Published in final edited form as:

J R Stat Soc Series B Stat Methodol. 2018 March ; 80(2): 433–452. doi:10.1111/rssb.12258.

Testing for Marginal Linear Effects in Quantile Regression

Huixia Judy Wang,

Associate Professor, Department of Statistics, George Washington University, Washington, District of Columbia 20052, USA

Ian W. McKeague, and

Professor, Department of Biostatistics, Columbia University, New York, NY 20032, USA

Min Qian

Assistant Professor, Department of Biostatistics, Columbia University, New York, NY 20032, USA

Abstract

This paper develops a new marginal testing procedure to detect the presence of significant predictors associated with the conditional quantiles of a scalar response. The idea is to fit the marginal quantile regression on each predictor one at a time, and then base the test on the t -statistics associated with the most predictive predictors. A resampling method is devised to calibrate this test statistic, which has non-regular limiting behavior due to the selection of the most predictive variables. Asymptotic validity of the procedure is established in a general quantile regression setting in which the marginal quantile regression models can be misspecified. Even though a fixed dimension is assumed to derive the asymptotic results, the proposed test is applicable and computationally feasible for large-dimensional predictors. The method is more flexible than existing marginal screening test methods based on mean regression, and has the added advantage of being robust against outliers in the response. The approach is illustrated using an application to an HIV drug resistance dataset.

Keywords

Bootstrap calibration; inference; marginal regression; non-standard asymptotics; quantile regression

1 Introduction

Consider a scalar response Y and a p -dimensional predictor $\mathbf{X} = (X_1, \dots, X_p)^T$. We are interested in testing the presence of significant predictors that affect the conditional quantile of Y at a given quantile level τ or across multiple quantiles. To answer this question, we develop a new inference procedure based on marginal linear quantile regression. For

Correspondence to: Min Qian.

Supplementary Material

The supplementary material available online contains the proofs of Lemmas 1–4, Corollary 1 and Lemmas 5–6, and some additional simulation results.

theoretical development, throughout we assume that p is finite, though the proposed test is applicable to cases with a large number of predictors.

Quantile regression has drawn increasing attention in recent years, mainly due to the following attractive features: (1) robustness against outliers in the response, especially in the case of median regression; (2) ability to capture heterogeneity in the set of important predictors at different quantile levels of the response distribution caused by, for instance, heteroscedastic variance.

To test for the presence of significant predictors at the π th conditional quantile of Y , one natural approach is to carry out a hypothesis test comparing the null model of no predictors and the full model consisting of all p predictors, by using either Wald-type, quasilielihood-ratio tests or bootstrap methods; see related discussions in Koenker (2005, Chapter 3), Kocherginsky et al. (2005), and Feng et al. (2011). Unfortunately, these omnibus-type tests are based on fitting a model with all p predictors, which quickly becomes prohibitive and leads to less powerful tests for large p .

Recently McKeague and Qian (2015) proposed an adaptive resampling test, which provides valid post-selection inference for detecting the presence of significant predictors based on marginal linear regression. As in correlation learning or sure independence screening, the idea of marginal regression is to regress the response on each predictor one at a time (Fan and Lv, 2008; Genovese et al., 2012). The adaptive resampling test successfully controls the family-wise error rate by accounting for the variation caused by variable selection.

Partly inspired by this work, here we propose a new method for testing the existence of marginal effects in quantile regression. The new approach allows the most active predictors to vary at different quantile levels of the response distribution, so it is more flexible than mean regression, and has the added advantage of being robust against outliers in the response. The proposed test is based on a maximum-type t -statistic associated with the selected most informative predictor at the quantile level of interest. Valid statistical inference has to take into account the uncertainty involved in the selection step; see related discussion in Leeb and Pötscher (2006) and Belloni et al. (2014). We show that the limiting distribution of the proposed post-selected statistic changes abruptly in the proximity of the null hypothesis of no effect. To adapt to the non-regular asymptotic behavior of the test statistic caused by variable selection, we develop a modified bootstrap procedure using a “pre-test” through thresholding. The proposed test is the first inference tool developed for quantile regression that scales in a computationally practical way with dimension. Unlike the omnibus-type tests that require fitting a model with all p predictors, the proposed method is based on fitting p linear quantile regression models with a single predictor, so its computational cost grows only linearly with p . Even though we assume a fixed dimension to derive the asymptotic results, there is numerical evidence that the proposed test continues to work in cases with $p > n$, but a thorough investigation for the high dimensional case remains open.

This article makes a novel contribution that is distinct from McKeague and Qian (2015). First of all, the test in McKeague and Qian (2015) is based on selecting the predictor that is

maximally correlated with the response Y . However, correlation is not useful in the quantile regression setup. Instead, we propose a new selection rule, which selects the most informative predictor at the quantile of interest as the one that minimizes an empirical asymmetric L_1 -loss function. Secondly, our proposed test uses a scale-invariant t -statistic, while that in McKeague and Qian (2015) is based on the maximum slope estimator. In quantile regression, the scale of the slope estimator depends not only on covariates but also on the quantile level. Therefore, a scale-invariant statistic is more desirable for multiple-quantile analysis, since pre-standardization of covariates is not sufficient to make the type of statistic as in McKeague and Qian (2015) scale free. Thirdly, unlike least squares regression, quantile regression allows the analysis over a set of quantile levels to capture the population heterogeneity. The flexibility of such globally concerned quantile regression has been discussed in Zheng et al. (2015). With the established convergence results across quantiles, our developed method can be used to detect the significance of predictors not only at a single quantile level but also at multiple quantiles jointly.

For quantile regression, developing asymptotic theory to calibrate the test statistic is enormously more challenging due to three main obstacles. Firstly, unlike in mean regression, the quantile estimator $\hat{\theta}_n(\tau)$ of the slope parameter for the selected predictor has no explicit form. Secondly, the asymmetric L_1 -loss function is not differentiable everywhere. Lastly, since each marginal quantile regression model is (possibly) misspecified, there exists a prediction bias caused by omitting the correct predictors, and the bias takes a complicated form in the quantile regression setting.

To overcome these challenges and to study the non-regular limiting behavior of $\hat{\theta}_n(\tau)$, we consider a general quantile regression model indexed by a (unidentifiable) local parameter that represents uncertainty at the $\sqrt{n}$ -scale close to the null hypothesis. We establish the asymptotic properties of the t -statistic $T_n(\tau) = n^{1/2} \hat{\theta}_n(\tau) / \hat{\sigma}_n(\tau)$, and obtain a quadratic expansion for the non-differentiable loss function using empirical process tools. In addition, we assess the bias of the marginal quantile regression estimator due to model misspecification under the local model by adapting the results in Angrist et al. (2006). Based on the developed asymptotic theory, we devise a nonparametric bootstrap procedure that adapts to the non-regular asymptotic behavior of $T_n(\tau)$ using a “pre-test” that involves thresholding. We establish the bootstrap consistency of this procedure under the general local quantile regression model.

The current paper is closely related to the post-selection inference literature. In particular, one can view the determination of the most informative predictor as a selection step, and the inference on the slope associated with the selected predictor in marginal quantile regression as the inference after selection. Lee and Taylor (2014) developed a post-selection inference method for marginal screening in linear regression, which selects the top k predictors most correlated with Y . Our proposed test solves a similar problem for the special case of $k = 1$ in quantile regression setup. On the other hand, the method in Lee and Taylor (2014) relies on a strong model assumption that the regression errors are normally distributed with constant variance, and thus cannot be applied to quantile regression. Among very few related works in the quantile regression literature, Belloni et al. (2015) and Belloni et al. (2014) proposed post-selection inference methods for median and quantile regression models, respectively.

However, inference in these papers is focused on the slope parameter of a single pre-specified predictor (such as a treatment indicator), while the remaining predictors are selected through L_1 -penalization. Thus their approach does not apply to marginal screening-type test in which no predictor is singled out *a priori*.

Recently, Zhang et al. (2017) proposed a test based on the martingale difference divergence to detect the dependence of the conditional quantile of the response on covariates in a model-free setting. The current paper differs from Zhang et al. (2017) from several perspectives. When the interest is on the dependence at a single quantile level τ , our proposed method does not require any model assumptions on the quantiles near τ , while the test calibration in Zhang et al. (2017) requires a stronger local quantile independence assumption. In addition, our proposed method can be used to conduct joint tests across multiple quantiles, and identify the most predictive variables, while that in Zhang et al. (2017) can only be applied to assess the overall dependence at a given single quantile level.

The rest of the paper is organized as follows. In Section 2, we formulate the problem, establish the non-regular local asymptotic distribution of $T_n(\tau)$, and develop the proposed test procedure. We assess the performance of the approach through a simulation study in Section 3, and apply it to an HIV drug susceptibility dataset in Section 4. Some concluding remarks are made in Section 5. Proofs are collected in the Appendix and the online Supplementary Material. The developed R program and the HIV data are available at <http://blogs.gwu.edu/judywang/software/QMET/>.

2 Proposed Method

2.1 Marginal Quantile Regression

Suppose that $\{(y_i, \mathbf{x}_i), i = 1, \dots, n\}$ is a random sample of $(Y, \mathbf{X})$, where $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^T$. Let $\mathcal{T} \subset (0, 1)$ be a set with L prespecified quantile levels, where L is finite. We assume the following linear quantile regression model:

$$Q_\tau(Y|\mathbf{X}) = \alpha_0(\tau) + \mathbf{X}^T \boldsymbol{\beta}_0(\tau), \tau \in \mathcal{T}, \quad (1)$$

where $\alpha_0(\tau)$ and $\boldsymbol{\beta}_0(\tau) \in \mathbb{R}^p$ are the unknown quantile coefficients, $Q_\tau(Y|\mathbf{X}) = \inf\{y : F_Y(y|\mathbf{X}) \geq \tau\}$ is the τ th conditional quantile of Y given $\mathbf{X}$, and $F_Y(\cdot|\mathbf{X})$ is the distribution function of Y given $\mathbf{X}$. Model (1) only assumes that the conditional quantile of Y is linear in $\mathbf{X}$ at quantiles of interest, and thus is broader than the linear regression model as in McKeague and Qian (2015). Unlike least squares regression, quantile regression analysis allows one to study at multiple quantiles. We aim to develop a formal test of whether any component of $\mathbf{X}$ has an effect on either a given quantile or at multiple quantiles of Y . Throughout we assume that all predictors are standardized before the data analysis. For notational simplicity we will first present the proposed test at a single quantile level $\tau \in \mathcal{T}$, and discuss testing across multiple quantiles in Section 2.5.

