

where $\|\cdot\|_{\mathcal{F}} = \sup_{f \in \mathcal{F}} |\cdot|$ and C is a universal constant.

Proof. See Corollary 5.1 in [10]. \square

In particular, if we take $\sigma^2 = PF^2$, then using the inequality $B \leq \sqrt{n}\|F\|_{P,2}$, we also have

$$\mathbb{E} \left[\left\| \sum_{i=1}^n \{f(X_i) - Pf\} \right\|_{\mathcal{F}} \right] \leq 2C\sqrt{n}\|F\|_{P,2}V \log A. \quad (10)$$

The right hand side on (10) can be improved to $\|F\|_{P,2}\sqrt{V \log A}$ up to a universal constant (cf. Theorem 2.14.1 in [51]), but this does not matter to the proof of Theorem 1.

Lemma 2. For i.i.d. random variables ζ_1, ζ_2, \dots , $\mathbb{E}[\max_{1 \leq i \leq n} |\zeta_i|] = o(n)$ if and only if $\mathbb{E}[|\zeta_1|] < \infty$.

Proof. This is a well known result in probability theory, but we provide its proof for the sake of completeness. The “only if” direction is trivial, and so we prove the “if” direction. Suppose that $\mathbb{E}[|\zeta_1|] < \infty$. Then the strong law

of large numbers yields that $\max_{1 \leq i \leq n} |\zeta_i|/n \leq \sum_{i=1}^n |\zeta_i|/n \rightarrow \mathbb{E}[|\zeta_1|]$ almost surely, which also implies that $\max_{1 \leq i \leq n} |\zeta_i|/n \rightarrow 0$ almost surely (in general for a sequence of real numbers $\{a_i\}_{i=1}^\infty$, if $n^{-1} \sum_{i=1}^n a_i$ converges as $n \rightarrow \infty$, then $\max_{1 \leq i \leq n} |a_i| = o(n)$). The the desired result follows from the generalized dominated convergence theorem (cf. Problem 4.3.12 in [11]). \square

A.2. Proof of Theorem 1

The proof of Theorem 1 depends on the following Bahadur representation of the quantile regression estimator $\widehat{\beta}(\tau)$.

Lemma 3 (Bahadur representation of $\widehat{\beta}(\tau)$). *Under Assumption 1, we have*

$$\widehat{\beta}(\tau) - \beta(\tau) = J(\tau)^{-1} \left[\frac{1}{n} \sum_{i=1}^n \{\tau - I(Y_i \leq X_i^T \beta(\tau))\} X_i \right] + R_n(\tau),$$

where $J(\tau) = \mathbb{E}[f(X^T \beta(\tau) | X) X X^T]$ and $\|R_n\|_{[\varepsilon/2, 1-\varepsilon/2]} = o_{\mathbb{P}}(n^{-3/4} \log n)$. In addition,

$$\left\| \frac{1}{n} \sum_{i=1}^n \{\tau - I(Y_i \leq X_i^T \beta(\cdot))\} X_i \right\|_{[\varepsilon/2, 1-\varepsilon/2]} = O_{\mathbb{P}}(n^{-1/2}). \quad (11)$$

The conclusion of the lemma is partly known in the literature, but we include the proof of the lemma since we could not find a right reference that exactly establishes the conclusion of the lemma under our assumption. We defer the proof of this lemma after the proof of Theorem 1.

Proof of Theorem 1. We divide the proof into several steps.

Step 1. We first expand the objective function $\widehat{s}_x(\tau)$ using the Bahadur representation of $\widehat{\beta}(\tau)$. Let $F(y | X)$ denote the conditional distribution function of Y given X , and let $U_i = F(Y_i | X_i)$ for $i = 1, \dots, n$. The variable U_i follows the uniform distribution on $(0, 1)$ independent of X_i for each $i = 1, \dots, n$. Since

$$Y_i \leq X_i^T \beta(\tau) \Leftrightarrow U_i \leq \tau$$

under our assumption (recall that $X_i^T \beta(\tau)$ is the conditional τ -quantile of Y_i given X_i), we also have

$$\widehat{\beta}(\tau) - \beta(\tau) = J(\tau)^{-1} \left[\frac{1}{n} \sum_{i=1}^n \{\tau - I(U_i \leq \tau)\} X_i \right] + R_n(\tau). \quad (12)$$

Using the Bahadur representation (12) along with some calculations, we have that

$$\begin{aligned} \widehat{s}_x(\tau) &= s_{x,n}(\tau) + x^T J(\tau)^{-1} \left[\frac{1}{n} \sum_{i=1}^n \{1 - I(U_i \in (\tau - h, \tau + h]) / (2h)\} X_i \right] \\ &\quad + \underbrace{O_{\mathbb{P}}(n^{-1/2}) + o_{\mathbb{P}}(n^{-3/4} h^{-1} \log n)}_{=o_{\mathbb{P}}((nh^2)^{-2/3})}, \end{aligned}$$

where $s_{n,x} = \{Q_x(\tau+h) - Q_x(x-h)\}/(2h)$ and the $o_{\mathbb{P}}$ and $O_{\mathbb{P}}$ terms are uniform in $\tau \in [\varepsilon, 1 - \varepsilon]$.

Now, let $K(u) = I(u \in (-1, 1])/2$ and $K_h(u) = h^{-1}K(u/h)$. Define

$$g_{n,\tau}(U, X) = s_{x,n}(\tau) + x^T J(\tau)^{-1} X \{1 - K_h(U - \tau)\}.$$

Denoting by \mathbb{P}_n the empirical probability measure for $\{(U_i, X_i)\}_{i=1}^n$, we have

$$\widehat{s}_x(\tau) = \mathbb{P}_n g_{n,\tau} + o_{\mathbb{P}}((nh^2)^{-2/3}),$$

where the $o_{\mathbb{P}}$ term is uniform in $\tau \in [\varepsilon, 1 - \varepsilon]$, and so $\widehat{\tau}_x$ satisfies that

$$\mathbb{P}_n g_{n,\widehat{\tau}_x} \leq \inf_{\tau \in [\varepsilon, 1 - \varepsilon]} \mathbb{P}_n g_{n,\tau} + o_{\mathbb{P}}((nh^2)^{-2/3}). \quad (13)$$

In what follows, we denote by P the joint distribution of (U, X) .

Step 2. Next, we show consistency of $\widehat{\tau}_x$. To this end, consider the function class $\mathcal{G}_n = \{g_{n,\tau} : \tau \in [\varepsilon, 1 - \varepsilon]\}$. It is seen that there exists a constant C_1 (independent of n) such that $\sup_{\tau \in [\varepsilon, 1 - \varepsilon]} |g_{n,\tau}(U, X)| \leq C_1(1 + \|X\|/h) =: G_n(U, X)$. Then there exist constants A_1 and V_1 independent of n such that

$$\sup_Q N(\mathcal{G}_n, \|\cdot\|_{Q,2}, \eta \|G_n\|_{Q,2}) \leq (A_1/\eta)^{V_1}, \quad 0 < \forall \eta \leq 1,$$

where the \sup_Q is taken over all finitely discrete distributions on $(0, 1) \times \mathcal{X}$. This follows from a small modification to the proof of Lemma 3.1 in [17] and so we omit the detailed proof. In addition, it is seen that $\sup_{\tau \in [\varepsilon, 1 - \varepsilon]} P g_{n,\tau}^2 = O(h^{-1})$, $P G_n^2 = O(h^{-2})$, and $\mathbb{E}[\max_{1 \leq i \leq n} G_n^2(U_i, X_i)] = o(n^{1/2}h^{-2})$ by Lemma 2.

Now, applying the maximal inequality of Lemma 1, we have

$$\mathbb{E} [\|\mathbb{P}_n g_{n,\tau} - P g_{n,\tau}\|_{[\varepsilon, 1 - \varepsilon]}] = \underbrace{O((nh)^{-1/2} \sqrt{\log n})}_{=o(1)} + o(n^{-3/4}h^{-1} \log n), \quad (14)$$

which implies that $\|\mathbb{P}_n g_{n,\tau} - P g_{n,\tau}\|_{[\varepsilon, 1 - \varepsilon]} = o_{\mathbb{P}}(1)$ by Markov's inequality. Further, $P g_{n,\tau} = s_{x,n}(\tau) = s_x(\tau) + o(1)$ uniformly in $\tau \in [\varepsilon, 1 - \varepsilon]$ and $s_x(\tau)$ is uniquely minimized at $\tau = \tau_x$ by assumption. Hence, by Theorem 5.7 in [50], we have $\widehat{\tau}_x \xrightarrow{\mathbb{P}} \tau_x$.

