Robust Inference for Time Series Quantile Regression: A Dependent Wild Bootstrap-Based Approach

Zongwu Cai* Wei Long†

April 3, 2026

Abstract

Quantile regression is widely used to study heterogeneous effects, but inference in time series settings remains challenging when regression errors are serially correlated. Building on the dependent wild bootstrap, we develop an inference procedure for linear time series quantile regression that reweights the restricted quantile score with tapered multipliers and employs a one-step bootstrap update together with the HAC-based studentization. The procedure avoids repeated solution of a non-smooth quantile regression problem within each bootstrap draw while targeting the same inferential object as robust HAC testing. Under strong mixing and standard smoothness and bandwidth conditions, we establish asymptotic validity of the bootstrap test and derive its local power under Pitman alternatives. Monte Carlo results indicate improved size control relative to conventional and robust HAC methods, especially under strong dependence, with only modest differences in power. An application to the determinants of U.S. housing prices over the past four decades illustrates the practical usefulness of the method.

Keywords: Time series quantile regression; Dependent wild bootstrap; HAC inference

*Department of Economics, University of Kansas, Lawrence, KS 66045.
†Department of Economics, Tulane University, New Orleans, LA 70118. Corresponding author.

1 Introduction

Quantile regression is a standard tool for studying heterogeneous effects, and its time series applications are broad. Examples include, but not limited to, quantile autoregressive models (Koenker and Xiao, 2004, 2006), value-at-risk and systemic-risk measurement (Engle and Manganelli, 2004; Adrian and Brunnermeier, 2016), directional predictability (Linton and Whang, 2007), structural change (Oka and Qu, 2011), cointegrating relationships (Xiao, 2009), tail dependence in multivariate systems (White et al., 2015), and predictive quantile regressions in asset returns (Lee, 2016; Fan and Lee, 2019; Cai et al., 2023, 2026). In these settings, interest often centers on inference for slope coefficients or other linear restrictions across the conditional distribution. However, when regression disturbance becomes serially correlated, that inference becomes considerably more challenging. Conventional quantile regression standard errors treat the score process as if it were conditionally independent and therefore understate sampling uncertainty when the data are persistent.

Recent work has made important progress along two complementary directions. In covariance estimation, Galvao and Yoon (2024) develop a systematic heteroscedasticity and autocorrelation consistent (HAC) covariance matrix estimator for time series quantile regression, providing a quantile analogue of the robust covariance estimators of Newey and West (1987) and Andrews (1991). Their method combines standard estimation of the Jacobian term (Powell, 1991; Kato, 2012) with a lag-window estimator of the long-run covariance matrix of the quantile score, and it delivers both asymptotic validity and a feasible bandwidth-selection rule. In resampling, Gregory et al. (2018), building on Fitzenberger (1998) and Gregory et al. (2015), develop a smooth extended tapered block bootstrap (SETBB) for dependent quantile regression. Their procedure resamples blocks to preserve serial dependence, tapers the block weights to mitigate edge effects, and smooths individual observations so that the bootstrap more closely reproduces the smooth features entering the limiting distribution of the quantile regression estimator. Their simulation evidence shows noticeable improvements over the unsmoothed moving block bootstrap and over normal approximations based on Powell-type kernel variance estimation. These methodological developments offer credible approaches for time series quantile regression inference that account for serial dependence rather than assuming the score process is independent.

Even with that progress, however, the practical inferential problem is not fully settled. On one hand, the feasible HAC statistic remains a normal approximation based on an estimated long-run variance, and its finite-sample performance depends on how accurately that long-run variance can be estimated from a non-smooth score process. On the other hand, existing block bootstrap procedures for dependent quantile regression are instead designed to approximate the distribution of the quantile regression estimator itself. For testing a linear restriction, that approach typically requires resolving a non-smooth quantile regression problem in every bootstrap replication and selecting additional tuning parameters such as a block length, a taper, and a smoothing bandwidth. These demands become more consequential when serial dependence is strong, when the interested quantile level lies in the tails, or when the specification contains many persistent covariates. In such settings, the long-run variance estimate can be noisy and the bootstrap procedure can be computationally burdensome. The question addressed in this paper is therefore how to improve finite-sample inference for time series quantile regression while preserving the HAC inferential target and avoiding full bootstrap re-estimation.