Our proposed test is based on fitting the working marginal quantile regression models by regressing Y on X_k , $k = 1, \dots, p$, for each k separately. That is, for each k , the working marginal quantile regression solves the population minimization problem to obtain

$$(\alpha_k(\tau), \theta_k(\tau)) = \underset{\alpha, \theta}{\operatorname{argmin}} E[\rho_\tau(Y - \alpha - \theta X_k) - \rho_\tau(Y)], \quad (2)$$

where $\rho_\tau(u) = \{\tau - \mathbb{I}(u < 0)\}u$ is the quantile loss function (Koenker, 2005), and here we (implicitly) assume integrability. The coefficient $\theta_k(\tau)$ approximates the linear effect of X_k on the τ th quantile of Y , and we refer to it as the *quantile marginal linear effect* of X_k hereafter. However, it is worth noting that the marginal quantile regression models are in general misspecified, and consequently $\alpha_k(\tau) + \theta_k(\tau)X_k$ may differ from $Q_\tau(Y/X_k)$, the conditional quantile of Y given a single predictor X_k , and $\theta_k(\tau)$ may differ from the derivative of $Q_\tau(Y/X)$ with respect to X_k . Since the seminal work of Fan and Lv (2008), marginal regression has been used for feature screening in various models. For conditional quantile screening, He et al. (2013) proposed a screening method based on marginal nonparametric quantile regression, and Shao and Zhang (2014) developed a model-free approach based on martingale difference correlation. Different from these works, our focus is not on screening but on testing the existence of overall covariate effect through marginal regression, and the proposed test can be used as a first step before variable selection.

We define the index of the most informative predictor at the τ th quantile as

$$k_0(\tau) = \underset{k=1, \dots, p}{\operatorname{argmin}} E[\rho_\tau\{Y - \alpha_k(\tau) - \theta_k(\tau)X_k\} - \rho_\tau(Y)], \quad (3)$$

and denote the corresponding slope parameter by $\theta_0(\tau) = \theta_{k_0(\tau)}(\tau)$. We focus on testing $H_0: \theta_0(\tau) = 0$ versus $H_a: \theta_0(\tau) \neq 0$. Note that the rejection of H_0 implies that at least one of the p predictors has an effect on the τ th conditional quantile of Y .

Based on the random sample, we define

$$\begin{aligned} \hat{k}_n(\tau) &= \underset{k=1, \dots, p}{\operatorname{argmin}} \mathbb{P}_n[\rho_\tau\{Y - \hat{\alpha}_k(\tau) - \hat{\theta}_k(\tau)X_k\}] \\ &\doteq \underset{k=1, \dots, p}{\operatorname{argmin}} n^{-1} \sum_{i=1}^n \rho_\tau\{y_i - \hat{\alpha}_k(\tau) - \hat{\theta}_k(\tau)x_{i,k}\}, \end{aligned} \quad (4)$$

where $(\hat{\alpha}_k(\tau), \hat{\theta}_k(\tau)) = \underset{\alpha, \theta}{\operatorname{argmin}} \mathbb{P}_n\{\rho_\tau(Y - \alpha - \theta X_k)\}$ are the sample quantile coefficient estimators obtained by regressing Y on X_k . The marginal quantile regression estimator of $\theta_0(\tau)$ is then given by $\hat{\theta}_n(\tau) = \hat{\theta}_{\hat{k}_n(\tau)}(\tau)$.

For any given $k = 1, \dots, p$, by Theorem 3 of Angrist et al. (2006), the asymptotic covariance of $(\hat{\alpha}_k(\tau), \hat{\theta}_k(\tau))$ is

$\sum_k(\tau) \doteq \mathbf{J}_k(\tau)^{-1} E[\{\psi_\tau(Y - \alpha_k(\tau) - \theta_k(\tau)X_k)\}^2 \tilde{\mathbf{X}}_k \tilde{\mathbf{X}}_k^T] \mathbf{J}_k(\tau)^{-1}$, where $\mathbf{J}_k(\tau) = E[f_Y\{\alpha_k(\tau) + \theta_k(\tau)X_k | X_k\} \tilde{\mathbf{X}}_k \tilde{\mathbf{X}}_k^T]$, $\psi_\tau(u) = \tau - I(u < 0)$ and $\tilde{\mathbf{X}}_k = (1, X_k)^T$. We can estimate $\mathbf{J}_k(\tau)$ by the kernel-based estimator suggested by Powell (1991),

$$\hat{\mathbf{J}}_k(\tau) = (nh_n)^{-1} \sum_{i=1}^n \mathcal{K}\{(y_i - \hat{\alpha}_k(\tau) - \hat{\theta}_k(\tau)x_{i,k})/h_n\} \tilde{\mathbf{x}}_{i,k} \tilde{\mathbf{x}}_{i,k}^T, \quad (5)$$

where $\mathcal{K}(\cdot)$ is a kernel function, and h_n is the positive bandwidth satisfying $h_n \rightarrow 0$ and $\sqrt{nh_n} \rightarrow \infty$. By Lemmas 1–3 in the Appendix and Theorem 3 of Powell (1991), it is easy to show that $\hat{\mathbf{J}}_k(\tau) - \mathbf{J}_k(\tau) = o_p(1)$ uniformly in τ . Alternatively we can also estimate $\mathbf{J}_k(\tau)$ by using the difference quotient method in Hendricks and Koenker (1991). Consequently, we can estimate $\Sigma_k(\tau)$ consistently by

$\hat{\Sigma}_k(\tau) = \hat{\mathbf{J}}_k^{-1}(\tau) \mathbb{P}_n(\{\tau - I(Y < \hat{\alpha}_k(\tau) + \hat{\theta}_k(\tau)X_k)\}^2 \tilde{\mathbf{X}}_k \tilde{\mathbf{X}}_k^T) \hat{\mathbf{J}}_k^{-1}(\tau)$. For any $k = 1, \dots, p$, denote the lower right diagonal element of $\hat{\Sigma}_k(\tau)$ as $\hat{\sigma}_k^2(\tau)$.

Our proposed test statistic is defined as

$$T_n(\tau) = \frac{n^{1/2} \hat{\theta}_n(\tau)}{\hat{\sigma}_n(\tau)}, \quad (6)$$

where $\hat{\sigma}_n(\tau) = \hat{\sigma}_{k_n(\tau)}(\tau)$. The test statistic $T_n(\tau)$ is a maximum-type statistic, which is a natural choice for testing H_0 . However, when none of the components of $\mathbf{X}$ have an effect on the τ th conditional quantile of Y , $k_0(\tau)$ is unidentifiable—it can be any of the p indices. In addition, the distribution of $n^{1/2}\{\hat{\theta}_n(\tau) - \theta_0(\tau)\}/\hat{\sigma}_n(\tau)$ does not converge uniformly with respect to $\theta_0(\tau)$ in the neighborhood of $\theta_0(\tau) = 0$, so the normal limiting distribution that holds away from $\theta_0(\tau) = 0$ cannot be used to construct rejection regions. To construct a suitable test procedure for H_0 , it is important to study the asymptotic behavior of $T_n(\tau)$ under local alternatives.

2.2 Local Model

In the local model, we replace the slope parameter $\beta_0(\tau)$ in model (1) by

$$\beta_n(\tau) = \beta_0(\tau) + n^{-1/2} \mathbf{b}_0(\tau), \quad (7)$$

where $\mathbf{b}_0(\tau) \in \mathbb{R}^p$ is the local parameter. When $\beta_0(\tau) = \mathbf{0}$, the quantile effect of $\mathbf{X}$ is $n^{-1/2} \mathbf{X}^T \mathbf{b}_0(\tau)$ and it vanishes asymptotically. Let $\varepsilon(\tau) = Y - \alpha_0 - \mathbf{X}^T \beta_n(\tau)$ denote the quantile regression residuals. It is clear that $Q_\tau\{\varepsilon(\tau) | \mathbf{X}\} = 0$. Throughout, we assume that the

distributions of $\mathbf{X}$ and $\varepsilon(\tau)$ are fixed and only the distribution of Y depends on n , but we suppress n in the notation of Y for notational simplicity.

Under the local model, we define

$$\bar{k}_n(\tau, \mathbf{b}_0) = \arg\min_{k=1, \dots, p} \min_{\alpha, \theta} E[\rho_\tau(Y - \alpha - X_k \theta) - \rho_\tau(Y)], \quad (8)$$

and rewrite $\theta_0(\tau)$ as $\theta_n(\tau) \doteq \theta_{\bar{k}_n(\tau, \mathbf{b}_0)}(\tau)$. Note that when $\mathbf{b}_0(\tau) = \mathbf{0}$, $\bar{k}_n(\tau, \mathbf{0})$ coincides with $k_0(\tau)$ defined in (3) under the global model. If $\beta_0(\tau) \neq \mathbf{0}$ and $k_0(\tau)$ is unique, then $\bar{k}_n(\tau, \mathbf{b}_0) \rightarrow k_0(\tau)$ and $\theta_n(\tau)$ is asymptotically bounded away from zero representing a non-local alternative case. On the other hand if $\beta_0(\tau) = \mathbf{0}$, then $\theta_n(\tau)$ is in the neighborhood of zero representing a local alternative case. However, if $\beta_0(\tau) = \mathbf{b}_0(\tau) = \mathbf{0}$, then $\bar{k}_n(\tau, \mathbf{b}_0)$ is not well defined and $\theta_n(\tau) = 0$, representing the null case. That is, under the local model (7), even though $\bar{k}_n(\tau, \mathbf{b}_0)$ is unidentifiable at $\beta_0(\tau) = \mathbf{b}_0(\tau) = \mathbf{0}$, it is still “weakly identifiable” when $\mathbf{b}_0(\tau) \neq \mathbf{0}$. By analyzing the local model, we can study the (non-regular) asymptotic behavior of the test statistic $T_n(\tau)$.

2.3 Limiting Behavior Under the Local Model

Under the local model, in which the response Y depends on n , $(\alpha_k(\tau), \theta_k(\tau))$ defined in (2) also depends on n , but throughout the paper we suppress this dependence to simplify notation. We make the following assumptions.

- A1** The conditional quantile regression model (1) with slope parameters defined in (7) hold, where $\alpha_0(\tau)$, $\beta_0(\tau)$ and $\mathbf{b}_0(\tau)$ are in the interior of compact sets.
- A2** Under the local model, for each $n \geq 1$, $k = 1, \dots, p$ and any $\tau \in \mathcal{T}$, there exist unique $(\alpha_k(\tau), \theta_k(\tau)) = \arg\min_{(\alpha, \theta)} E[\rho_\tau(Y - \alpha - \theta X_k) - \rho_\tau(Y)]$.
- A3** The covariate $\mathbf{X}$ has a compact support.
- A4** The conditional density $f_Y(\cdot|\mathbf{X})$ is bounded and uniformly continuous with a bounded first derivative, over the support of $\mathbf{X}$.
- A5** Under the local model, for $k = 1, \dots, p$, the matrices

$$\mathbf{J}_k\{\tau, \beta_0(\tau)\} = \lim_{n \rightarrow \infty} E \left[f_{\varepsilon(\tau)}\{e_k(\tau)|\mathbf{X}\} \tilde{\mathbf{X}}_k \tilde{\mathbf{X}}_k^T \right] \doteq \begin{pmatrix} \pi_k\{\tau, \beta_0(\tau)\} & \mu_k\{\tau, \beta_0(\tau)\} \\ \mu_k\{\tau, \beta_0(\tau)\} & \mu_{kk}\{\tau, \beta_0(\tau)\} \end{pmatrix}$$

and $\lim_{n \rightarrow \infty} E[(\psi_\tau\{\varepsilon(\tau) - e_k(\tau)\})^2 \tilde{\mathbf{X}}_k \tilde{\mathbf{X}}_k^T]$ exist and are positive definite for all $\tau \in \mathcal{T}$, where $e_k(\tau) = \alpha_k(\tau) + X_k \theta_k(\tau) - \alpha_0(\tau) - \mathbf{X}^T \beta_n(\tau)$ is the prediction bias due to marginal regression of Y on X_k at the quantile level τ .