Step 3. The aim of this step is to show that $\widehat{\tau}_x = \tau_x + O_{\mathbb{P}}((nh^2)^{-1/3})$. We divide this step into three sub-steps.

Step 3-(a). We begin with observing that, for any $\delta = \delta_n \rightarrow 0$, $P g_{n,\tau}$ can be expanded as

$$P g_{n,\tau} = s_{x,n}(\tau) = s_{x,n}(\tau_x) + s'_{n,x}(\tau_x)(\tau - \tau_x) + (s''_{x,n}(\tau_x)/2 + o(1))(\tau - \tau_x)^2$$

uniformly in $|\tau - \tau_x| < \delta$, and $s'_{n,x}(\tau_x) = \{Q'_x(\tau_x+h) - Q'_x(\tau_x-h)\}/(2h) = O(h^2)$, where we have used the fact that $Q''_x(\tau_x) = s'_x(\tau_x) = 0$ (recall that τ_x is a

minimizer of $s_x(\tau)$. Indeed, recalling that $Q_x(\tau)$ is four times continuously differentiable in τ , we have

$$Q'_x(\tau_x + h) = Q'_x(\tau_x) + \underbrace{Q''_x(\tau_x)}_{=0} h + \frac{Q'''_x(\tau_x)}{2} h^2 + O(h^3), \text{ and likewise}$$

$$Q'_x(\tau_x - h) = Q'_x(\tau_x) + \frac{Q'''_x(\tau_x)}{2} h^2 + O(h^3),$$

which implies that $\{Q'_x(\tau_x + h) - Q'_x(\tau_x - h)\}/(2h) = O(h^2)$. Since $h^2 = o((nh^2)^{-1/3})$, using the inequality $|ab| \leq (a^2 + b^2)/2$, we have

$$|s'_{n,x}(\tau_x)(\tau - \tau_x)| \leq o(1)(\tau - \tau_x)^2 + o((nh^2)^{-2/3}).$$

Further, $s''_{x,n}(\tau_x) = s''_x(\tau_x) + o(1)$, and so we have

$$P(g_{n,\tau} - g_{n,\tau_x}) = (v_x + o(1))(\tau - \tau_x)^2 + o((nh^2)^{-2/3}) \tag{15}$$

uniformly in $|\tau - \tau_x| < \delta$, where $v_x = s''_x(\tau_x)/2 > 0$.

Step 3-(b). Next, for given $\delta > 0$, consider the function class $\mathcal{G}_{n,\delta} = \{g_{n,\tau} - g_{n,\tau_x} : \tau \in [\varepsilon, 1 - \varepsilon], |\tau - \tau_x| < \delta\}$. It is seen that there exists a constant C_2 independent of n and δ such that, whenever $|\tau - \tau_x| < \delta$,

$$\begin{aligned} & |g_{n,\tau}(U, X) - g_{n,\tau_x}(U, X)| \\ & \leq C_2 [\{(1 + \|X\|/h)\delta + (\|X\|/h)\{I(|U - \tau_x + h| \leq \delta) + I(|U - \tau_x - h| \leq \delta)\} \\ & =: G_{n,\delta}(U, X). \end{aligned} \tag{16}$$

Then there exist constants A_2 and V_2 independent of n and δ such that

$$\sup_Q N(\mathcal{G}_{n,\delta}, \|\cdot\|_{Q,2}, \eta \|G_{n,\delta}\|_{Q,2}) \leq (A_2/\eta)^{V_2}, \quad 0 < \forall \eta \leq 1. \tag{17}$$

Again, this follows from a small modification to the proof of Lemma 3.1 in [17].

Step 3-(c). Finally, by consistency of $\widehat{\tau}_x$, there exists $\delta = \delta_n \rightarrow 0$ such that $\mathbb{P}(|\widehat{\tau}_x - \tau_x| < \delta_n) \rightarrow 1$. In view of the expansion (15), for sufficiently large n , we have

$$P(g_{n,\tau} - g_{n,\tau_x}) \geq v_x(\tau - \tau_x)^2/2 - o((nh^2)^{-2/3})$$

uniformly in $|\tau - \tau_x| < \delta$. Further, by the covering number estimate of Step 3-(b) together with the maximal inequality (10), we have

$$\mathbb{E} \left[\|\mathbb{P}_n g - P g\|_{\mathcal{G}_{n,\delta}} \right] = O(n^{-1/2} h^{-1} \delta^{1/2}),$$

where we have used the fact that $P G_{n,\delta}^2 = O(h^{-2}\delta)$. Now, a small modification to the proof of Theorem 3.2.5 in [51] shows that $|\widehat{\tau}_x - \tau_x| = O_{\mathbb{P}}(r_n^{-1})$, where r_n satisfies $r_n^2 h^{-1} r_n^{-1/2} = n^{1/2}$, i.e., $r_n = (nh^2)^{1/3}$. This completes Step 3.

Step 4. Let $a_n = (nh^2)^{1/3}$, and define

$$\check{g}_{n,t} = \begin{cases} n^{1/6} h^{4/3} (g_{n,\tau_x+t/a_n} - g_{n,\tau_x}) & \text{if } \tau_x + t/a_n \in [\varepsilon, 1 - \varepsilon] \\ 0 & \text{otherwise} \end{cases}.$$

Consider the empirical process

$$\mathbb{G}_n \check{g}_{n,t} := \sqrt{n}(\mathbb{P}_n \check{g}_{n,t} - P \check{g}_{n,t}), \quad t \in \mathbb{R}.$$

Recall that $\sigma_x^2 = \mathbb{E}[(x^T J(\tau_x)^{-1} X)^2]/2$. The aim of this step is to show weak convergence of the empirical process $\{\mathbb{G}_n \check{g}_{n,t} : t \in \mathbb{R}\}$ to $\{\sigma_x B(t) : t \in \mathbb{R}\}$ in $\ell_{\text{loc}}^\infty(\mathbb{R})$, where $\ell_{\text{loc}}^\infty(\mathbb{R})$ is the space of locally bounded functions on \mathbb{R} equipped with the metric $d(f, g) = \sum_{N=1}^\infty 2^{-N} (1 \wedge \|f - g\|_{[-N, N]})$; cf. Section 1.6 in [51]. This reduces to verifying (i) the finite dimensional convergence, i.e., for any $t_1, \dots, t_\ell \in \mathbb{R}$,

$$(\mathbb{G}_n \check{g}_{n,t_1}, \dots, \mathbb{G}_n \check{g}_{n,t_\ell}) \xrightarrow{d} (\sigma_x B(t_1), \dots, \sigma_x B(t_\ell));$$

and (ii) the asymptotic equicontinuity of the empirical process on $[-N, N]$ for each $N = 1, 2, \dots$, i.e., for any $\eta > 0$,

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} \mathbb{P} \left(\sup_{\substack{|t_1 - t_2| < \delta \\ t_1, t_2 \in [-N, N]}} |\mathbb{G}_n(\check{g}_{n,t_1} - \check{g}_{n,t_2})| > \eta \right) = 0. \tag{18}$$

To verify the finite dimensional convergence, we first compute the limit of the covariance of \check{g}_{n,t_1} and \check{g}_{n,t_2} for $t_1 \leq t_2$. To this end, let

$$\varphi_{n,t}(U, X) = n^{1/6} h^{4/3} x^T J(\tau_x)^{-1} X \{K_h(U - \tau_x) - K_h(U - \tau_x - t/a_n)\}.$$

Direct (but tedious) calculations show that $\text{Cov}_P(\check{g}_{n,t_1}, \check{g}_{n,t_2}) = P(\varphi_{n,t_1} \varphi_{n,t_2}) + o(1)$, where Cov_P denotes the covariance under P . Since X and U are independent, we focus on computing

$$\begin{aligned} & \mathbb{E}[\{K_h(U - \tau_x) - K_h(U - \tau_x - t_1/a_n)\} \{K_h(U - \tau_x) - K_h(U - \tau_x - t_2/a_n)\}] \\ &= \frac{1}{4h^2} \left(2h - |[(\tau_x + t_1/a_n) \pm h] \cap [\tau_x \pm h]| - |[(\tau_x + t_2/a_n) \pm h] \cap [\tau_x \pm h]| \right. \\ & \quad \left. + |[(\tau_x + t_1/a_n) \pm h] \cap [(\tau_x + t_2/a_n) \pm h]| \right), \end{aligned} \tag{19}$$

where $[a \pm b] = [a - b, a + b]$ and $|\cdot|$ denotes the Lebesgue measure. First, since $ha_n = (nh^5)^{1/3} \rightarrow \infty$, for sufficiently large n , we have

$$|[(\tau_x + t/a_n) \pm h] \cap [\tau_x \pm h]| = 2h - \frac{|t|}{a_n}.$$

Next, if $t_1 \leq t_2$, then for sufficiently large n , we have

$$|[(\tau_x + t_1/a_n) \pm h] \cap [(\tau_x + t_2/a_n) \pm h]| = 2h - \frac{t_2 - t_1}{a_n}.$$