Our approach is to adapt the dependent wild bootstrap (DWB) to inference on linear restrictions in time series quantile regression with serially correlated disturbances. The bootstrap literature for quantile regression includes i.i.d. procedures (Horowitz, 1998; Feng et al., 2011), cluster-robust score-based procedures (Hagemann, 2017), and block bootstrap methods for dependent quantile regression (Fitzenberger, 1998; Gregory et al., 2018). More generally, Künsch (1989), Paparoditis and Politis (2001), Paparoditis and Politis (2002), and Shao (2010) develop resampling schemes that mimic serial dependence for stationary statistics. Inference on coefficient restrictions in time series quantile regression combines features of several of these studies. Relative to the cross-sectional and clustered wild bootstrap settings, the quantile score is serially correlated and the observed statistic is HAC-studentized. Block bootstrap procedures for dependent quantile regression typically resample the full quantile regression estimator, whereas our objective is to approximate the null distribution of a studentized restriction test while avoiding a fresh solution of a non-smooth quantile regression problem in every bootstrap replication. This focus leads naturally to a bootstrap based on the quantile score evaluated under the null restriction.

The procedure proposed here is built around that null-imposed quantile score. We first estimate the restricted quantile regression, because the bootstrap intends to reproduce the null distribution of the test statistic rather than the unrestricted sampling distribution of the estimator. We then recenter the restricted score to remove finite-sample drift in its sample mean, and perturb it with a dependent multiplier sequence generated by a tapered moving average construction so that the bootstrap perturbation inherits the serial dependence relevant for long-run variance estimation. Instead of resolving a bootstrap quantile regression problem after each perturbation, we use a one-step update motivated by the Bahadur expansion, which captures the leading effect of the bootstrap shock on the coefficient estimator. We finally studentize each bootstrap draw with a Bartlett long-run variance estimate computed from the projected bootstrap score, so the bootstrap statistic is scaled in the same way as the observed restriction test. After the restricted model has been estimated once, a bootstrap draw no longer requires solving a new quantile regression problem; it only requires generating dependent multipliers, applying the linear update, and computing a one-dimensional Bartlett variance estimate. This keeps the procedure computationally light in practice.

The contribution of this study is threefold. First, we develop a dependent wild bootstrap for testing linear restrictions in time series quantile regression. The procedure is built around the null-imposed score and targets the same asymptotic object as feasible HAC inference. In that respect, it differs from block bootstrap methods for dependent quantile regression, which approximate the distribution of the estimator itself, and from HAC procedures that rely on asymptotic normal critical values. Second, we provide the theory needed to validate the bootstrap-t statistic used in practice. We derive a restricted Bahadur expansion and show first-order equivalence between the restricted projected-score statistic and the feasible HAC statistic. We then establish consistency of the observed and bootstrap Bartlett variance estimators, conditional weak convergence of the multiplier statistic under the null, and non-trivial local power under Pitman alternatives. Third, the projected-score representation also yields a data-driven, MSE-motivated plug-in selector for the bootstrap dependence bandwidth. We show that a clipped consistent implementation remains within the admissible class required for first-order validity, so both the HAC bandwidth for feasible studentization and the dependence bandwidth for the bootstrap can be selected from the data. As a result, bootstrap critical values remain asymptotically valid under feasible implementation.

We evaluate the finite-sample performance of the proposed method through Monte Carlo simulations that compare it with conventional quantile regression inference, the robust HAC procedure, and SETBB. When serial dependence is weak, all methods perform similarly. As dependence becomes stronger, the conventional procedure becomes severely oversized, and although the HAC and SETBB procedures reduce that distortion, they still over-reject in the more persistent cases. In our designs, the dependent wild bootstrap is typically closest to the nominal level, particularly in upper-tail cases. After size adjustment, the power differences across methods are generally modest, so the improved size control does not appear to come at a material power cost. The similar pattern also appears in the multivariate model and in models with many covariates. The empirical application studies the determinants of U.S. housing prices using monthly data over 1988–2025. The results show that conventional confidence bands are noticeably narrower than the HAC, SETBB, and bootstrap bands. The HAC and bootstrap bands often lead to similar qualitative conclusions, but the bootstrap bands are usually wider, especially in the tails, and thus yield more cautious conclusion.