The assumption A2 is an identifiability condition needed to ensure that the population quantile coefficient vector $(\alpha_k(\tau), \theta_k(\tau))$ obtained by regressing Y on X_k is unique. The assumptions A3–A5 concern $\mathbf{X}$ and $\varepsilon(\tau)$, and are needed to apply a result of Angrist et al.

(2006, Theorem 3) to obtain the asymptotic properties of the marginal regression estimator $(\hat{\alpha}_k(\tau), \hat{\theta}_k(\tau))$ under misspecification caused by omitting the correct predictors. Assumption A3 is needed to obtain the approximate representation of $(\alpha_k(\tau), \theta_k(\tau))$ when $\beta_0(\tau) = 0$. This condition can be relaxed by assuming some boundedness condition on the higher moments of $\mathbf{X}$, but this would greatly complicate the proofs. In theory, the boundedness in the support of $\mathbf{X}$ is naturally needed to avoid quantile crossing if multiple linear quantile regression functions are assumed. When $\beta_0(\tau) = 0$, $(\alpha_k(\tau), \theta_k(\tau)) \rightarrow (\alpha_0(\tau), 0)$ by the proof of Lemma 2, so the first limiting matrix in A5 is $\mathbf{J}_k(\tau, 0) = E \left\{ f_{\varepsilon(\tau)}(0|\mathbf{X}) \tilde{\mathbf{X}}_k \tilde{\mathbf{X}}_k^T \right\}$ with $\pi_k(\tau, 0) \equiv \pi(\tau, 0) = E\{f_{\varepsilon(\tau)}(0|\mathbf{X})\}$, $\mu_k(\tau, 0) = E\{f_{\varepsilon(\tau)}(0|\mathbf{X})X_k\}$ and $\mu_{kk}(\tau, 0) = E\{f_{\varepsilon(\tau)}(0|\mathbf{X})X_k^2\}$, and the second limit is $\tau(1-\tau)E(\tilde{\mathbf{X}}_k \tilde{\mathbf{X}}_k^T)$.

Theorem 1 gives the asymptotic representation of $n^{1/2}\{\hat{\theta}_n(\tau) - \theta_n(\tau)\}/\hat{\sigma}_n(\tau)$ when $\beta_0(\tau) \neq 0$ and $\beta_0(\tau) = 0$ separately. The asymptotic representation goes through a phase transition at $\beta_0(\tau) = 0$, with a different form for $\beta_0(\tau) \neq 0$ (in which case $\bar{k}_n(\tau, \mathbf{b}_0)$ is identifiable). We first fix some notation. Let $\mathbf{X} = (X_1, \dots, X_p)^T$ and $\mathbb{X}_j$ be its j th element for $j = 1, \dots, 2p$. In addition, for $k = 1, \dots, p$, let $V_k(\tau, \beta_0) = |\mathbf{J}_k(\tau, \beta_0)|$, $\mathbf{C}_k(\tau) = E\{f_{\varepsilon(\tau)}(0|\mathbf{X})\} E\{f_{\varepsilon(\tau)}(0|\mathbf{X})X_k\mathbf{X}\} - E\{f_{\varepsilon(\tau)}(0|\mathbf{X})X_k\} E\{f_{\varepsilon(\tau)}(0|\mathbf{X})\mathbf{X}\}$ and $\mathbf{B}_k(\tau) = E\left\{f_{\varepsilon(\tau)}(0|\mathbf{X})\mathbf{X}\tilde{\mathbf{X}}_k^T\right\}$. Let $\mathbf{M}(\tau, \beta_0(\tau)) \doteq (M_1(\tau, \beta_0(\tau)), \dots, M_{2p}(\tau, \beta_0(\tau)))^T$ be a Gaussian process with mean zero and covariance function $\Omega(\tau, \tau')$, whose (j, j') th element is $\Omega_{jj'}(\tau, \tau') = \lim_{n \rightarrow \infty} E(\mathbb{X}_j \mathbb{X}_{j'} \{ \tau - I\{\varepsilon(\tau) < e_j(\tau)\} \} [\tau' - I\{\varepsilon(\tau') < e_{j'}(\tau')\}])$, where $e_j(\tau) = e_{j-p}(\tau)$ for $j = p+1, \dots, 2p$.

Theorem 1—Suppose that assumptions A1–A5 hold, and for each $\tau \in \mathcal{T}$, $\bar{k}_n(\tau, 0) \equiv k_0(\tau)$ is unique when $\beta_0(\tau) \neq 0$, and $\bar{k}_n\{\tau, \mathbf{b}_0(\tau)\} \rightarrow \kappa_\tau\{\mathbf{b}_0(\tau)\} \in \{1, \dots, p\}$ when $\beta_0(\tau) = 0$ and $\mathbf{b}_0(\tau) \neq 0$. Then jointly over $\tau \in \mathcal{T}$, we have

$$\frac{n^{1/2}\{\hat{\theta}_n(\tau) - \theta_n(\tau)\}}{\hat{\sigma}_n(\tau)} \xrightarrow{d} \begin{cases} \frac{M_{p+k_0(\tau)}(\tau, \beta_0(\tau))\pi_{k_0(\tau)}(\tau, \beta_0(\tau)) - M_{k_0(\tau)}(\tau, \beta_0(\tau))\mu_{k_0(\tau)}(\tau, \beta_0(\tau))}{V_{k_0(\tau)}(\tau, \beta_0(\tau))\sigma_{k_0(\tau)}(\tau)} & \text{if } \beta_0(\tau) \neq 0, \\ \frac{M_{p+K(\tau)}(\tau, 0)\pi(\tau, 0) - M_{K(\tau)}(\tau, 0)\mu_{K(\tau)}(\tau, 0)}{V_{K(\tau)}(\tau, 0)\sigma_{K(\tau)}(\tau)} + \left(\frac{\mathbf{C}_{K(\tau)}(\tau)}{V_{K(\tau)}(\tau, 0)} - \frac{\mathbf{C}_{\kappa_\tau\{\mathbf{b}_0(\tau)\}}(\tau)}{V_{\kappa_\tau\{\mathbf{b}_0(\tau)\}}(\tau, 0)} \right)^T \frac{\mathbf{b}_0(\tau)}{\sigma_{K(\tau)}(\tau)} & \text{if } \beta_0(\tau) = 0, \end{cases}$$

where

$$K(\tau) = \operatorname{argmax}_{k=1, \dots, p} \{ \mathbf{M}_k(\tau) + \mathbf{B}_k(\tau)^T \mathbf{b}_0(\tau) \}^T \mathbf{J}_k(\tau, 0)^{-1} \{ \mathbf{M}_k(\tau) + \mathbf{B}_k(\tau)^T \mathbf{b}_0(\tau) \} \quad (9)$$

with $\mathbf{M}_k(\tau) = (M_k(\tau, 0), M_{k+p}(\tau, 0))^T$.

In Theorem 1, the convergence condition on $\bar{k}_n\{\tau, \mathbf{b}_0(\tau)\}$ is a weak continuity condition on the joint distribution of $\varepsilon(\tau)$ and $\mathbf{X}$. Due to the prediction bias $e_k(\tau)$ from marginal regression, $\Omega_{jj'}(\tau, \tau')$ does not have an explicit expression. However, when $\beta_0(\tau) = \beta_0(\tau') = 0$, $e_k(\tau)$ and $e_k(\tau')$ go to zero under the local model for all $k = 1, \dots, p$, and consequently $\Omega_{jj'}(\tau, \tau') = \{\min(\tau, \tau') - \tau\tau'\} E(\mathbb{X}_j \mathbb{X}_{j'})$.

We refer to the special case in which $f_{\varepsilon(\tau)}(\cdot|\mathbf{X}) = f_{\varepsilon(\tau)}(\cdot)$ is the same across all values of $\mathbf{X}$ as *homoscedastic*. That is, the distribution of the regression error $\varepsilon(\tau)$ does not depend on the covariates. In this case, $V_k(\tau, \mathbf{0}) = f_{\varepsilon(\tau)}^2(0)\text{Var}(X_k)$, $\mathbf{C}_k(\tau) = f_{\varepsilon(\tau)}^2(0)\text{Cov}(X_k, \mathbf{X})$, and the limiting distribution in Theorem 1 when $\boldsymbol{\beta}_0(\tau) = \mathbf{0}$ can be simplified as in the following Corollary.

Corollary 1—Under the assumptions of Theorem 1, for the homoscedastic case, we have that jointly over all $\tau \in \mathcal{T}$ such that $\boldsymbol{\beta}_0(\tau) = \mathbf{0}$,

$$\frac{n^{1/2}\{\hat{\theta}_n(\tau) - \theta_n(\tau)\}}{\hat{\sigma}_n(\tau)} \xrightarrow{d} \frac{M_{p+K(\tau)}(\tau) - M_{K(\tau)}(\tau)E(X_K)}{f_{\varepsilon(\tau)}(0)\text{Var}(X_{K(\tau)})\sigma_{K(\tau)}(\tau)} + \left(\frac{\text{Cov}(X_K, \mathbf{X})}{\text{Var}(X_K)} - \frac{\text{Cov}(X_{K(\tau)\{\mathbf{b}_0(\tau)\}}, \mathbf{X})}{\text{Var}(X_{K(\tau)\{\mathbf{b}_0(\tau)\}})} \right)^T \frac{\mathbf{b}_0(\tau)}{\sigma_{K(\tau)}(\tau)},$$

where $K(\tau) = \arg\max_{k=1, \dots, p} \{M_{p+k}(\tau) - M_k(\tau)E(X_k) + f_{\varepsilon(\tau)}(0)\mathbf{b}_0^T(\tau)\text{Cov}(X_k, \mathbf{X})\}^2 / \text{Var}(X_k)$
with $M_k(\tau) = M_k(\tau, \mathbf{0})$.

2.4 Adaptive Bootstrap for Testing at a Single Quantile

By Theorem 1, we can obtain the asymptotic critical values for the test statistic $T_n(\tau)$ by simulating its asymptotic representation under the null hypothesis (with $\boldsymbol{\beta}_0(\tau) = \mathbf{b}_0(\tau) = \mathbf{0}$). This approach, however, requires estimating the weighted covariance matrix of $\mathbf{X}$ with weights to accommodate heteroscedasticity, and thus does not perform well in finite samples with large p . For practical purposes, we propose to adopt the idea in McKeague and Qian (2015) and develop an adaptive bootstrap procedure.

Denote $\{(y_i^*, \mathbf{x}_i^*), i = 1, \dots, n\}$ as the bootstrap sample obtained by sampling the observed data with replacement. The conventional bootstrap version of $R_n(\tau) \doteq n^{1/2}\{\hat{\theta}_n(\tau) - \theta_n(\tau)\} / \hat{\sigma}_n(\tau)$ is $R_n^*(\tau) \doteq n^{1/2}\{\hat{\theta}_n^*(\tau) - \hat{\theta}_n(\tau)\} / \hat{\sigma}_{n^*}(\tau)$, where $\hat{\theta}_n^*(\tau)$ is the bootstrap versions of $\hat{\theta}_n(\tau)$, and $\hat{\sigma}_{n^*}(\tau) = \hat{\sigma}_{\hat{k}_n^*(\tau)}(\tau)$ with $\hat{k}_n^*(\tau)$ as the bootstrap counterpart of $\hat{k}_n(\tau)$.