Combining these estimates leads to

$$\begin{aligned} & 2h - |[(\tau_x + t_1/a_n) \pm h] \cap [\tau_x \pm h]| - |[(\tau_x + t_2/a_n) \pm h] \cap [\tau_x \pm h]| \\ & \quad + |[(\tau_x + t_1/a_n) \pm h] \cap [(\tau_x + t_2/a_n) \pm h]| \\ &= \begin{cases} \frac{2t_1}{a_n} & \text{if } 0 \leq t_1 \leq t_2 \\ \frac{-2t_2}{a_n} & \text{if } t_1 \leq t_2 \leq 0 \\ 0 & \text{if } t_1 \leq 0 \leq t_2 \end{cases} \end{aligned}$$

Since $a_n h^2 = n^{1/3} h^{8/3}$, we conclude that

$$\lim_{n \rightarrow \infty} \text{Cov}_P(\check{g}_{n,t_1}, \check{g}_{n,t_2}) = \sigma_x^2 \mathbb{E}[B(t_1)B(t_2)].$$

The rest is to verify the Lindeberg condition, and to this end it is enough to verify that for any $t \in \mathbb{R}$ and $\eta > 0$,

$$n^{1/3} h^{8/3} P G_{n,|t|/a_n}^2 I(n^{1/6} h^{4/3} G_{n,|t|/a_n} > \eta \sqrt{n}) \rightarrow 0,$$

where $G_{n,\delta}$ is given in (16). After a few more calculations, we see that the problem boils down to showing that

$$a_n \mathbb{E} \left[\|X\|^2 I(|U - \tau_x \pm h| \leq |t|/a_n) I(\|X\| > \eta n^{1/3} h^{-1/3}) \right] \rightarrow 0. \tag{20}$$

However, since X and U are independent, the left hand side on (20) is

$$\underbrace{a_n \mathbb{P}(|U - \tau_x \pm h| \leq |t|/a_n)}_{=O(1)} \underbrace{\mathbb{E} \left[\|X\|^2 I(\|X\| > \eta n^{1/3} h^{-1/3}) \right]}_{=o(1)} \rightarrow 0.$$

Therefore, we have proved the finite dimensional convergence.

To verify the asymptotic equicontinuity (18), consider the function class

$$\check{G}_{n,\delta} = \{\check{g}_{n,t_1} - \check{g}_{n,t_2} : |t_1 - t_2| < \delta, t_1, t_2 \in [-N, N]\}.$$

We will apply Lemma 1 to the function class $\check{G}_{n,\delta}$. First, an envelope function for $\check{G}_{n,\delta}$ is given by $\check{G}_n = 2n^{1/6} h^{4/3} G_{n,N/a_n}$. Observe that, using independence between U and X , we have $P \check{G}_n^2 = O(1)$ and

$$\mathbb{E} \left[\max_{1 \leq i \leq n} \check{G}_n^2(U_i, X_i) \right] \leq O(n^{1/3} h^{2/3}) \mathbb{E} \left[\max_{1 \leq i \leq n} \|X_i\|^2 \right] = o(n^{5/6} h^{3/2}) = o(n),$$

where we have used $\mathbb{E} [\max_{1 \leq i \leq n} \|X_i\|^2] = o(n^{1/2})$, which follows from Lemma 2.

Next, from the covering number estimate (17), there exist constants A_3 and V_3 independent of n and δ such that

$$\sup_Q N(\check{G}_{n,\delta}, \|\cdot\|_{Q,2}, \eta \| \check{G}_n \|_{Q,2}) \leq (A_3/\eta)^{V_3}, \quad 0 < \forall \eta \leq 1.$$

Finally, it is seen that there exists a constant C_3 independent of n such that

$$P(g_{n,\tau_1} - g_{n,\tau_2})^2 \leq C_3 |\tau_1 - \tau_2|/h^2, \quad \forall \tau_1, \tau_2 \in [\varepsilon, 1 - \varepsilon],$$

which implies that

$$P(\check{g}_{n,t_1} - \check{g}_{n,t_2})^2 \leq C_3 |t_1 - t_2|, \quad \forall t_1, t_2 \in [-N, N]$$

for sufficiently large n .

Therefore, applying Lemma 1 to the function class $\check{\mathcal{G}}_{n,\delta}$, we conclude that there exists a constant C_4 independent of n and δ such that

$$\mathbb{E} \left[\sup_{\substack{|t_1-t_2|<\delta \\ t_1, t_2 \in [-N, N]}} |\mathbb{G}_n(\check{g}_{n,t_1} - \check{g}_{n,t_2})| \right] \leq C_4 \sqrt{\delta \log(1/\delta)} + o(1) \log(1/\delta)$$

for sufficiently small δ , where the $o(1)$ term is independent of δ . This leads to the asymptotic equicontinuity (18) by Markov's inequality.

Step 5. We derive the limit distribution of $\hat{\tau}_x$ by applying Theorem 2.7 in [27]. The optimality condition (13) implies that the rescaled estimator $\hat{t} = (nh^2)^{1/3}(\hat{\tau}_x - \tau_x)$ satisfies

$$\sqrt{n}P_n(-\check{g}_{n,\hat{t}}) \geq \sup_{t \in \mathbb{R}} \sqrt{n}P_n(-\check{g}_{n,t}) - o_{\mathbb{P}}(1).$$

In view of the expansion (15), we have

$$\sqrt{n}P\check{g}_{n,t} = v_x t^2 + o(1)$$

locally uniformly in $t \in \mathbb{R}$, i.e., uniformly in $t \in [-N, N]$ for each $N = 1, 2, \dots$

From the weak convergence result of Step 4, together with the fact that $B \stackrel{d}{=} -B$, the non-centered empirical process $\{\sqrt{n}P_n(-\check{g}_{n,t}) : t \in \mathbb{R}\}$ converges weakly to the process $\{\sigma_x B(t) - v_x t^2 : t \in \mathbb{R}\}$ in $\ell_{\text{loc}}^\infty(\mathbb{R})$, and the limit process concentrates on $C_{\max}(\mathbb{R})$ (as defined in [27]) by Lemmas 2.5 and 2.6 in [27]. Further, $\hat{t} = O_{\mathbb{P}}(1)$ by Step 3. Therefore, by Theorem 2.7 in [27], we have

$$\hat{t} = (nh^2)^{1/3}(\hat{\tau}_x - \tau_x) \stackrel{d}{\rightarrow} \arg \max_{t \in \mathbb{R}} \{\sigma_x B(t) - v_x t^2\}.$$

The right hand side is equal in distribution to $(\sigma_x/v_x)^{2/3}Z$ by Problem 3.2.5 in [51], where $Z = \arg \max_{t \in \mathbb{R}} \{B(t) - t^2\}$. This leads to the first result (5) of the theorem.