The remainder of the paper is organized as follows. Section 2 introduces the model and reviews the robust HAC framework. Section 3 presents the proposed bootstrap procedure and introduces the plug-in rule for the dependence bandwidth. Section 4 states the regularity conditions and develops the asymptotic theory. Section 5 reports the simulation evidence, and Section 6 presents the empirical application. We conclude in Section 7. Technical proofs and supplementary simulation and empirical results are deferred to Online Appendix.

2 Econometric Model and HAC Estimation

2.1 Model Setup

Let y_t be a real-valued dependent variable and $x_t^* = (1, x_t')' \in \mathbb{R}^{p+1}$ be a vector of covariates that includes an intercept. Fix a compact quantile index set $\mathcal{T} = [\underline{\tau}, \bar{\tau}] \subset (0, 1)$. For $\tau \in \mathcal{T}$ and $t = 1, \ldots, T$, we assume that the conditional τ-quantile of y_t given x_t^* is linear $Q_{y_t}(\tau \mid x_t^*) = x_t^{*\prime}\beta_0(\tau)$, where $\beta_0(\tau) \in \mathbb{R}^{p+1}$ is unknown. Define the quantile error $u_t(\tau) = y_t - x_t^{*\prime}\beta_0(\tau)$, so that $Q_{u_t(\tau)|x_t^*}(\tau \mid x_t^*) = 0$ almost surely. Unlike the standard cross-sectional setting, the underlying stochastic process $\{(y_t, x_t)\}_{t\in\mathbb{Z}}$ is allowed to exhibit serial dependence, and the observed sample $\{(y_t, x_t)\}_{t=1}^{T}$ is drawn from that process. We assume throughout that the data arise from a strictly stationary and weakly dependent sequence; the precise moment, mixing, and smoothness conditions will be provided in Assumptions 1–3 in Section 4.

Define $\psi_\tau(u) = \tau - \mathbf{1}(u \leq 0)$, where $\mathbf{1}(\cdot)$ denotes the indicator function, and $\rho_\tau(u) = u\psi_\tau(u)$, which is the so-called check function. The parameter $\beta_0(\tau)$ is estimated by the quantile regression estimator

$$\hat{\beta}(\tau) = \underset{b\in\mathbb{R}^{p+1}}{\operatorname{argmin}} \sum_{t=1}^{T} \rho_\tau(y_t - x_t^{*\prime}b)\,.$$

Under regularity conditions for dependent data, $\sqrt{T}\big(\hat{\beta}(\tau) - \beta_0(\tau)\big) \xrightarrow{d} N\big(0, \Sigma(\tau)\big)$ as $T \to \infty$, where the asymptotic covariance matrix takes the familiar sandwich form $\Sigma(\tau) = D(\tau)^{-1}J(\tau)D(\tau)^{-1}$

(Koenker and Bassett, 1978; Koenker, 2005). The outer term $D(\tau) = E\left[f_{u_t(\tau)|x_t^*}(0 \mid x_t^*)\, x_t^* x_t^{*\prime}\right]$ is the Jacobian term, which depends on the conditional density of the quantile error evaluated at zero and accommodates heteroscedasticity. Its estimation is standard using kernel methods (Powell, 1991; Kato, 2012), and the bandwidth conditions needed for the feasible Jacobian estimator are absorbed into Assumption 6 below. The inner term $J(\tau)$ is the long-run variance of the quantile score process. Define $z_t(\tau) = x_t^* \psi_\tau(u_t(\tau))$ and let $\Gamma_j(\tau) = \mathrm{Cov}(z_t(\tau), z_{t-j}(\tau))$ denote its j-th order autocovariance matrix. Then, $J(\tau) = \sum_{j=-\infty}^{\infty} \Gamma_j(\tau)$. In the cross-sectional i.i.d. case, this simplifies to $\tau(1-\tau)E[x_t^* x_t^{*\prime}]$. However, in the time series setting, $J(\tau)$ depends on the full autocovariance structure of the score process and therefore requires HAC estimation. The central difficulty is that $z_t(\tau)$ involves the non-differentiable indicator $\mathbf{1}(u_t(\tau) \le 0)$, so estimating $J(\tau)$ is more challenging than in mean regression.

For the statistical inference, we focus on inference about a scalar linear restriction $\mathbb{H}_0 : R\beta_0(\tau) = r_0$, where $R \in \mathbb{R}^{1\times(p+1)}$ is a fixed non-zero row vector and r_0 is known.