However, since $\hat{\theta}_n(\tau)$ is a non-regular estimator, i.e., its distribution does not converge uniformly in a $n^{-1/2}$ neighborhood of the null, the conventional bootstrap would fail to estimate the null distribution. To account for the discontinuity in the limiting distribution of $R_n(\tau)$, we propose to compare $|T_n(\tau)|$ with some threshold $\lambda_n(\tau) > 0$ to capture the different behaviors of $R_n(\tau)$ in two scenarios: $\boldsymbol{\beta}_0(\tau) \neq \mathbf{0}$ (away from the null) and $\boldsymbol{\beta}_0(\tau) = \mathbf{0}$ (in the $n^{-1/2}$ neighborhood of the null). We consider the following modified bootstrap version of $R_n(\tau)$,

$$R_n^*\{\tau, \mathbf{b}_0(\tau)\} = \begin{cases} \frac{n^{1/2}\{\hat{\theta}_n^*(\tau) - \hat{\theta}_n(\tau)\}}{\hat{\sigma}_{n^*}(\tau)} & \text{if } |T_n(\tau)| > \lambda_n(\tau) \text{ or } |T_n^*(\tau)| > \lambda_n(\tau), \\ \frac{n^{1/2}\{\hat{\theta}_n^*(\tau) - \hat{\theta}_n(\tau)\}}{\hat{\sigma}_{n^*}(\tau)} & \text{otherwise,} \end{cases} \quad (10)$$

where $T_n^*(\tau) = n^{1/2} \hat{\theta}_n^*(\tau) / \hat{\sigma}_{n*}(\tau)$, ${}_{n*}^*(\tau, \mathbf{b})$ is the bootstrap version of $\mathbb{V}_n(\tau, \mathbf{b})$, a process in $\mathbf{b} \in \mathbb{R}^p$ defined in (11) that denotes the asymptotic representation of $R_n(\tau)$ in the local model. The definition of ${}_{n*}^*(\tau, \mathbf{b}(\tau))$ can be found in (13). For computational convenience, we use the sample estimator $\hat{\sigma}_{n*}(\tau)$ in the bootstrap to avoid recalculating $\hat{\mathbf{J}}_k(\tau)$ for each bootstrap sample and each k , but one can also use the bootstrap variance estimator to achieve potential second-order accuracy (Efron and Tibshirani, 1993; Hall and Kang, 2001).

The idea of the modified bootstrap is as follows. If $\lambda_n(\tau) = \alpha(n^{1/2})$ and $\lambda_n(\tau) \rightarrow \infty$, Theorem 1 indicates that $P\{|T_n| > \lambda_n(\tau)\} \rightarrow \mathbb{I}(\beta_0 = 0)$, thus $R_n(\tau)$ can be bootstrapped consistently by the naive bootstrap. The challenge lies in bootstrapping $R_n(\tau)$ when $\beta_0(\tau) = 0$ since $\hat{k}_n(\tau)$ does not converge to $k_0(\tau)$ in this case. From Lemma 4 and (S.15) in the Supplementary Material, we know that when $\beta_0(\tau) = 0$, the estimated asymptotic representation of $R_n(\tau)$ is $\mathcal{V}_n\{\tau, \mathbf{b}_0(\tau), \hat{k}_n(\tau)\}$, where

$$\mathcal{V}_n(\tau, \mathbf{b}, k) = \left[\frac{(-\hat{\mu}_k(\tau), \hat{\pi}(\tau)) \mathbb{G}_n[\tilde{\mathbf{X}}_k \psi_\tau\{\varepsilon(\tau)\}]}{\hat{V}_k(\tau) \sigma_k(\tau)} + \left(\frac{\hat{\mathbf{C}}_k(\tau)}{\hat{V}_k(\tau)} - \frac{\hat{\mathbf{C}}_{\bar{k}_n(\tau, \mathbf{b})}}{\hat{V}_{\bar{k}_n(\tau, \mathbf{b})}(\tau)} \right)^T \frac{\mathbf{b}}{\sigma_k(\tau)} \right],$$

$\mathbb{G}_n = n^{1/2}(\mathbf{P}_n - P_n)$ with P_n being the distribution of $(Y, \mathbf{X})$, $\hat{V}_k(\tau) = |\hat{\mathbf{J}}_k(\tau)|$, $(\hat{\pi}(\tau), \hat{\mu}_k(\tau))$ are the elements in the first row of $\hat{\mathbf{J}}_k$, and $\hat{\mathbf{C}}_k(\tau)$ is the kernel-based estimator of $\mathbf{C}_k(\tau)$.

Therefore, we define the $\mathbb{V}_n(\tau, \mathbf{b})$ process as

$$\begin{aligned} \mathbb{V}_n(\tau, \mathbf{b}) &= \mathcal{V}_n\{\tau, \mathbf{b}, \mathbb{K}_n(\mathbf{b})\}, \\ \mathbb{K}_n(\tau, \mathbf{b}) &= \underset{k=1, \dots, p}{\operatorname{argmax}} \left(\mathbb{G}_n[\tilde{\mathbf{X}}_k \psi_\tau\{\varepsilon(\tau)\}] + {}^T_k(\tau) \mathbf{b} \right)^T \hat{\mathbf{J}}_k^{-1}(\tau) \left(\mathbb{G}_n[\tilde{\mathbf{X}}_k \psi_\tau\{\varepsilon(\tau)\}] + \hat{\mathbf{B}}_k^T(\tau) \mathbf{b} \right), \end{aligned}$$

(11)

where $\hat{\mathbf{B}}_k(\tau)$ is the kernel-based estimator of $\mathbf{B}_k(\tau)$.

The following Theorem 2 shows that the adaptive bootstrap version $R_n^*\{\tau, \mathbf{b}_0(\tau)\}$ provides a consistent estimation of the distribution of $R_n(\tau)$ jointly over $\tau \in \mathcal{T}$.

Theorem 2—Suppose that assumptions in Theorem 1 hold, and $\lambda_n(\tau) \rightarrow \infty$ and $\lambda_n(\tau) = \alpha(n^{1/2})$ for all $\tau \in \mathcal{T}$. Then $R_n^*\{\tau, \mathbf{b}_0(\tau)\}$ converges to the limiting distribution of $R_n(\tau)$ jointly over $\tau \in \mathcal{T}$ conditionally (on the observed data) in probability.

For the calibration of the test statistic, we only need compute $R_n^*(\tau) \doteq R_n^*(\tau, \mathbf{0})$, the bootstrap statistic with $\mathbf{b} = \mathbf{0}$. Let $R_{n(j)}^*(\tau)$, $j = 1, \dots, m$, be the bootstrap statistics from m bootstrap repetitions. For a level- γ test, we can calculate the lower and upper critical values of $T_n(\tau)$ by the $(\gamma/2)$ th and $(1 - \gamma/2)$ th sample quantiles of $\{R_{n(j)}^*(\tau), j = 1, \dots, m\}$. If $T_n(\tau)$

is smaller than the lower or larger than the upper critical values, we will reject the null hypothesis and declare that there is at least one predictor that has significant effect on the τ th quantile of Y . Alternatively, we can also compute the p -value as

$$2 \min \left(m^{-1} \sum_{j=1}^m I\{R_{n(j)}^*(\tau) > T_n(\tau)\}, m^{-1} \sum_{j=1}^m I\{R_{n(j)}^*(\tau) < T_n(\tau)\} \right).$$

2.5 Testing Across Multiple Quantiles

One attractive feature of quantile regression is that it allows us to assess the relationship between Y and $\mathbf{X}$ at multiple quantiles. We now discuss testing across multiple quantiles through marginal quantile regression. Let $\mathcal{T} = \{\tau_1, \dots, \tau_L\}$ be a set of quantile levels of interest. To test whether any component of $\mathbf{X}$ has an effect on either one of the L quantiles, we can consider testing $\tilde{H}_0: \theta_0(\tau_1) = \dots = \theta_0(\tau_L) = 0$ versus $\tilde{H}_a$: at least one of the $\theta_0(\tau_l)$'s is nonzero, where $\theta_0(\tau_l)$ is the slope of the most informative predictor at the τ_l th quantile as defined in (3).

To pool information across quantiles, we propose to consider the sum-type test statistic:

$S_n = \sum_{l=1}^L T_n^2(\tau_l)$, where $T_n(\tau)$ is defined in (6). By the joint convergence result in Theorem 2 over $\tau \in \mathcal{T}$, we can extend the adaptive bootstrap for the test calibration of S_n .

Specifically, define the modified bootstrap statistic as $S_n^* = \sum_{l=1}^L R_n^{*2}(\tau_l, 0)$. The p -value of the multiple-quantile test can then be calculated by the proportion of modified bootstrap statistics that are larger than S_n .

One may also consider the maximum-type test statistic $\max_{1 \leq l \leq L} |T_n(\tau_l)|$ to combine information across quantiles. We choose the sum-type test statistic as it has good power against dense alternatives, and in quantile regression $\theta_0(\tau)$ is more likely to be nonzero in an interval of τ and thus dense when the alternative is true. The joint test for $\tilde{H}_0$ is an omnibus test and may not be informative in applications where the linear effect exists at some but not all quantile levels. In such cases, one may apply the test at several individual quantile levels after the joint test rejects $\tilde{H}_0$ to get a more comprehensive picture.

2.6 Selection of the Tuning Parameter $\lambda_n(\tau)$

The tuning parameter $\lambda_n(\tau)$ is involved in the pre-test with criterion $|T_n(\tau)| > \lambda_n(\tau)$ to determine whether a conventional bootstrap $R_n^*(\tau)$ or a bootstrap of the asymptotic representation $\mathbb{V}_n(\tau, \mathbf{b})$ should be used. In the extreme case with $\lambda_n(\tau) = 0$, the proposed bootstrap reduces to the naive bootstrap, which gives inflated Type I error rate. When $\lambda_n(\tau) = \alpha(n^{1/2})$ and $\lambda_n(\tau) \rightarrow \infty$, satisfying the conditions in Theorem 2, the pre-test has asymptotically negligible Type I error as $\lim_{n \rightarrow \infty} P\{|T_n(\tau)| > \lambda_n(\tau) | \theta_n(\tau) = 0\} = 0$, so that only bootstrap of $\mathbb{V}_n(\tau, \mathbf{0})$ is used under the null hypothesis that $\theta_n(\tau) = 0$. In finite samples, however, if $\lambda_n(\tau)$ is overly large, the test will be too conservative. Our empirical

investigation shows that for large samples, $\lambda_n(\tau) = c \sqrt{\tau(1-\tau)\log(n)}$ with $c \in [4, 10]$ provides a good choice, where the chosen rate is proportional to the expected value of the maximum of n sub-Gaussian variables with mean zero. In finite samples, we propose to use

a double bootstrap procedure to choose the constant c . For illustration, we describe the procedure for testing at a single quantile τ as follows.