Finally, observe that

$$\hat{m}(x) - m(x) = \hat{Q}_x(\hat{\tau}_x) - Q_x(\tau_x) = \hat{Q}_x(\hat{\tau}_x) - Q_x(\hat{\tau}_x) + Q_x(\hat{\tau}_x) - Q_x(\tau_x).$$

By Lemma 3,

$$|\hat{Q}_x(\hat{\tau}_x) - Q_x(\hat{\tau}_x)| \leq \|\hat{Q}_x - Q_x\|_{[\varepsilon, 1-\varepsilon]} \leq \|x\| \|\hat{\beta} - \beta\|_{[\varepsilon, 1-\varepsilon]} = O_{\mathbb{P}}(n^{-1/2}).$$

Applying the delta method, we have

$$\begin{aligned} & (nh^2)^{1/3}(\hat{m}(x) - m(x)) \\ &= (nh^2)^{1/3}(Q_x(\hat{\tau}_x) - Q_x(\tau_x)) + o_{\mathbb{P}}(1) \stackrel{d}{\rightarrow} s_x(\tau_x)(\sigma_x/v_x)^{2/3}Z. \end{aligned}$$

This completes the proof. \square

Proof of Lemma 3. The results (11) and $\|\widehat{\beta} - \beta\|_{[\varepsilon/2, 1-\varepsilon/2]} = O_{\mathbb{P}}(n^{-1/2})$ follow from Theorem 3 in [2]. By the first order condition for the quantile regression problem (2), we have

$$\left\| \sum_{i=1}^n \{\tau - I(Y_i \leq X_i^T \widehat{\beta}(\tau))\} X_i \right\| \leq \text{Card}(\{i \in \{1, \dots, n\} : Y_i = X_i^T \widehat{\beta}(\tau)\}) \max_{1 \leq i \leq n} \|X_i\|, \text{ and} \tag{21}$$

$$\sup_{\tau \in [\varepsilon/2, 1-\varepsilon/2]} \text{Card}(\{i \in \{1, \dots, n\} : Y_i = X_i^T \widehat{\beta}(\tau)\}) \leq d \text{ almost surely.} \tag{22}$$

The first result (21) follows from a modification to the proof of Lemma 2.1 in [15]; see Lemma 4 ahead. The second result (22) follows from the following observation. Pick any subset $I \subset \{1, \dots, n\}$ such that $\text{Card}(I) \geq d + 1$. Conditionally on $X_1^n = \{X_1, \dots, X_n\}$, consider the set

$$S_I = \{(X_i^T \beta)_{i \in I} : \beta \in \mathbb{R}^d\} \subset \mathbb{R}^{\text{Card}(I)},$$

which is a linear subspace of dimension at most d . If there exists $\tau \in [\varepsilon/2, 1-\varepsilon/2]$ such that $Y_i = X_i^T \widehat{\beta}(\tau)$ for all $i \in I$, then $(Y_i)_{i \in I} \in S_I$, so that

$$\begin{aligned} &\mathbb{P}(\text{ there exists } \tau \in [\varepsilon/2, 1-\varepsilon/2] \text{ such that } Y_i = X_i^T \widehat{\beta}(\tau) \text{ for all } i \in I \mid X_1^n) \\ &\leq \mathbb{P}((Y_i)_{i \in I} \in S_I \mid X_1^n). \end{aligned} \tag{23}$$

However, since the distribution of $(Y_i)_{i \in I}$ conditionally on X_1^n is absolutely continuous, the conditional probability on the right hand side is 0. By Fubini's theorem, the unconditional probability of the event inside the conditional probability on the left hand side of (23) is 0. Now,

$$\begin{aligned} &\mathbb{P}\left(\sup_{\tau \in [\varepsilon/2, 1-\varepsilon/2]} \text{Card}(\{i \in \{1, \dots, n\} : Y_i = X_i^T \widehat{\beta}(\tau)\}) \geq d + 1\right) \\ &\leq \sum_{\substack{I \subset \{1, \dots, n\} \\ \text{Card}(I) \geq d+1}} \mathbb{P}(\text{ there exists } \tau \in [\varepsilon/2, 1-\varepsilon/2] \text{ such that } Y_i = X_i^T \widehat{\beta}(\tau) \text{ for all } i \in I) = 0, \end{aligned}$$

which leads to the result (22).

Since $\mathbb{E}[\|X\|^4] < \infty$, we have $\max_{1 \leq i \leq n} \|X_i\| = o_{\mathbb{P}}(n^{1/4})$ (cf. Lemma 2), and so

$$\left\| \frac{1}{n} \sum_{i=1}^n \{\tau - I(Y_i \leq X_i^T \widehat{\beta}(\cdot))\} X_i \right\|_{[\varepsilon/2, 1-\varepsilon/2]} = o_{\mathbb{P}}(n^{-3/4}).$$

We will expand $n^{-1} \sum_{i=1}^n \{\tau - I(Y_i \leq X_i^T \widehat{\beta}(\tau))\} X_i$. Observe that

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \{\tau - I(Y_i \leq X_i^T \widehat{\beta}(\tau))\} X_i \\ &= \frac{1}{n} \sum_{i=1}^n \{\tau - I(Y_i \leq X_i^T \beta(\tau))\} X_i + \mathbb{E}[\{\tau - I(Y \leq X^T \beta)\} X]_{\beta=\widehat{\beta}(\tau)} \\ & \quad + \frac{1}{n} \sum_{i=1}^n \{I(Y_i \leq X_i^T \beta(\tau)) - I(Y_i \leq X_i^T \widehat{\beta}(\tau))\} X_i - \mathbb{E}[\{\tau - I(Y \leq X^T \beta)\} X]_{\beta=\widehat{\beta}(\tau)} \end{aligned}$$

The Taylor expansion yields that

$$\mathbb{E}[\{\tau - I(Y \leq X^T \beta)\} X]_{\beta=\widehat{\beta}(\tau)} = -J(\tau)(\widehat{\beta}(\tau) - \beta(\tau)) + O_{\mathbb{P}}(n^{-1})$$

uniformly in $\tau \in [\varepsilon/2, 1 - \varepsilon/2]$. It remains to show that

$$\begin{aligned} & \left\| \frac{1}{n} \sum_{i=1}^n \{I(Y_i \leq X_i^T \beta(\tau)) - I(Y_i \leq X_i^T \widehat{\beta}(\tau))\} X_i \right. \\ & \quad \left. - \mathbb{E}[\{\tau - I(Y \leq X^T \beta)\} X]_{\beta=\widehat{\beta}(\tau)} \right\|_{[\varepsilon/2, 1 - \varepsilon/2]} = o_{\mathbb{P}}(n^{-3/4} \log n). \end{aligned} \tag{24}$$

Since $\|\widehat{\beta} - \beta\|_{[\varepsilon/2, 1 - \varepsilon/2]} = O(n^{-1/2})$, for any $M_n \rightarrow \infty$ sufficiently slowly, $\mathbb{P}(\|\widehat{\beta} - \beta\|_{[\varepsilon/2, 1 - \varepsilon/2]} \leq M_n n^{-1/2}) \rightarrow 1$. Consider the function class

$$\begin{aligned} \mathcal{F}_n &= \left\{ (y, x) \mapsto \{I(y \leq x^T \beta) - I(y \leq x^T (\beta + \delta))\} \alpha^T x \right. \\ & \quad \left. : \beta \in \mathbb{R}^d, \|\delta\| \leq M_n n^{-1/2}, \alpha \in \mathbb{S}^{d-1} \right\}, \end{aligned}$$

where $\mathbb{S}^{d-1} = \{x \in \mathbb{R}^d : \|x\| = 1\}$. Then the left side on (24) is bounded by

$$\left\| \frac{1}{n} \sum_{i=1}^n f(Y_i, X_i) - \mathbb{E}[f(Y, X)] \right\|_{\mathcal{F}_n} \tag{25}$$

with probability approaching one. Since the function class $\{(y, x) \mapsto I(y \leq x^T \beta) \alpha^T x : \beta \in \mathbb{R}^d, \alpha \in \mathbb{S}^{d-1}\}$ (that is independent of n) is a VC subgraph class with envelope $F(y, x) = \|x\|$, there exist constants A and V independent of n such that

$$\sup_Q N(\mathcal{F}_n, \|\cdot\|_{Q,2}, \eta \|F\|_{Q,2}) \leq (A/\eta)^V, \quad 0 < \forall \eta \leq 1.$$

See Section 2.6 in [51]. Simple calculations show that

$$\begin{aligned} & \sup_{f \in \mathcal{F}_n} \mathbb{E}[f^2(Y, X)] = O(M_n n^{-1/2}) \quad \text{and} \\ & \mathbb{E} \left[\max_{1 \leq i \leq n} F^2(Y_i, X_i) \right] = \mathbb{E} \left[\max_{1 \leq i \leq n} \|X_i\|^2 \right] = o(n^{1/2}) \end{aligned}$$

by Lemma 2. Therefore, applying Lemma 1 to the function class \mathcal{F}_n shows that the expectation of the term (25) is bounded by

$$O(n^{-3/4}\sqrt{M_n \log n}) + o(n^{-3/4} \log n).$$

Choosing $M_n \rightarrow \infty$ sufficiently slowly, we obtain the desired result. \square

Lemma 4. Let $(y_1, x_1), \dots, (y_n, x_n) \in \mathbb{R} \times \mathbb{R}^d$ be pairs of outcome variables and regressors. Consider to solve the quantile regression problem:

$$\min_{\beta \in \mathbb{R}^d} \sum_{i=1}^n \rho_\tau(y_i - x_i^T \beta), \quad (26)$$

where $\tau \in (0, 1)$ is fixed. Let β^* be an optimal solution to (26) and let $I^* = \{i \in \{1, \dots, n\} : y_i = x_i^T \beta^*\}$. Then there exist $a_i \in [-1, 0]$ for $i \in I^*$ such that

$$\sum_{i=1}^n \{\tau - I(y_i \leq x_i^T \beta^*)\} x_i = \sum_{i \in I^*} a_i x_i.$$

Hence we have $\|\sum_{i=1}^n \{\tau - I(y_i \leq x_i^T \beta^*)\} x_i\| \leq \text{Card}(I^*) \max_{1 \leq i \leq n} \|x_i\|$.