2.2 Robust HAC Inference

Let us review the HAC covariance matrix estimator of Galvao and Yoon (2024), which serves as the benchmark for the proposed bootstrap procedure in the next section. Under the dependent-data Bahadur expansion for $\hat{\beta}(\tau)$ and the HAC regularity conditions stated below, valid inference requires consistent estimation of the long-run covariance matrix of the quantile score process. In doing so, let $\hat{u}_t(\tau) = y_t - x_t^{*\prime}\hat{\beta}(\tau)$ be the unrestricted residuals and $\hat{z}_t(\tau) = x_t^* \psi_\tau(\hat{u}_t(\tau))$ the unrestricted score vector. For $j = 0, 1, \ldots, T-1$, define $\hat{\Gamma}_j(\tau) = \frac{1}{T}\sum_{t=j+1}^{T} \hat{z}_t(\tau)\, \hat{z}_{t-j}(\tau)'$, and set $\hat{\Gamma}_{-j}(\tau) = \hat{\Gamma}_j(\tau)'$ for $j \ge 1$. Galvao and Yoon (2024) estimate $J(\tau)$ by the lag-window estimator

$$\hat{J}_T(\tau) = \sum_{j=-(T-1)}^{T-1} k\left(\frac{j}{S_T}\right) \hat{\Gamma}_j(\tau),$$

where $k(\cdot)$ belongs to the class $\mathcal{K}$ of even, continuous kernel functions $k : \mathbb{R} \to [-1, 1]$ with $k(0) = 1$, $\int_{-\infty}^{\infty} k^2(z)\, dz < \infty$, and $\int_0^{\infty} \bar{k}(z)\, dz < \infty$ (with $\bar{k}(z) = \sup_{x \ge z} |k(x)|$), as in Andrews (1991), and S_T denotes the bandwidth. This class includes the Bartlett, Parzen, and quadratic spectrum (QS) kernels. The outer term $D(\tau)$ is estimated by the Powell kernel density estimator (Powell, 1991):

$$\hat{D}(\tau) = \frac{1}{Tm_T} \sum_{t=1}^{T} K\left(\frac{\hat{u}_t(\tau)}{m_T}\right) x_t^* x_t^{*\prime},$$

where $K(\cdot)$ is a symmetric kernel and the density estimation bandwidth m_T satisfies $m_T \to 0$ and $Tm_T \to \infty$. Assumption 6 in Section 4 strengthens these kernel-rate conditions so that plug-in Jacobian estimation is $o_p(T^{-1/4})$ and therefore negligible in both the observed and bootstrap studentizers. Kato (2012) establishes consistency and asymptotic normality of $\hat{D}(\tau)$ under conditions similar to those imposed here. The feasible HAC covariance matrix estimator is then

$$\hat{\Sigma}_r(\tau) = \hat{D}(\tau)^{-1} \hat{J}_T(\tau)\, \hat{D}(\tau)^{-1}.$$

The resulting feasible HAC t-statistic for $\mathbb{H}_0 : R\beta_0(\tau) = r_0$ is

$$T_T^r(\tau) = \frac{\sqrt{T}\left(R\hat{\beta}(\tau) - r_0\right)}{\hat{s}_T^r(\tau)} \quad \text{with} \quad \hat{s}_T^r(\tau) = \sqrt{R\hat{\Sigma}_r(\tau)R'}. \tag{1}$$

Under the null hypothesis and standard regularity conditions, the results of Galvao and Yoon (2024) imply $T_T^r(\tau) \xrightarrow{d} N(0,1)$. The rate conditions on S_T and the feasible selector requirement $\hat{S}_T/S_T \xrightarrow{p} 1$ will be included in Assumption 4, and Galvao and Yoon (2024) provide an optimal bandwidth selector in terms of mean squared error (MSE).