For a bootstrap sample $\{y_i^*, x_i^*, i = 1, \dots, n\}$, obtain the first-level bootstrap statistics $\hat{\theta}_n^*(\tau)$ and $\hat{\sigma}_{n^*}(\tau)$, and m double bootstrap statistics $R_{(1)}^{**}, \dots, R_{(m)}^{**}$, the analogies of $R_n^*(\tau, 0)$ as defined in (10), associated with a candidate c value. For each c in the grid, we calculate the rejection rate as the proportion of bootstrap samples for which $n^{1/2}\{\hat{\theta}_n^*(\tau) - \hat{\theta}_n(\tau)\} / \hat{\sigma}_{n^*}(\tau)$ exceeds the double bootstrap critical values, determined by the lower bound of the $(\gamma/2)$ th and the upper bound of the $(1 - \gamma/2)$ th sample quantiles of $\{R_{(1)}^{**}, \dots, R_{(m)}^{**}\}$. The double bootstrap then chooses the c value that gives the rejection rate closest to the nominal significance level γ . The double bootstrap could be time consuming, so in practice we recommend use double bootstrap with $m = 100$ for smaller samples of $n \leq 500$ and choose a rather conservative $c = 5$ for larger samples.

3 Simulation Study

3.1 Under Linear Regression Models

We first assess the performance of the proposed test in linear regression models by considering four cases. For Cases 1–3, data are generated from the model: $Y = \mathbf{X}^T \boldsymbol{\beta} + \varepsilon$, where $\varepsilon \sim N(0, 1)$ in Cases 1–2 and $\varepsilon \sim t_2$ in Case 3. Case 1 corresponds to the null hypothesis of no active predictors with $\boldsymbol{\beta} = \mathbf{0}$. In Case 2, we let $\boldsymbol{\beta} = (1/3, 0, \dots, 0)^T$ so there is a unique active predictor. In Case 3, we let $\beta_1 = \dots = \beta_5 = 0.25$, $\beta_6 = \dots = \beta_{10} = -0.15$ and $\beta_j = 0$ for $j = 10, \dots, p$, so there are ten active predictors and k_0 is not unique. For Case 4, data are generated from the model: $Y = \mathbf{X}^T \boldsymbol{\beta} + (1 + 0.45X_1)\varepsilon$, where $\boldsymbol{\beta} = \mathbf{0}$ and $\varepsilon \sim N(0, 1)$. Case 4 represents a heteroscedastic case, where the predictor X_1 is active at $\tau = 0.5$ but inactive at $\tau = 0.25$. In all cases, the covariate vector $\mathbf{X} = (X_1, \dots, X_p)^T$ is from the multivariate normal distribution with mean zero, variance one and an exchangeable correlation of 0.5, truncated at -2 and 2 . For each case, a random sample of size $n = 200$ of $(Y, X_1, \dots, X_p)$ is generated, and the simulation is repeated 500 times. Even though the asymptotic results assume a finite p , we shall demonstrate that the proposed test works empirically also for cases with $p > n$ through the analysis for five dimensions, $p = 10, 100, 200, 400$ and 1000 .

We apply the proposed quantile marginal effect test (QMET) at a single quantile level $\tau = 0.5$ and 0.75 , separately, and at three quantiles $\tau = \{0.25, 0.5, 0.75\}$ jointly to detect the predictor effects on the conditional quantiles of Y . For comparison, we also include three competing methods: (1) AOV; (2) BONF; (3) CPB for single-quantile testing. Each test is calibrated at the nominal level $\gamma = 0.05$. The AOV is a rank-score-type method for testing the null hypothesis that no predictors have significant effects against the alternative that at least one predictor has a significant effect, and this method is feasible only in cases where $p < n$. The AOV method is implemented in the function “anova.rq” of the R package *quantreg*. The BONF is a method of multiple comparison with Bonferroni adjustments, where a rank-score test is carried out in the marginal quantile regression of Y on each X_j , $j = 1, \dots, p$, and the null hypothesis is rejected if the minimum of p resulting p -values is smaller than γ . The rank-score-based tests are chosen for AOV and BONF since such tests were shown to be

more stable than the Wald-type tests in the quantile regression literature (Kocherginsky et al., 2005). The CPB is the centered percentile bootstrap method based on the conventional bootstrap statistic $R_n^*(\tau)$. For both CPB and QMET, 200 bootstrap samples are used. For the QMET method, we let the threshold $\lambda_n(\tau) = c \sqrt{\tau(1-\tau)\log(n)}$ and choose $c \in (0, 6)$ by double bootstrap with 100 double bootstrap samples.

Table 1 summarizes the rejection rates of different methods against p . The rejection rates for Case 1 with $\tau = 0.5, 0.75$ and $\tau = \{0.25, 0.5, 0.75\}$, and for Case 4 with $\tau = 0.5$ correspond to Type I errors where the null hypothesis is true, while they represent power in the other scenarios. In all three null scenarios, the CPB method gives high Type I errors, confirming that the conventional bootstrap procedure fails to control the family-wise error rate as it does not account for the uncertainty involved in the variable selection process. The AOV test is not feasible in cases with $p = n$, and it is highly conservative with power close to zero for $p = 100$. The Bonferroni correction method has similar performance with QMET for $p = 10$, but the QMET method is in general more powerful for larger p . Compared to the other methods, the QMET method with double bootstrap maintains the level reasonably well and it provides relatively high power across all scenarios considered. Under alternative models with homoscedastic errors, the QMET across quantiles tends to be more powerful than the single-quantile test especially at the tail quantiles. However, in heteroscedastic models (Case 4) where the signal has different magnitudes at different quantiles, it may be more advantageous to apply the test at a single quantile level with stronger signal than the omnibus test across quantiles.

3.2 Under Misspecified Nonlinear Regression Models

We consider two additional cases to assess the performance of the proposed test under the misspecification of linear models and for covariates with a different correlation structure.

The data is generated from the model: $Y = X_1^2/3 + bX_1 + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, 1)$, and the covariates $(X_1, \dots, X_p)$ are generated in the same way as in Cases 1–4 but with an AR(1) correlation structure of parameter 0.5. We let $b = 0$ in Case 5, and $b = 1/3$ in Case 6. Table 2 summarizes the rejection rates of the QMET and BONF methods at $\tau = 0.5$ in Cases 5–6. Note that the proposed test is only for detecting linear covariate effects. When there exists no linear relationship between $\mathbf{X}$ and Y at any quantiles (Case 5), the QMET has difficulty identifying the nonlinear relationship. However, if there exists some linear trend in a misspecified model such as in Case 6, the QMET can still identify the covariate effects with higher power than the Bonferroni method. The idea of the QMET method may be extended to detect nonlinear covariate effects by marginally regressing Y on some polynomial or basis functions of each covariate separately, and we leave this to future research.

3.3 Sensitivity Against the Bandwidth h_n

The calculation of $\hat{\sigma}_k(\tau)$ involves a bandwidth parameter h_n . In our implementation, we follow the suggestion in Hall and Sheather (1988) and choose $h_n = an^{-1/3} \{\Phi^{-1}(1 - \gamma/2)\}^{2/3} (1.5\phi^2\{\Phi^{-1}(\tau)\}/[2\{\Phi^{-1}(\tau)\}^2 + 1])^{2/3}$ with $a = 1$, where Φ is the distribution function of $\mathcal{N}(0, 1)$. To assess the sensitivity of the proposed test against h_n , we plot the rejection rates of

QMET at $\tau = 0.5$ versus a in Cases 1–4 in Figure 1. Results suggest that the performance of the proposed test is quite stable for $a \in [0.5, 1.5]$.

4 Application to the Study of HIV Drug Resistance

We illustrate the proposed method by analyzing a human immunodeficiency virus (HIV) drug susceptibility dataset from the HIV Drug Resistance Database (<http://hivdb.stanford.edu>), a public resource for the study of sequence variation and mutations in the molecular targets of HIV drugs (Rhee et al., 2003). After a patient starts antiretroviral therapy, the infecting HIV virus can form new mutations. Some mutations may not respond to existing drugs, a characteristic known as drug resistance or reduced drug susceptibility, meaning that the drugs become less effective at preventing the virus from multiplying. Researchers have estimated that in an untreated HIV-infected subject, every possible single point mutation occurs between 10^4 and 10^5 times per day (Coffin, 1995). Drug resistance has become a major obstacle to the success of HIV therapy. Therefore, understanding the impact of mutations on drug resistance is an important research topic.

We analyze the susceptibility data for the drug efavirenz (EFV). After excluding rare mutations, the data set includes 1472 HIV isolates and 197 locations of mutations. The susceptibility of a HIV sample is defined as the fold decrease in susceptibility of a single virus isolate compared with the susceptibility of a wildtype control isolate, that is, the virus that has never been challenged by drugs. We focus on predicting $\log_{10}$ -susceptibility, denoted by Y , to EFV based on X_k , $k = 1, \dots, p = 197$, indicating the presence of a mutation of interest in the k th viral sequence position. The susceptibility data are highly non-normal, even after log transformation (see Figure 3), so quantile regression provides a valuable way of analyzing these data. In addition, analysis at the upper quantiles of susceptibility is of particular interest, being associated with stronger drug resistance. In this analysis, we consider two quantile levels: $\tau = 0.5$ and 0.75 .

The proposed test can be used as a stopping rule in forward regression to select multiple significant predictors. Specifically, suppose that the test detects a significant predictor in the first step. We can then use the residuals $Y - \hat{\theta}_n(\tau)X_{k_1}(\tau)$ as a new outcome variable and carry out the marginal quantile regression over the remaining predictors. Repeat the procedure and record the p -values in the sequential tests as $p_1, p_2, \dots$. The procedure will stop at the m th step, where $m = \inf\{j: p_j > \gamma\}$ and γ is the nominal significance level. To account for the sequential testing involved in the forward regression, we further employ a multiple test adjustment in the style of Holm (1979), and finally choose the covariates identified in the first $\tilde{m}$ steps, where $\tilde{m} = 1$ if $p_1 > \gamma/(m-1)$, and otherwise $\tilde{m} = \max_{j=1, \dots, m-1} \{j: p_j \leq \gamma/(m-j)\}$ for all $j = 1, \dots, m$.

To assess the performance of the proposed method in settings with $n < p$, we randomly split the data into a training set of size $n = 190$ and a testing set of size 1282. For each split, we carry out 20 steps of forward quantile marginal effect testing (QMET) and standard forward selection quantile regression (FWD) using the training data, and use the model selected in each step to predict the τ th quantile of log susceptibility of the testing data.

The left panel of Figure 2 plots the training-set p -values ($\pm$ MAD) for the newly entered predictor at each step across 50 random splits at $\tau = 0.5$ and $\tau = 0.75$, and the right panel plots the corresponding prediction errors (median $\pm$ MAD) in the test sets, where MAD is the median absolute deviation. At a quantile level τ , the quantile prediction error is defined as the average of quantile loss $\rho_{\tau}(Y - \hat{Y})$, where Y is the true response and $\hat{Y}$ is the predicted value for subjects from the testing data. At the 0.05 significance level, the QMET method selects about one predictor at median and three at level $\tau = 0.75$, while the FWD method selects more than 20 predictors. The prediction error plots suggest that with the QMET method, the improvement in prediction error becomes negligible after the first three predictors enter the model. In contrast, the prediction accuracy of the FWD method improves much more slowly. Generally speaking, at both quantile levels, the FWD enters about 15 predictors to achieve the same prediction accuracy as the model selected by QMET with just three predictors.