Proof. Let $y = (y_1, \dots, y_n)^T$ and $\mathbb{X} = [x_1, \dots, x_n]^T$. The optimization problem (26) reduces to the following linear programming problem:

$$\begin{aligned} \min_{u, v \in \mathbb{R}^n, \beta \in \mathbb{R}^d} \quad & \tau 1_n^T u + (1 - \tau) 1_n^T v \\ \text{s.t.} \quad & u - v = y - \mathbb{X}\beta, \quad u \geq 0_n, \quad v \geq 0_n, \end{aligned} \quad (27)$$

where $1_n = (1, \dots, 1)^T \in \mathbb{R}^n$ and $0_n = (0, \dots, 0)^T \in \mathbb{R}^n$. The inequalities $u \geq 0_n$ and $v \geq 0_n$ are interpreted coordinatewise. Let $u_i^* = \max\{y_i - x_i^T \beta^*, 0\}$ and $v_i^* = \max\{-y_i + x_i^T \beta^*, 0\}$. Then $u^* - v^* = y - \mathbb{X}\beta^*$ and (u^*, v^*, β^*) is an optimal solution to the problem (27). Defining

$$\begin{aligned} f(u, v, \beta) &= \tau 1_n^T u + (1 - \tau) 1_n^T v, \\ g(u, v, \beta) &= (g_1(u, v, \beta), \dots, g_{2n}(u, v, \beta))^T = (-u^T, -v^T)^T, \\ h(u, v, \beta) &= (h_1(u, v, \beta), \dots, h_n(u, v, \beta))^T = u - v - y + \mathbb{X}\beta, \end{aligned}$$

the problem (27) can be written as

$$\begin{aligned} \min_{u, v \in \mathbb{R}^n, \beta \in \mathbb{R}^d} \quad & f(u, v, \beta) \\ \text{s.t.} \quad & g(u, v, \beta) \leq 0_{2n}, \quad h(u, v, \beta) = 0_n. \end{aligned}$$

Let $e_i \in \mathbb{R}^n$ denote the vector of which only the i -th element is 1 and the other elements are all zero. Then the gradient vectors of $f(u, v, \beta)$, $g_i(u, v, \beta)$, $g_{n+i}(u, v, \beta)$, and $h_i(u, v, \beta)$ are given by

$$\nabla f(u, v, \beta) = \begin{pmatrix} \tau 1_n \\ (1 - \tau) 1_n \\ 0_d \end{pmatrix}, \quad \nabla g_i(u, v, \beta) = \begin{pmatrix} -e_i \\ 0_n \\ 0_d \end{pmatrix},$$

$$\nabla g_{n+i}(u, v, \beta) = \begin{pmatrix} 0 \\ -e_i \\ 0_d \end{pmatrix}, \quad \nabla h_i(u, v, \beta) = \begin{pmatrix} e_i \\ -e_i \\ x_i \end{pmatrix}, \quad i = 1, \dots, n.$$

Since all the constraints are linear, by the Karush-Kuhn-Tucker theorem (cf. [4], Proposition 3.3.7), there exist $\mu_1, \dots, \mu_{2n} \geq 0$ and $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ such that

$$\begin{pmatrix} \tau 1_n \\ (1 - \tau) 1_n \\ 0_d \end{pmatrix} + \sum_{i=1}^n \mu_i \begin{pmatrix} -e_i \\ 0_n \\ 0_d \end{pmatrix} + \sum_{i=1}^n \mu_{n+i} \begin{pmatrix} 0_n \\ -e_i \\ 0_d \end{pmatrix} + \sum_{i=1}^n \lambda_i \begin{pmatrix} e_i \\ -e_i \\ x_i \end{pmatrix} = 0_{2n+d}, \tag{28}$$

$$\mu_i u_i^* = 0, \quad \text{and} \quad \mu_{n+i} v_i^* = 0, \quad i = 1, \dots, n. \tag{29}$$

Recall that $I^* = \{i \in \{1, \dots, n\} : y_i = x_i^T \beta^*\}$. Let $I_+^* = \{i \in \{1, \dots, n\} : y_i > x_i^T \beta^*\}$ and $I_-^* = \{i \in \{1, \dots, n\} : y_i < x_i^T \beta^*\}$. Observe that from the complementary slack condition (29),

$$\begin{aligned} i \in I_+^* &\Rightarrow u_i^* > 0 \Rightarrow \mu_i = 0 \Rightarrow \lambda_i = -\tau \quad \text{and} \\ i \in I_-^* &\Rightarrow v_i^* > 0 \Rightarrow \mu_{n+i} = 0 \Rightarrow \lambda_i = 1 - \tau. \end{aligned}$$

The last d equations in (28) imply that $\sum_{i=1}^n \lambda_i x_i = 0$, which can be rearranged as $\tau \sum_{i \in I_+^*} x_i + (\tau - 1) \sum_{i \in I_-^*} x_i = \sum_{i \in I^*} \lambda_i x_i$. The left hand side is

$$\sum_{i \in I_+^* \cup I_-^*} \{\tau - I(y_i \leq x_i^T \beta^*)\} x_i = \sum_{i=1}^n \{\tau - I(y_i \leq x_i^T \beta^*)\} x_i + (1 - \tau) \sum_{i \in I^*} x_i,$$

so that

$$\sum_{i=1}^n \{\tau - I(y_i \leq x_i^T \beta^*)\} x_i = \sum_{i \in I^*} \underbrace{(\lambda_i - 1 + \tau)}_{=a_i} x_i.$$

For $i \in I^*$, we have by the first $2n$ equations of (28),

$$\begin{aligned} \tau - \mu_i + \lambda_i = 0 &\Rightarrow \lambda_i \geq -\tau \quad \text{and} \\ 1 - \tau - \mu_{n+i} - \lambda_i = 0 &\Rightarrow \lambda_i \leq 1 - \tau, \end{aligned}$$

so that $\lambda_i \in [-\tau, 1 - \tau]$ for $i \in I^*$, i.e, $a_i \in [-1, 0]$ for $i \in I^*$. This completes the proof. \square

A.3. Proof of Corollary 1

The second result follows from the delta method (see the proof of Theorem 1), so we focus on proving the first result. We will follow the notation used in the proof of Theorem 1, but to make the dependence on x explicit, let us write $g_{n,x,\tau}(U, X) = s_{x,n}(\tau) + x^T J(\tau)^{-1} X \{1 - K_h(U - \tau)\}$,

$$\check{g}_{n,x,t} = \begin{cases} n^{1/6} h^{4/3} (g_{n,x,\tau_x+t/a_n} - g_{n,x,\tau_x}) & \text{if } \tau_x + t/a_n \in [\varepsilon, 1 - \varepsilon] \\ 0 & \text{otherwise} \end{cases},$$

and $\varphi_{n,x,t}(U, X) = n^{1/6}h^{4/3}x^T J(\tau_x)^{-1}X\{K_h(U - \tau_x) - K_h(U - \tau_x - t/a_n)\}$. Recall that $a_n = (nh^2)^{1/3}$.