The obvious difference in the prediction accuracy suggests that two methods enter different predictors at each stage, which is not surprising given their different selection criteria. In each step of the forward QMET method, the QMET procedure is applied by treating residuals from the previous stage as new outcomes and identifies the predictor that gives the smallest quantile loss in marginal regression. In contrast, the standard FWD method identifies the predictor that gives the smallest Wald-type p -value conditional on the predictors that have entered prior to the current step.

To demonstrate the value of regression at different quantiles, we look into one example training set, for which the forward QMET method at $\tau = 0.75$ at the significance level of 0.05 selects three predictors, 103N, 190A and 230L, sequentially, while the method at median selects only the first predictor. The binary predictors 103N, 190A and 230L indicate the presence of substitution of amino acid Asparagine, Alanine and Leucine at positions 103, 190 and 230, respectively. Figure 3 shows the box plots of the $\log_{10}$ susceptibility of the isolates in the training set with and without the three mutations, respectively. The boxplots suggest that isolates with these three mutations are associated with higher drug resistance at both median and $\tau = 0.75$. After accounting for the effects of the first mutation, the effects of 190A and 230L become insignificant at median but remain significant at $\tau = 0.75$, that is, for the isolates that are more drug resistant; see Table 3.

5 Discussion

We have developed a new procedure for detecting the presence of marginal effects in quantile regression. Our simulation study suggests that the proposed test is effective and has stable performance, providing adequate control of family-wise error rate along with competitive power, even for cases with large p (although the asymptotic theory we used to calibrate the test assumes fixed p). The theoretical study for diverging p is an interesting and challenging topic that deserves further research.

An alternative approach to QMET beyond those we have considered would be to use higher criticism, a way of synthesizing a collection of p -values from multiple tests that was originally proposed by Tukey (1989) and later systematically developed by Donoho and Jin

(2004, 2015). However, as reported in McKeague and Qian (2015) for the mean regression version of our proposed test, higher criticism is typically anti-conservative unless the predictors are close to be uncorrelated, which would be a highly restrictive assumption in most applications.

Our proposed test statistic can be viewed as a maximum-type test statistic across p covariates. Similar to the discussion as in Chatterjee and Lahiri (2015) for mean regression, we may also consider an alternative statistic based on the sum of squared t -statistics,

$\sum_{k=1}^p \hat{\theta}_k^2(\tau) / \hat{\sigma}_k^2(\tau)$. Our empirical studies show that the maximum-type test is more powerful for detecting sparse signals, and the sum-type test has more power for dense alternatives; see results in the Supplementary Material. This observation agrees with the findings in mean regression (Cai et al., 2014; Gregory et al., 2015; Chen and Qin, 2010; Fan et al., 2015). Therefore, we would recommend the proposed test for scenarios with sparse signals. How to combine two types of test statistics to accommodate signals of unknown sparsity levels is an interesting future research topic.

The proposed test can be used as a first step in applications to assess the overall significance of covariates on quantiles of Y . If the null hypothesis is rejected, one can use existing model selection procedures to see how many and which variables should be included next. Alternatively, one may use the proposed test as a stopping rule in forward regression to select multiple significant predictors in a sequential manner, and our empirical study has shown some promising evidence. There exist some recent work for sequential forward regression on error rate control and the probabilistic bounds for the number of selected covariates (Fithian et al., 2015; Li and Barber, 2015; G'Sell et al., 2016; Kozbur, 2015; Tibshirani et al., 2016). However, due to the dependence among the new outcomes and the non-additive property of quantile functions, it would be challenging and also beyond the scope of the paper to study the theoretical properties of the forward regression procedure.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The research is partly supported by NSF grants DMS-1149355, DMS-1307838, DMS-1712760, NIH Grant 2R01GM095722, R21MH108999 and the OSR-2015-CRG4-2582 grant from KAUST. The authors would like to thank the editor, an associate editor and three referees for their constructive comments and suggestions.

References

- Angrist J, Chernozhukov V, Fernández-Val I. Quantile regression under misspecification, with an application to the u.s. wage structure. *Econometrica*. 2006; 74:539–563.
- Belloni A., Chernozhukov V., Kato K. Tech rep. Duke University; 2014. Valid post-selection inference in highdimensional approximately sparse quantile regression models.
- Belloni A, Chernozhukov V, Kato K. Uniform post selection inference for LAD regression and other Z-estimation problems. *Biometrika*. 2015; 102:77–94.
- Cai TT, Liu W, Xia Y. Two-sample test of high dimensional means under dependence. *Journal of the Royal Statistical Society, Ser B*. 2014; 76:349–372.

- Chatterjee A, Lahiri SN. Comment on “an adaptive resampling test for detecting the presence of significant predictors” by I. McKeague and M. Qian. *Journal of American Statistical Association*. 2015; 110:1434–1438.
- Chen SX, Qin Y. A two-sample test for high-dimensional data with applications to gene-set testing1. *The Annals of Statistics*. 2010; 38:808–835.
- Coffin JM. HIV population dynamics in vivo: implications for genetic variation, pathogenesis, and therapy. *Science*. 1995; 267:483–489. [PubMed: 7824947]
- Donoho D, Jin J. Higher criticism for detecting sparse heterogeneous mixtures. *The Annals of Statistics*. 2004; 32:962–994.
- Donoho D, Jin J. Higher criticism for large-scale inference, especially for rare and weak effects. *Statistical Science*. 2015; 30:1–25.
- Efron, B., Tibshirani, BJ. *An Introduction to the Bootstrap*. London: Chapman & Hall; 1993.
- Fan J, Liao Y, Yao J. Power enhancement in high-dimensional cross-sectional tests. *Econometrica*. 2015; 83:1497–1541. [PubMed: 26778846]
- Fan J, Lv J. Sure independence screening for ultra-high dimensional feature space (with discussion). *Journal of the Royal Statistical Society, Ser B*. 2008; 70:849–911.
- Feng X, He X, Hu J. Wild bootstrap for quantile regression. *Biometrika*. 2011; 98:995–999. [PubMed: 23049133]
- Fithian W, Taylor J, Tibshirani R, Tibshirani RJ. Selective sequential model selection. 2015 ArXiv e-prints.
- Genovese C, Jin J, Wasserman L, ZY. A comparison of the lasso and marginal regression. *Journal of Machine Learning Research*. 2012; 13:2107–2143.
- Gregory KB, Carroll RJ, Baladandayuthapani V, Lahiri SN. A twosample test for equality of means in high dimension. *Journal of the American Statistical Association*. 2015; 110:837–849. [PubMed: 26279594]
- G’Sell MG, Wager S, Chouldechova A, Tibshirani R. Sequential selection procedures and false discovery rate control. *Journal of the Royal Statistical Society, Ser B*. 2016; 78:423–444.
- Hall P, Kang KH. Bootstrapping nonparametric density estimators with empirically chosen bandwidths. *The Annals of Statistics*. 2001; 29:1443–1468.
- Hall P, Sheather S. On the distribution of a studentized quantile. *Journal of the Royal Statistical Society, Series B*. 1988; 50:381–391.
- He X, Wang L, Hong HG. Quantile-adaptive model-free variable screening for high-dimensional heterogeneous data. *The Annals of Statistics*. 2013; 41:342–369.
- Hendricks W, Koenker R. Hierarchical spline models for conditional quantiles and the demand for electricity. *Journal of the American Statistical Association*. 1991; 87:58–68.
- Holm S. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*. 1979; 6:65–70.
- Kocherginsky M, He X, Mu Y. Practical confidence intervals for regression quantiles. *Journal of Computational and Graphical Statistics*. 2005; 14:41–55.
- Koenker, R. *Quantile Regression*. Cambridge: Cambridge University Press; 2005.
- Kosorok, MK. *Introduction to Empirical Processes and Semiparametric Inference*. New York: Springer; 2008.
- Kozbur D. Testing-based forward model selection. 2015 ArXiv e-prints.
- Lee JD, Taylor JE. Exact post model selection inference for marginal screening. 2014 arXiv: 1402.5596.
- Leeb H, Pötscher BM. Can one estimate the conditional distribution of postmodel- selection estimators? *The Annals of Statistics*. 2006; 34:2554–2591.
- Li A, Barber RF. Accumulation tests for fdr control in ordered hypothesis testing. 2015 ArXiv e-prints.
- McKeague IW, Qian M. An adaptive resampling test for detecting the presence of significant predictors. *Journal of American Statistical Association*. 2015; 110:1422–1433.
- Powell, JL. Estimation of monotonic regression models under quantile restrictions. In: Barnett, W.Powell, J., Tauchen, G., editors. *Nonparametric and Semiparametric Methods in Econometrics*. Cambridge: Cambridge University Press; 1991.

- Rhee SY, Gonzales MJ, Kantor R, Betts BJ, Ravela J, Shafer RW. Human immunodeficiency virus reverse transcriptase and protease sequence database. *Nucleic Acids Research*. 2003; 31:298–303. [PubMed: 12520007]
- Shao X, Zhang J. Martingale difference correlation and its use in highdimensional variable screening. *Journal of American Statistical Association*. 2014; 109:1302–1318.
- Tibshirani RJ, Taylor J, Lockhart R, Tibshirani R. Exact post- selection inference for sequential regression procedures. *Journal of the American Statistical Association*. 2016; 111:600–620.
- Tukey, JW. Internal working paper. Princeton University; 1989. Higher criticism for individual significances in several tables or parts of tables.
- Zhang X, Yao S, Shao X. Conditional mean and quantile dependence testing in high dimension. *The Annals of Statistics*. 2017 to appear.
- Zheng Q, Peng L, He X. Globally adaptive quantile regression with ultra-high dimensional data. *The Annals of Statistics*. 2015; 43:2225–2258. [PubMed: 26604424]

Appendix

A.1. Proof of Theorem 1

The proof of Theorem 1 follows immediately from the following Lemmas 1 and 4, which establish the joint asymptotic representations of $n^{1/2}\{\hat{\theta}_n(\tau) - \theta_n(\tau)\}/\hat{\sigma}_n(\tau)$ over $\tau \in \mathcal{T}$ for $\beta_0(\tau) \neq 0$ and $\beta_0(\tau) = 0$, separately. Lemma 2 builds an approximate connection between $(\alpha_k(\tau), \theta_k(\tau))$ and $(\alpha_0(\tau), \beta_0(\tau))$ under the local model with $\beta_0(\tau) = 0$, which together with Lemma 3 are needed for proving Lemma 4. The proofs for Lemmas 1–4 are provided in the Supplementary Material.

We first fix some notation. For any vector $\mathbf{v} \in \mathbb{R}^p$, let v_k denote its k th element, and $\mathbf{v}_{(-k)}$ denotes its subset excluding the k th element, for any $k = 1, \dots, p$. For notational simplicity, we write $\bar{k}_n\{\tau, \mathbf{b}_0(\tau)\}$ as $k_n(\tau)$, and omit the argument τ in various expressions such as $M_k(\tau, \cdot)$, $\pi_k(\tau, \cdot)$, $V_k(\tau, \cdot)$, $\mathbf{J}_k(\tau, \cdot)$ and etc. when needed.