We begin with observing that for $\hat{t}_j = (nh^2)^{1/3}(\hat{\tau}_{x^j} - \tau_{x^j})$, $j = 1, \dots, L$,

$$\begin{aligned} \sqrt{n}\mathbb{P}_n \left(-\sum_{j=1}^L \check{g}_{n,x^j, \hat{t}_j} \right) &\geq \sup_{(t_1, \dots, t_L)^T \in \mathbb{R}^L} \sqrt{n}\mathbb{P}_n \left(-\sum_{j=1}^L \check{g}_{n,x^j, t_j} \right) - o_{\mathbb{P}}(1), \quad \text{and} \\ \sqrt{n}P \left(\sum_{j=1}^L \check{g}_{n,x^j, t_j} \right) &= \sum_{j=1}^L v_{x^j} t_j^2 + o(1) \quad \text{locally uniformly in } (t_1, \dots, t_L)^T \in \mathbb{R}^L. \end{aligned}$$

In addition, from Theorem 1, we know that $\hat{t}_j = O_{\mathbb{P}}(1)$ for each $j = 1, \dots, L$. Hence, in view of Theorem 2.7 in [27], we only have to verify the following. Let $\ell_{\text{loc}}^{\infty}(\mathbb{R}^L)$ denote the space of all locally bounded functions on \mathbb{R}^L equipped with the metric $d(f, g) = \sum_{N=1}^{\infty} 2^{-N} (1 \wedge \|f - g\|_{[-N, N]^L})$. Recall that $\mathbb{G}_n g = \sqrt{n}(\mathbb{P}_n g - P g)$.

- (i) There exists a continuous version of \mathbb{B}_k for each $k = 1, \dots, M$, and the stochastic process $\{\mathbb{G}_n(\sum_{j=1}^L \check{g}_{n,x^j, t_j}) : (t_1, \dots, t_L)^T \in \mathbb{R}^L\}$ converges weakly to the process $\{\sum_{k=1}^M \mathbb{B}_k((t_j)_{j \in S_k}) : (t_1, \dots, t_L)^T \in \mathbb{R}^L\}$ in $\ell_{\text{loc}}^{\infty}(\mathbb{R}^L)$, where $\mathbb{B}_1, \dots, \mathbb{B}_M$ are independent.
- (ii) For each $k = 1, \dots, M$, the process

$$(t_j)_{j \in S_k} \mapsto \mathbb{B}_k((t_j)_{j \in S_k}) - \sum_{j \in S_k} v_{x^j} t_j^2$$

admits a unique maximizer almost surely.

The latter (ii) follows from Lemmas 2.5 and 2.6 in [27], so we focus on verifying the weak convergence (i). By Section 1.6 in [51], this boils down to verifying the finite dimensional convergence together with the asymptotic equicontinuity on each $[-N, N]^L$, i.e., for any $\eta > 0$,

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} \mathbb{P} \left(\sup_{\substack{|t_j - t'_j| < \delta \\ t_j, t'_j \in [-N, N], j=1, \dots, L}} \left| \mathbb{G}_n \left(\sum_{j=1}^L (\check{g}_{n,x^j, t_j} - \check{g}_{n,x^j, t'_j}) \right) \right| > \eta \right) = 0. \quad (30)$$

As we will see, the finite dimensional convergence and the asymptotic equicontinuity automatically imply the existence of a continuous version of \mathbb{B}_k for each $k = 1, \dots, M$.

The asymptotic equicontinuity (30) follows from the fact that $|\mathbb{G}_n(\sum_{j=1}^L (\check{g}_{n,x^j, t_j} - \check{g}_{n,x^j, t'_j}))| \leq \sum_{j=1}^L |\mathbb{G}_n(\check{g}_{n,x^j, t_j} - \check{g}_{n,x^j, t'_j})|$ and what we have proved in Step 4 in the proof of Theorem 1. It remains to prove the finite dimensional convergence. Direct calculations show that

$$\text{Cov}_P \left(\sum_{i=1}^L \check{g}_{n,x^i, t_i}, \sum_{j=1}^L \check{g}_{n,x^j, t'_j} \right) = \sum_{i,j=1}^L P(\varphi_{n,x^i, t_i} \varphi_{n,x^j, t'_j}) + o(1)$$

for any $(t_1, \dots, t_L)^T, (t'_1, \dots, t'_L)^T \in \mathbb{R}^L$. Consider first the case where $\tau_{x^i} = \tau_{x^j} = \tau_{(k)}$ for some $k = 1, \dots, M$. Then, from the calculation done in Step 4 in the proof of Theorem 1, we see that

$$\lim_{n \rightarrow \infty} P(\varphi_{n, x^i, t_i} \varphi_{n, x^j, t'_j}) = \frac{1}{2} (x^i)^T J(\tau_{(k)})^{-1} \mathbb{E}[X X^T] J(\tau_{(k)})^{-1} x^j \mathbb{E}[B(t_i) B(t'_j)].$$

Next, consider the case where $\tau_{x^i} \neq \tau_{x^j}$. Then, the intervals $[\tau_{x^i} \pm h]$ and $[(\tau_{x^i} + t_i/a_n) \pm h]$ have empty intersections with $[\tau_{x^j} \pm h]$ and $[(\tau_{x^j} + t_j/a_n) \pm h]$ for sufficiently large n , so that

$$\lim_{n \rightarrow \infty} P(\varphi_{n, x^i, t_i} \varphi_{n, x^j, t'_j}) = 0.$$

Conclude that

$$\begin{aligned} & \lim_{n \rightarrow \infty} \text{Cov}_P \left(\sum_{i=1}^L \check{g}_{n, x^i, t_i}, \sum_{j=1}^L \check{g}_{n, x^j, t'_j} \right) \\ &= \frac{1}{2} \sum_{k=1}^M \sum_{i, j \in S_k} (x^i)^T J(\tau_{(k)})^{-1} \mathbb{E}[X X^T] J(\tau_{(k)})^{-1} x^j \mathbb{E}[B(t_i) B(t'_j)]. \end{aligned} \quad (31)$$

The Lindeberg condition can be verified in a similar way to Step 4 in the proof of Theorem 1, so we have proved the finite dimensional convergence.

Now, for each $k = 1, \dots, M$, since $\mathbb{G}_n(\sum_{j=1}^L \check{g}_{n, x^j, t_j})|_{t_j=0, j \notin S_k} = \mathbb{G}_n(\sum_{j \in S_k} \check{g}_{n, x^j, t_j})$, we see that the process $(t_j)_{j \in S_k} \mapsto \mathbb{G}_n(\sum_{j \in S_k} \check{g}_{n, x^j, t_j})$ is asymptotically equicontinuous (with respect to the Euclidean metric) on $[-N, N]^{s_k}$ for each $N = 1, 2, \dots$ and the finite dimensional distributions converge weakly to those of \mathbb{B}_k . By the final paragraph in Section 1.6 of [51], the limit process (in $\ell_{\text{loc}}^\infty(\mathbb{R}^{s_k})$) is a version of \mathbb{B}_k with continuous paths.

We have already seen that the process $\{\mathbb{G}_n(\sum_{j=1}^L \check{g}_{n, x^j, t_j}) : (t_1, \dots, t_L)^T \in \mathbb{R}^L\}$ is weakly convergent in $\ell_{\text{loc}}^\infty(\mathbb{R}^L)$. The rest is to verify that the limit process is $\{\sum_{k=1}^M \mathbb{B}_k((t_j)_{j \in S_k}) : (t_1, \dots, t_L)^T \in \mathbb{R}^L\}$ where $\mathbb{B}_1, \dots, \mathbb{B}_M$ are independent, which however follows from the fact that the right hand side on (31) is identical to $\text{Cov}(\sum_{k=1}^M \mathbb{B}_k((t_i)_{i \in S_k}), \sum_{k=1}^M \mathbb{B}_k((t'_j)_{j \in S_k}))$. This completes the proof. \square

A.4. Proof of Proposition 1

The consistency of $\widehat{s}_x(\widehat{\tau}_x)$ follows from the uniform consistency of $\widehat{s}_x(\tau)$ on $[\varepsilon, 1 - \varepsilon]$, i.e., $\|\widehat{s}_x - s_x\|_{[\varepsilon, 1 - \varepsilon]} \xrightarrow{\mathbb{P}} 0$, which is established in Steps 1 and 2 in the proof of Theorem 1, together with the consistency of $\widehat{\tau}_x$. Next, $\widehat{\Sigma}$ is trivially consistent, and $\widehat{J}(\tau)$ is uniformly consistent on $[\varepsilon, 1 - \varepsilon]$ by Section A.4 in [2]. Together with the consistency of $\widehat{\tau}_x$ and continuity of the map $\tau \mapsto J(\tau)$, we obtain the consistency of $\widehat{\sigma}_x^2$. Finally, observe that $\Delta_h^3 \widehat{Q}_x(\tau) = \Delta_h^2 \widehat{s}_x(\tau)$, and $\widehat{s}_x(\tau) = \Delta_h Q_x(\tau) + O_{\mathbb{P}}((nh)^{-1/2} \sqrt{\log n})$ uniformly in $\tau \in [2\varepsilon/3, 1 - 2\varepsilon/3]$ by (14), so that $\Delta_h^3 \widehat{Q}_x(\tau) = \Delta_h^3 Q_x(\tau) + O_{\mathbb{P}}((nh^5)^{-1/2} \sqrt{\log n})$ uniformly in $\tau \in [\varepsilon, 1 - \varepsilon]$.