Lemma 1

Suppose that assumptions A1–A5 hold. For all τ 's in $\mathcal{T}$ for which $\beta_0(\tau) \neq 0$ and $k_0(\tau)$ is unique, we have $\hat{k}_n(\tau) \xrightarrow{a.s.} k_0(\tau)$ and

$$\frac{n^{1/2}\{\hat{\theta}_n(\tau) - \theta_n(\tau)\}}{\hat{\sigma}_n(\tau)} \xrightarrow{d} \frac{M_{p+k_0(\tau)}\{\beta_0(\tau)\}\pi_{k_0(\tau)}\{\beta_0(\tau)\} - M_{k_0(\tau)}\{\beta_0(\tau)\}\mu_{k_0(\tau)}\{\beta_0(\tau)\}}{V_{k_0(\tau)}\{\beta_0(\tau)\}\sigma_{k_0(\tau)}(\tau)}.$$

Lemma 2

If assumptions A1–A5 hold, we have

$$\begin{pmatrix} \alpha_k(\tau) \\ \theta_k(\tau) \end{pmatrix} = \begin{pmatrix} \alpha_0(\tau) \\ \beta_{n,k}(\tau) \end{pmatrix} + \{\mathbf{J}_k^{-1}(\tau, \mathbf{0})\} \mathbf{A}_k^T(\tau) \beta_{n,(-k)}(\tau) + o(n^{-1/2}) \quad (12)$$

uniformly over τ for which $\beta_0(\tau) = \mathbf{0}$, where $\mathbf{A}_k(\tau) = (E\{f_{\varepsilon(\tau)}(0|\mathbf{X})\mathbf{X}_{(-k)}\}, E\{f_{\varepsilon(\tau)}(0|\mathbf{X})X_k\mathbf{X}_{(-k)}\})$.

Remark 1

For the homoscedastic case where $f_{\varepsilon(\tau)}(\cdot|\mathbf{X})$ is common across $\mathbf{X}$, by Lemma 2, we have that uniformly in $\tau \in \mathcal{T}$ for which $\beta_0(\tau) = \mathbf{0}$,

$$\theta_k(\tau) = \beta_{n,k}(\tau) + \frac{\text{Cov}(X_k, \mathbf{X}_{(-k)}^T) \beta_{(-k)}(\tau)}{\text{Var}(X_k)} + o(n^{-1/2}) = \frac{\text{Cov}(X_k, \mathbf{X}^T) \beta_n(\tau)}{\text{Var}(X_k)} + o(n^{-1/2}).$$

Lemma 3

If assumptions A1–A5 hold, we have

$$\left(n^{1/2} \begin{pmatrix} \hat{\alpha}_k(\tau) - \alpha_0(\tau) \\ \hat{\theta}_k(\tau) - \beta_{n,k}(\tau) \end{pmatrix} \right)_{k=1}^p \xrightarrow{d} \left(\mathbf{J}_k^{-1}(\tau, \mathbf{0}) \left\{ \begin{pmatrix} M_k(\tau, \mathbf{0}) \\ M_{p+k}(\tau, \mathbf{0}) \end{pmatrix} + \mathbf{A}_k^T(\tau) \mathbf{b}_{0,(-k)}(\tau) \right\} \right)_{k=1}^p$$

uniformly in $\tau \in \mathcal{T}$ for which $\beta_0(\tau) = \mathbf{0}$.

Lemma 4

Under assumptions A1–A5, we have

$$\frac{n^{1/2} \{\hat{\theta}_n(\tau) - \theta_n(\tau)\}}{\hat{\sigma}_n(\tau)} \xrightarrow{d} \frac{M_{p+K(\tau)}(\tau, \mathbf{0}) \pi(\tau, \mathbf{0}) - M_{K(\tau)}(\tau, \mathbf{0}) \mu_{K(\tau)}(\tau, \mathbf{0})}{V_{K(\tau)}(\tau, \mathbf{0}) \sigma_{K(\tau)}(\tau)} + \left\{ \frac{\mathbf{C}_{K(\tau)}(\tau)}{V_{K(\tau)}(\tau, \mathbf{0})} - \frac{\mathbf{C}_{\kappa_{\tau}\{\mathbf{b}_0(\tau)\}}(\tau)}{V_{\kappa_{\tau}\{\mathbf{b}_0(\tau)\}}(\tau, \mathbf{0})} \right\}^T \frac{\mathbf{b}_0(\tau)}{\sigma_{K(\tau)}(\tau)}$$

uniformly over $\tau \in \mathcal{T}$ for which $\beta_0(\tau) = \mathbf{0}$, where

$$K(\tau) = \arg\max_{k=1, \dots, p} \{ \mathbf{M}_k(\tau) + \mathbf{B}_k^T(\tau) \mathbf{b}_0(\tau) \}^T \mathbf{J}_k^{-1}(\tau, \mathbf{0}) \{ \mathbf{M}_k(\tau) + \mathbf{B}_k^T(\tau) \mathbf{b}_0(\tau) \}$$

with $\mathbf{M}_k(\tau) = (M_k(\tau, \mathbf{0}), M_{k+p}(\tau, \mathbf{0}))^T$.

A.2. Proof of Corollary 1

The proof is provided in the Supplementary Material.

A.3. Proof of Theorem 2

Let $\mathbb{P}_n^*$ denote bootstrap average, and $\mathbb{P}_n^* = \sqrt{n}(\mathbb{P}_n^* - \mathbb{P}_n)$. Define $\hat{\varepsilon}_n(\tau) = Y - \hat{\alpha}_n(\tau) - \hat{\theta}_n(\tau) X_{\hat{k}_n(\tau)}$. The bootstrapped process $\mathbb{P}_n^*(\tau, \mathbf{b}(\tau))$ is defined as

$$*_n(\tau, \mathbf{b}(\tau)) = \frac{(-\hat{\mu}_{\mathbb{K}_n^*(\tau, \mathbf{b})}(\tau), \hat{\pi}(\tau))^* [\tilde{\mathbf{X}}_{\mathbb{K}_n^*(\tau, \mathbf{b})} \psi_\tau\{\hat{\varepsilon}_n(\tau)\}]}{\hat{V}_{\mathbb{K}_n^*(\tau, \mathbf{b})}(\tau) \hat{\sigma}_{\mathbb{K}_n^*(\tau, \mathbf{b})}(\tau)} + \left(\frac{\hat{\mathbf{C}}_{\mathbb{K}_n^*(\tau, \mathbf{b})}(\tau)}{\hat{V}_{\mathbb{K}_n^*(\tau, \mathbf{b})}(\tau)} - \frac{\hat{\mathbf{C}}_{\mathbb{K}_n(\tau, \mathbf{b})}}{\hat{V}_{\mathbb{K}_n(\tau, \mathbf{b})}(\tau)} \right)^T \frac{\mathbf{b}}{\hat{\sigma}_{\mathbb{K}_n^*(\tau, \mathbf{b})}(\tau)}, \quad (13)$$

where $\mathbb{K}_n(\tau, \mathbf{b}) = \arg \min \mathbb{L}_{n,k}(\tau, \mathbf{b})$, $*_n(\tau, \mathbf{b}) = \arg \max_{k_{n,k}} *_n(\tau, \mathbf{b})$,

$$\mathbb{L}_{n,k}(\tau, \mathbf{b}) = \min_{\alpha, \theta} \mathbb{P}_n[\rho_\tau\{\varepsilon(\tau) + \alpha_0(\tau) + \mathbf{X}^T(\boldsymbol{\beta}_0(\tau) + n^{-1/2}\mathbf{b}) - \alpha - \theta X_k\}],$$

$$\text{and } *_n(\tau, \mathbf{b}) = \left(*_n[\tilde{\mathbf{X}}_k \psi_\tau\{\hat{\varepsilon}_n(\tau)\}] + \hat{\mathbf{B}}_k^T(\tau) \mathbf{b} \right)^T \hat{\mathbf{J}}_k^{-1}(\tau) \left(*_n[\tilde{\mathbf{X}}_k \psi_\tau\{\hat{\varepsilon}_n(\tau)\}] + \hat{\mathbf{B}}_k^T(\tau) \mathbf{b} \right).$$

Let E^M denote expectation conditional on the data, and let P^M be the corresponding probability measure. We will show that

$$I\{|T_n^*(\tau)| > \lambda_n(\tau) \text{ or } |T_n(\tau)| > \lambda_n(\tau)\} \xrightarrow{P^M} I(\boldsymbol{\beta}_0(\tau) \neq \mathbf{0}) \text{ and}$$

$$I\{|T_n^*(\tau)| \leq \lambda_n(\tau)\} I\{|T_n(\tau)| \leq \lambda_n(\tau)\} \xrightarrow{P^M} I\{\boldsymbol{\beta}_0(\tau) = \mathbf{0}\} \text{ for all } \tau \in \mathcal{T} \text{ conditionally (on the data) in probability. This together with Lemmas 5 and 6 below implies the result.}$$

For $k = 1, \dots, p$, the bootstrapped marginal regression coefficients satisfy

$$(\hat{\alpha}_k^*(\tau), \hat{\theta}_k^*(\tau)) = \arg \min_{\alpha, \theta} \mathbb{P}_n^*\{\rho_\tau(Y - \alpha - \theta X_k)\}. \text{ By the first order condition, we have}$$

$\mathbb{P}_n^*[\psi_\tau(Y - \hat{\alpha}_k^*(\tau) - \hat{\theta}_k^*(\tau)X_k)\tilde{\mathbf{X}}_k] = 0$ for $\tau \in \mathcal{T}$. Similarly, by the definition of $(\alpha_k(\tau), \theta_k(\tau))$ in (2), $E[\psi_\tau\{Y - \alpha_k(\tau) - \theta_k(\tau)X_k\}\tilde{\mathbf{X}}_k] = 0$. Under Assumptions listed in Section 2.3, it can be verified that for any $\eta > 0$, the class of functions $\{\psi_\tau(Y - \alpha - \theta X_k)\tilde{\mathbf{X}}_k : \tau \in \mathcal{T}, (\alpha, \theta) \in \mathbb{R}^2, \sup_{\tau \in \mathcal{T}} \|(\alpha, \theta) - (\alpha_k(\tau), \theta_k(\tau))\| \leq \eta\}$ is P-Donsker, and $E\|\psi_\tau(Y - \alpha_n - \theta_n X_k) - \psi_\tau(Y - \alpha_k(\tau) - \theta_k(\tau)X_k)\tilde{\mathbf{X}}_k\|^2 \rightarrow 0$ for any sequence (α_n, θ_n) such that $\|(\alpha_n, \theta_n) - (\alpha_k(\tau), \theta_k(\tau))\| \rightarrow 0$ for all $\tau \in \mathcal{T}$.