The consistency of \hat{v}_x then follows from the condition that $nh^5/\log n \rightarrow \infty$, continuity of the third derivative of $Q_x(\tau)$ at $\tau = \tau_x$, and the consistency of $\hat{\tau}_x$. This completes the proof. \square

Appendix B: Convergence of maximum of Chernoff random variables

In this appendix, we consider weak convergence of the maximum of independent Chernoff random variables. Let Z_1, \dots, Z_n be independent Chernoff random variables, and let $Z_{(n)} = \max_{1 \leq i \leq n} Z_i$ and $|Z|_{(n)} = \max_{1 \leq i \leq n} |Z_i|$. Chernoff's distribution is known to be absolutely continuous, and denote its density by f_Z . In addition, let F_Z denote the distribution function of Chernoff's distribution. An explicit form of f_Z is unknown, but by Corollary 3.4 of [19], the tail behavior of f_Z is given by

$$f_Z(z) \sim 2\lambda|z|e^{-\frac{2}{3}|z|^3 - \kappa|z|}, \quad |z| \rightarrow \infty, \quad (32)$$

where λ and κ are positive constants whose explicit values can be found in [19]. The precise meaning of (32) is that the ratio of the left and right hand sides approaches one as $|z| \rightarrow \infty$. This implies that

$$1 - F_Z(z) \sim \frac{\lambda}{z} e^{-\frac{2}{3}z^3 - \kappa z}, \quad z \rightarrow \infty. \quad (33)$$

Cf. Lemma 2.1 in [22]. The following lemma shows that both $Z_{(n)}$ and $|Z|_{(n)}$ converge in distribution to the Gumbel distribution as $n \rightarrow \infty$ after normalization. This lemma gives a supporting result for Remark 8, but is of independent interest. Recall that the (standard) Gumbel distribution is a distribution on \mathbb{R} with distribution function $\Lambda(z) = e^{-e^{-z}}$.

Lemma 5. *Let*

$$a_n = 3 \left(\frac{2}{3} \right)^{1/3} (\log n)^{2/3},$$

$$b_n = \left(\frac{3}{2} \log n \right)^{1/3} - \frac{1}{a_n} \left[\kappa \left(\frac{3}{2} \log n \right)^{1/3} + \frac{1}{3} \log \log n + \frac{1}{3} \log \frac{3}{2} - \log \lambda \right],$$

and define b'_n by replacing λ by 2λ in the definition of b_n . Then we have for any $z \in \mathbb{R}$,

$$\lim_{n \rightarrow \infty} \mathbb{P}(a_n(Z_{(n)} - b_n) \leq z) = e^{-e^{-z}} \quad \text{and} \quad \lim_{n \rightarrow \infty} \mathbb{P}(a_n(|Z|_{(n)} - b'_n) \leq z) = e^{-e^{-z}}.$$

We note that [22] already point out that Chernoff's distribution is in the domain of attraction of the Gumbel distribution (see [22] p.219), but they do not derive explicit norming constants.

The proof follows from the tail behavior of the Chernoff survival function (33) combined with the following lemma.

Lemma 6. Let $X_1, X_2, \dots \sim F$ i.i.d. for some distribution function F , and let $X_{(n)} = \max_{1 \leq i \leq n} X_i$. For a given constant $\tau \geq 0$ and a given sequence u_n , we have

$$n(1 - F(u_n)) \rightarrow \tau \Leftrightarrow \mathbb{P}(X_{(n)} \leq u_n) \rightarrow e^{-\tau}.$$

Proof of Lemma 6. See [33] Theorem 1.5.1. \square

Proof of Lemma 5. We first consider $Z_{(n)}$. Fix any $z \in \mathbb{R}$ and define u_n by $n(1 - F_Z(u_n)) = e^{-z}$. Then by the preceding lemma we have $\lim_{n \rightarrow \infty} \mathbb{P}(Z_{(n)} \leq u_n) = e^{-e^{-z}}$. We will find an explicit value of u_n . By (33), u_n satisfies

$$\frac{n\lambda}{u_n} e^{z - \frac{2}{3}u_n^3 - \kappa z u_n} \rightarrow 1.$$

Taking logarithms of both sides, we have

$$\log n + \log \lambda + z - \frac{2}{3}u_n^3 - \kappa z u_n - \log u_n = o(1). \quad (34)$$

Among the last three terms on the left hand side of (34), $\frac{2}{3}u_n^3$ is the dominant term, so that

$$\frac{\frac{2}{3}u_n^3}{\log n} \rightarrow 1. \quad (35)$$

Taking logarithms of both sides, we also have

$$\log u_n = \frac{1}{3} \left[\log \log n + \log \frac{3}{2} \right] + o(1).$$

Plugging this into (34), we have

$$\frac{2}{3}u_n^3 = \log n + z - \kappa z u_n - \frac{1}{3} \log \log n - \frac{1}{3} \log \frac{3}{2} + \log \lambda + o(1).$$

In addition, (35) also implies that

$$u_n = \left(\frac{3}{2} \log n \right)^{1/3} + \delta_n \quad \text{with } \delta_n = o((\log n)^{1/3}).$$

Plugging this into the preceding equation, using the identity $(a + b)^3 = a^3 + 3a^2b + 3ab^2 + b^3$, and comparing the orders, we see that $\delta_n = o(1)$. Conclude that

$$u_n^3 = \left(\frac{3}{2} \log n \right) \left[1 + \frac{z - \kappa \left(\frac{3}{2} \log n \right)^{1/3} - \frac{1}{3} \log \log n - \frac{1}{3} \log \frac{3}{2} + \log \lambda}{\log n} + o((\log n)^{-1}) \right].$$

Using $(1+x)^{1/3} = 1 + x/3 + O(x^2)$ as $x \rightarrow 0$, we have

$$u_n = \left(\frac{3}{2} \log n\right)^{1/3} \left[1 + \frac{z - \kappa \left(\frac{3}{2} \log n\right)^{1/3} - \frac{1}{3} \log \log n - \frac{1}{3} \log \frac{3}{2} + \log \lambda}{3 \log n} + o((\log n)^{-1}) \right] = a_n^{-1} z + b_n + o(a_n^{-1}).$$

Therefore, we have $\mathbb{P}(Z_{(n)} \leq u_n) = \mathbb{P}(a_n(Z_{(n)} - b_n) \leq z + o(1))$, which leads to the desired result for $Z_{(n)}$.

The proof for $|Z|_{(n)}$ is completely analogous, since by the symmetry of Chernoff's distribution, the distribution function G_Z of $|Z|$ is $G_Z(z) = 2F_Z(z) - 1$, so that $1 - G_Z(z) = 2(1 - F_Z(z))$. \square

Acknowledgments

The authors would like to thank the Editor Domenico Marinucci, an AE, and an anonymous referee for their careful review and constructive comments that helped improve on the quality of the paper.

References

- [1] J. Abrevaya and J. Huang. On the bootstrap of the maximum score estimator. *Econometrica*, 73(4):1175–1204, 2005. [MR2149245](#)
- [2] J. Angrist, V. Chernozhukov, and I. Fernández-Val. Quantile regression under misspecification, with an application to the US wage structure. *Econometrica*, 74(2):539–563, 2006. [MR2207400](#)
- [3] A. Belloni, V. Chernozhukov, and K. Kato. Valid post-selection inference in high-dimensional approximately sparse quantile regression models. *Journal of the American Statistical Association*, 2018. To appear. [MR3963177](#)
- [4] D. Bertsekas. *Nonlinear Programming (2nd Edition)*. Athena Scientific, 1999. [MR3444832](#)
- [5] J.E. Chacón. The modal age of statistics. arXiv:1807.02789, 2018.
- [6] P. Chaudhuri, K. Doksum, and A. Samarov. On average derivative quantile regression. *Annals of Statistics*, 25:715–744, 1997. [MR1439320](#)
- [7] Y.-C. Chen. Modal regression using kernel density estimation: A review. arXiv:1710.07004, 2017. [MR3826098](#)
- [8] Y.-C. Chen, C.R. Genovese, R.J. Tibshirani, and L. Wasserman. Non-parametric modal regression. *Annals of Statistics*, 44(2):489–514, 2016. [MR3476607](#)
- [9] H. Chernoff. Estimation of the mode. *Annals of the Institute of Statistical Mathematics*, 16(1):31–41, 1964. [MR0172382](#)
- [10] V. Chernozhukov, D. Chetverikov, and K. Kato. Gaussian approximation of suprema of empirical processes. *Annals of Statistics*, 42(4):1564–1597, 2014. [MR3262461](#)

- [11] R.M. Dudley. *Real Analysis and Probability*. Cambridge University Press, 2002. [MR1932358](#)
- [12] C. Durot, V.N. Kulikov, and H.P. Lopuhaä. The limit distribution of the L^∞ -error of Grenander type estimators. *Annals of Statistics*, 40:1578–1608, 2012. [MR3015036](#)
- [13] C. Durot and H.P. Lopuhaä. Limit theory in monotone function estimation. *Statistical Science*, 33:547–567, 2018. [MR3881208](#)
- [14] J. Einbeck and G. Tutz. Modelling beyond regression functions: an application of multimodal regression to speed–flow data. *Journal of the Royal Statistical Society Series C*, 55(4):461–475, 2006. [MR2242274](#)
- [15] R.A. El-Attar, M. Vidyasagar, and S.P.K. Dutta. An algorithm for l_1 -norm minimization with application to nonlinear l_1 -approximation. *SIAM Journal on Numerical Analysis*, 16(1):70–86, 1979. [MR0518685](#)
- [16] Y. Feng, J. Fan, and J.A.K. Suykens. A statistical learning approach to modal regression. arXiv:1702.05960, 2017.
- [17] S. Ghosal, A. Sen, and A.W. van der Vaart. Testing monotonicity of regression. *Annals of Statistics*, 28(4):1054–1082, 2000. [MR1810919](#)
- [18] U. Grenander. On the theory of mortality measurement: Part II. *Scandinavian Actuarial Journal*, 39:125–153, 1956. [MR0093415](#)
- [19] P. Groeneboom. Brownian motion with a parabolic drift and airy functions. *Probability Theory and Related Fields*, 81:79–109, 1989. [MR0981568](#)
- [20] P. Groeneboom and J.A. Wellner. Computing Chernoff’s distribution. *Journal of Computational and Graphical Statistics*, 10(2):388–400, 2001. [MR1939706](#)
- [21] C. Ho, P. Damien, and S. Walker. Bayesian mode regression using mixtures of triangular densities. *Journal of Econometrics*, 197(2):273–283, 2017. [MR3614201](#)
- [22] G. Hooghiemstra and H.P. Lopuhaä. An extremal limit theorem for the argmax process of Brownian motion minus a parabolic drift. *Extreme*, 1:215–240, 1998. [MR1814624](#)
- [23] J.L. Horowitz. A smoothed maximum score estimator for the binary response model. *Econometrica*, 60:505–531, 1992. [MR1162997](#)
- [24] H. Kaya, P. Tufekci, and S.F. Gurgun. Local and global learning methods for predicting power of a combined gas & steam turbine. *Proceedings of the International Conference on Emerging Trends in Computer and Electronics Engineering ICETCEE 2012*, pages 13–18, 2012.
- [25] G.C. Kemp and J. Santos-Silva. Regression towards the mode. *Journal of Econometrics*, 170(1):92–101, 2012. [MR2955942](#)
- [26] S. Khardani and A.F. Yao. Non linear parametric mode regression. *Communications in Statistics-Theory and Methods*, 46(6):3006–3024, 2017. [MR3579782](#)
- [27] J. Kim and D. Pollard. Cube root asymptotics. *Annals of Statistics*, 18(1):191–219, 1990. [MR1041391](#)
- [28] R. Koenker. *Quantile Regression*. Cambridge University Press, 2005. [MR2268657](#)
- [29] R. Koenker and G. Bassett. Regression quantiles. *Econometrica*, 46(1):33–

- 50, 1978. [MR0474644](#)
- [30] R. Koenker and J.A.F. Machado. Goodness of fit and related inference processes for quantile regression. *Journal of the American Statistical Association*, 94(448):1296–1310, 1999. [MR1731491](#)
- [31] M. Kosorok. Bootstrapping the Grenander estimator. In N. Balakrishnan, E. Pena, and M. Silvapulle, editors, *Beyond Parametrics in Interdisciplinary Research: Festschrift in Honour of Professor Pranab K. Sen*, pages 282–292. IMS, 2008. [MR2462212](#)
- [32] J.M. Krief. Semi-linear mode regression. *Econometrics Journal*, 20(2):149–167, 2017. [MR3685645](#)
- [33] M.R. Leadbetter, G. Lindgren, and H. Rootzén. *Extremes and Related Properties of Random Sequences and Processes*. Springer, 1983. [MR0691492](#)
- [34] M.-J. Lee. Mode regression. *Journal of Econometrics*, 42(3):337–349, 1989. [MR1040748](#)
- [35] M.-J. Lee. Quadratic mode regression. *Journal of Econometrics*, 57(1-3):1–19, 1993. [MR1237497](#)
- [36] C. Léger and B. MacGibbon. On the bootstrap in cube root asymptotics. *Canadian Journal of Statistics*, 34(1):29–44, 2006. [MR2267708](#)
- [37] J. Lei, M. G’Sell, A. Rinaldo, R. Tibshirani, and L. Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113:1094–1111, 2018. [MR3862342](#)
- [38] S. Ma and X. He. Inference for single-index quantile regression models with profile optimization. *Annals of Statistics*, 44:1234–1268, 2016. [MR3485959](#)
- [39] C.F. Manski. Maximal score estimation of the stochastic utility model of choice. *Journal of Econometrics*, 27(3):205–228, 1975. [MR0436905](#)
- [40] E. Parzen. On estimation of a probability density function and mode. *Annals of Mathematical Statistics*, 33(3):1065–1076, 1962. [MR0143282](#)
- [41] D.N. Politis and J.P. Romano. Large sample confidence regions based on subsamples under minimal conditions. *Annals of Statistics*, 22(4):2031–2050, 1994. [MR1329181](#)
- [42] D.N. Politis, J.P. Romano, and M. Wolf. *Subsampling*. Springer, 1999. [MR1707286](#)
- [43] J.L. Powell. Censored regression quantiles. *Journal of Econometrics*, 32(1):143–155, 1986. [MR0853049](#)
- [44] J.P. Romano. On weak convergence and optimality of kernel density estimates of the mode. *Annals of Statistics*, 16(2):629–647, 1988. [MR0947566](#)
- [45] T.W. Sager and R.A. Thisted. Maximum likelihood estimation of isotonic modal regression. *Annals of Statistics*, 10(3):690–707, 1982. [MR0663426](#)
- [46] H. Sasaki, Y. Ono, and M. Sugiyama. Modal regression via direct log-density derivative estimation. In *International Conference on Neural Information Processing*, pages 108–116, 2016. [MR3867784](#)
- [47] B. Sen, M. Banerjee, and M. Woodroffe. Inconsistency of bootstrap: The Grenander estimator. *Annals of Statistics*, 38(4):1953–1977, 2010. [MR2676880](#)
- [48] M.H. Seo and T. Otsu. Local M-estimation with discontinuous criterion for dependent and limited observations. *Annals of Statistics*, 46(1):344–369,

2018. [MR3766955](#)
- [49] P. Tufekci. Prediction of full load electrical power output of a base load operated combined cycle power plant using machine learning methods. *International Journal of Electrical Power & Energy Systems*, 60:126–140, 2014.
- [50] A.W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 2000. [MR1652247](#)
- [51] A.W. van der Vaart and J.A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer, 1996. [MR1385671](#)
- [52] T.Z. Wu, K. Yu, and Y. Yu. Singe-index quantile regression. *Journal of Multivariate Analysis*, 101:1607–1621, 2010. [MR2610735](#)
- [53] W. Yao and L. Li. New regression model: modal linear regression. *Scandinavian Journal of Statistics*, 41(3):656–671, 2014. [MR3249422](#)
- [54] W. Yao, B.G. Lindsay, and R. Li. Local modal regression. *Journal of Nonparametric Statistics*, 24(3):647–663, 2012. [MR2968894](#)
- [55] H. Zhou and X. Huang. Nonparametric modal regression in the presence of measurement error. *Electronic Journal of Statistics*, 10(2):3579–3620, 2016. [MR3575565](#)