Using similar arguments as those in the proof of Theorem 10.16 of Kosorok (2008), we have

$$n^{1/2} \begin{pmatrix} \hat{\alpha}_k^*(\tau) - \alpha_k(\tau) \\ \hat{\theta}_k^*(\tau) - \theta_k(\tau) \end{pmatrix} = \mathbf{J}_k^{-1}\{\tau, \boldsymbol{\beta}_0(\tau)\} n^{1/2} \mathbb{P}_n^*[\psi_\tau\{Y - \alpha_k(\tau) - \theta_k(\tau)X_k\}\tilde{\mathbf{X}}_k] + o_p(1). \quad (14)$$

This together with (S.2) in the Supplementary Material implies that

$$n^{1/2} \begin{pmatrix} \hat{\alpha}_k^*(\tau) - \alpha_k(\tau) \\ \hat{\theta}_k^*(\tau) - \theta_k(\tau) \end{pmatrix} = \mathbf{J}_k^{-1}\{\tau, \boldsymbol{\beta}_0(\tau)\} *_n[\psi_\tau\{Y - \alpha_k(\tau) - \theta_k(\tau)X_k\}\tilde{\mathbf{X}}_k] + o_p(1) \quad (15)$$

for all $\tau \in \mathcal{T}$ for $k = 1, \dots, p$.

When $\beta_0(\tau) = \mathbf{0}$, note that

$$|T_n^*(\tau)| = \left| \frac{n^{1/2} \hat{\theta}_n^*(\tau)}{\hat{\sigma}_{n*}(\tau)} \right| \leq \max_{k=1, \dots, p} \left| \frac{n^{1/2} \{\hat{\theta}_k^*(\tau) - \hat{\theta}_k(\tau)\}}{\hat{\sigma}_k(\tau)} + \frac{n^{1/2} \hat{\theta}_k(\tau)}{\hat{\sigma}_k(\tau)} \right|$$

where $\hat{\sigma}_k^2(\tau)$ is the lower right diagonal element of $\hat{\Sigma}_k(\tau)$, the consistent estimator of $\Sigma_k(\tau)$, which is assumed in A5 to be bounded away from zero for all $\tau \in \mathcal{T}$ and k . This together with (15), bootstrap consistency of the sample mean, Lemma 3, and the condition that $\lambda_n(\tau) \rightarrow \infty$ implies that $|T_n^*(\tau)|/\lambda_n(\tau) = o_{PM}(1)$ conditionally in probability.

When $\beta_0(\tau) \neq \mathbf{0}$, it is easy to verify that $|\theta_n(\tau)| \rightarrow |\theta_{0,k_0}(\tau)| > 0$ under the condition that k_0 is unique, where $(\alpha_{0,k}(\tau), \theta_{0,k}(\tau)) = \arg \min_{a, \theta} E\{\rho_{\tau}(e(\tau) + a_0(\tau) + \mathbf{X}^T \beta_0(\tau) - a - \theta X_k)\}$. Thus

$$\begin{aligned} P^M\{|T_n^*(\tau)| \leq \lambda_n(\tau)\} &= P^M\{n^{1/2}|\{\hat{\theta}_n^*(\tau) - \hat{\theta}_n(\tau)\} + \{\hat{\theta}_n(\tau) - \theta_n(\tau)\} + \{\theta_n(\tau) - \theta_{0,k_0}(\tau)\} + \theta_{0,k_0}(\tau)| \leq \lambda_n(\tau) \hat{\sigma}_{n*}(\tau)\} \\ &\leq P^M\{|\theta_{0,k_0}(\tau)| \leq n^{-1/2} \lambda_n(\tau) \max_{k=1, \dots, p} \{\hat{\sigma}_k(\tau) + |\hat{\theta}_n^*(\tau) - \hat{\theta}_n(\tau)| + |\hat{\theta}_n(\tau) - \theta_n(\tau)| + |\theta_n(\tau) - \theta_{0,k_0}(\tau)|\}\} \end{aligned}$$

tends to zero in probability, where the convergence follows from Lemma 1, Lemma 5 below, and $\lambda_n(\tau) = o(\sqrt{n})$.

Therefore, for all $\tau \in \mathcal{T}$, we have

$$\begin{aligned} E^M |I\{|T_n^*(\tau)| \leq \lambda_n(\tau)\} - I\{\beta_0(\tau) = \mathbf{0}\}| &= E^M |I\{|T_n^*(\tau)| > \lambda_n(\tau)\} - I\{\beta_0(\tau) \neq \mathbf{0}\}| \\ &= P^M\{|T_n^*(\tau)| > \lambda_n(\tau), \beta_0(\tau) = \mathbf{0}(\tau)\} + P^M\{|T_n^*(\tau)| \leq \lambda_n(\tau), \beta_0(\tau) \neq \mathbf{0}\} \\ &= P^M\{|T_n^*(\tau)| > \lambda_n(\tau) | \beta_0(\tau) = \mathbf{0}\} I\{\beta_0(\tau) = \mathbf{0}\} + P^M\{|T_n^*(\tau)| \leq \lambda_n(\tau) | \beta_0(\tau) \neq \mathbf{0}\} I\{\beta_0(\tau) \neq \mathbf{0}\} \end{aligned}$$

tends to zero in probability. This implies that $I\{|T_n^*(\tau)| > \lambda_n(\tau)\} \xrightarrow{PM} I\{\beta_0(\tau) \neq \mathbf{0}\}$ and $I\{|T_n^*(\tau)| \leq \lambda_n(\tau)\} \xrightarrow{PM} I\{\beta_0(\tau) = \mathbf{0}\}$ conditionally in probability. By Lemmas 1 and 3, it is easy to verify that $I\{|T_n(\tau)| \leq \lambda_n(\tau)\} \xrightarrow{P} I\{\beta_0(\tau) = \mathbf{0}\}$ for all $\tau \in \mathcal{T}$. The result of Theorem 2 follows from Slutsky's lemma.

Lemma 5

Suppose the assumptions in Theorem 1 hold. Then $\hat{k}_n^(\tau) \xrightarrow{PM} k_0(\tau)$ conditionally (on the data) a.s. and*

$$\frac{n^{1/2}\{\hat{\theta}_n^*(\tau) - \hat{\theta}_n(\tau)\}}{\hat{\sigma}_{n*}(\tau)} \xrightarrow{d} \frac{M_{p+k_0(\tau)}\{\beta_0(\tau)\}\pi_{k_0(\tau)}\{\beta_0(\tau)\} - M_{k_0(\tau)}\{\beta_0(\tau)\}\mu_{k_0(\tau)}\{\beta_0(\tau)\}}{V_{k_0(\tau)}\{\beta_0(\tau)\}\sigma_{k_0(\tau)}(\tau)}$$

for all $\tau \in \mathcal{T}$ for which $\beta_0(\tau) = \mathbf{0}$, conditionally (on the data) in probability, where $M_j\{\beta_0(\tau)\} = M_j\{\tau, \beta_0(\tau)\}$.

Proof—The proof is provided in the Supplementary Material.

Lemma 6

Suppose all assumptions in Theorem 1 hold. Then ${}^*_n\{\tau, \mathbf{b}_0(\tau)\}$ converges to the same limiting distribution as $\sqrt{n}\{\hat{\theta}_n(\tau) - \theta_n(\tau)\}/\hat{\sigma}_n(\tau)$ for all $\tau \in \mathcal{T}$ for which $\beta_0(\tau) = \mathbf{0}$, conditionally (on the data) in probability.

Proof—The proof is provided in the Supplementary Material.

Figure 1. Rejection rates of QMET at $\tau = 0.5$ in Cases 1–4 against the constant a involved in the bandwidth parameter h_n .

Figure 2.

Left: the training-set p -values (median $\pm$ MAD) for the newly entered predictor at each step for the standard forward quantile regression (FWD) and the proposed forward quantile marginal effect testing (QMET) method across 50 random splits at $\tau = 0.5$ and 0.75 ; right: the corresponding prediction errors (median $\pm$ MAD) in the testing sets.

Figure 3.

Boxplots of $\log_{10}$ susceptibility of a training data versus three predictors (103N, 190A and 230L) that are selected by the forward QMET method at $\tau = 0.75$.

Table 1

Percentages of rejections from different methods in Cases 1–4 with $n = 200$. Case 1 with $\tau = 0.5$, 0.75 and Case 4 with $\tau = 0.5$ correspond to the null model, and the others correspond to the alternative model. The last column is the proposed test across three quantiles 0.25, 0.5 and 0.75.

Case	p	$\tau = 0.5$						$\tau = 0.75$			Mult. τ		
		CPB	AOV	BONF	QMET	CPB	AOV	BONF	QMET	QMET			
1	10	27.8	6.6	5.2	3.8	24.8	6.2	3.2	5.6	2.4			
	100	35.4	0.0	2.4	3.0	37.2	0.0	2.4	4.8	2.8			
	200	39.4		2.2	4.6	34.6		1.8	5.6	2.6			
	400	42.2		2.0	4.4	42.4		2.0	8.0	2.4			
	1000	48.8		1.6	4.6	42.0		1.2	4.8	1.6			
2	10	86.5	61.0	79.6	75.0	78.8	50.6	73.0	64.8	83.8			
	100	85.7	0.0	67.2	67.0	75.4	0.2	55.8	54.4	70.2			
	200	83.2		59.2	61.5	77.2		47.6	59.2	65.6			
	400	84.8		58.8	62.1	74.8		43.0	60.0	70.1			
	1000	85.7		50.6	64.6	75.6		38.2	61.2	65.4			
3	10	94.0	94.2	91.6	85.4	75.0	70.2	70.6	62.0	84.2			
	100	88.2	0.0	75.6	71.8	71.0	0.4	48.8	47.6	72.5			
	200	89.6		74.2	72.0	71.2		52.2	50.2	71.3			
	400	84.4		67.4	71.6	69.4		44.8	49.4	68.8			
	1000	87.4		64.0	74.4	70.2		39.0	52.6	70.8			
4	10	24.0	7.2	6.6	6.0	90.8	47.8	70.8	80.6	77.4			
	100	26.2	0.0	4.4	3.2	86.0	0.0	53.4	61.4	46.8			
	200	25.0		3.0	4.0	86.4		44.0	58.4	43.4			
	400	30.6		2.2	5.8	85.8		41.4	57.0	30.7			
	1000	28.6		2.2	6.6	87.4		38.2	58.2	26.8			

Table 2
Percentages of rejections of the QMET and Bonferroni methods at $\alpha = 0.5$ in Cases 5–6.

<i>p</i>	Case 5					Case 6				
	10	100	200	400	1000	10	100	200	400	1000
QMET	3.6	1.6	2.1	2.3	4.1	65.8	50.0	49.0	53.1	60.5
BONF	4.0	3.2	6.0	2.6	3.2	65.0	46.0	38.6	40.2	34.2

Table 3

The p -values from the forward QMET procedure for predictors 103N, 190A and 230L, and the corresponding estimated coefficients, standard errors and Wald-type p -values by regressing the log susceptibility on all three predictors jointly at $\tau = 0.5$ and 0.75 in the example training set.

τ	Variable	Forward QMET			Multiple quantile regression		
		p -value	Coefficient	Std. error	Wald p -value		
0.5	103N	0.0000	1.5593	0.1122	0.0000		
	190A	0.8525	1.6021	0.5290	0.0028		
	230L	0.7175	1.6372	0.8625	0.0592		
0.75	103N	0.0000	1.7249	0.1827	0.0000		
	190A	0.0125	1.8947	0.1900	0.0000		
	230L	0.0000	2.1871	0.0928	0.0000		