For comparison with the bootstrap-side studentizer introduced in Section 3, it is useful to express the observed scale in terms of a scalar projected score so that both the observed and bootstrap statistics can be written through the same generic lag-window functional, even when different kernels or bandwidths are used in implementation. This scalar representation is algebraically convenient for matching the bootstrap denominator, and demeaning the projected score changes the lag-window statistic only by an asymptotically negligible amount under the bandwidth conditions stated below. Define $\hat{g}_{t,u}(\tau) = R\hat{D}(\tau)^{-1}\hat{z}_t(\tau)$ where u denotes quantities formed from the unrestricted estimator. For any scalar series $\{v_t\}_{t=1}^T$, let $\bar{v}_T = T^{-1}\sum_{t=1}^T v_t$ and $v_t^\circ = v_t - \bar{v}_T$. We write

$$\mathcal{L}_T(\{v_t\}; k, S) := \frac{1}{T}\sum_{t=1}^{T}(v_t^\circ)^2 + 2\sum_{j=1}^{T-1} k\left(\frac{j}{S}\right) \frac{1}{T} \sum_{t=j+1}^{T} v_t^\circ v_{t-j}^\circ$$

for the scalar lag-window functional applied to the demeaned series. The associated projected-score studentizer is

$$\tilde{s}_T^r(\tau) = \sqrt{\mathcal{L}_T(\{\hat{g}_{t,u}(\tau)\}_{t=1}^T; k, S_T)},$$

with S_T replaced by $\hat{S}_T$ in the feasible implementation. The Bartlett studentizers used for the observed restricted statistic and for each bootstrap draw in Section 3 are special cases of this functional. For the bootstrap theory, the only property we need is the asymptotic equivalence $|\hat{s}_T^r(\tau) - \tilde{s}_T^r(\tau)| \xrightarrow{p} 0$, which will be formalized in Lemma 1 under Assumptions 1–7 in Section 4.

3 A Dependent Wild Bootstrap-Based Approach

This section describes an inference procedure using the dependent wild bootstrap developed by Shao (2010). The bootstrap reweights the restricted quantile score with a dependent multiplier sequence and is calibrated to a restricted projected-score statistic defined below. The steps described below impose the null restriction, preserve the serial dependence relevant for the long-run variance, and avoid repeated solution of a non-smooth quantile regression problem inside each bootstrap replication (Davidson and MacKinnon, 1999).

Setup and algorithm. Under $\mathbb{H}_0$, the restricted quantile estimator is

$$\hat{\beta}_0^r(\tau) = \underset{b \in \mathbb{R}^{p+1}:\, Rb = r_0}{\operatorname{argmin}} \sum_{t=1}^{T} \rho_\tau(y_t - x_t^{*\prime} b)\,;$$

see Zhao (2000) for the general theory of restricted regression quantiles. In practice, $\hat{\beta}_0^r(\tau)$ can be computed with any equality-constrained quantile regression solver. When the null hypothesis imposes restrictions on a subset of coefficients, the problem reduces to an ordinary quantile regression in the remaining unrestricted coefficients. Let $\hat{u}_{t,0}(\tau) = y_t - x_t^{*\prime}\hat{\beta}_0^r(\tau)$ denote the restricted residuals and $\hat{z}_{t,0}(\tau) = x_t^* \psi_\tau(\hat{u}_{t,0}(\tau))$ the restricted score vector. Under $\mathbb{H}_0$, the population restricted score has mean zero at the true restricted parameter, but the sample average $\bar{z}_0(\tau) = T^{-1}\sum_{t=1}^T \hat{z}_{t,0}(\tau)$ might not be exactly zero after replacing the true parameter by $\hat{\beta}_0^r(\tau)$. We therefore recenter the restricted score as $\tilde{z}_{t,0}(\tau) = \hat{z}_{t,0}(\tau) - \bar{z}_0(\tau)$. Since $\bar{z}_0(\tau) = O_p(T^{-1/2})$, recentering is asymptotically negligible and removes the sample-mean component from the bootstrap numerator. The resulting array $\{\tilde{z}_{t,0}(\tau)\}_{t=1}^T$ is the basic object to which the dependent wild bootstrap is applied. In the construction below, the Jacobian is estimated from the restricted residuals. Denote the resulting estimator by $\hat{D}_0(\tau)$. Under $\mathbb{H}_0$, $\hat{D}_0(\tau)$ and the unrestricted $\hat{D}(\tau)$ are asymptotically equivalent.

Serial dependence in the recentered score is mimicked through a triangular array $\{\xi_t^*\}_{t=1}^T$ of random variables generated independently of the data. Conditional on the sample,

$$E^*(\xi_t^*) = 0, \qquad E^*\left[(\xi_t^*)^2\right] = 1, \quad \text{and} \quad E^*(\xi_t^*\xi_s^*) = a\left(\frac{t-s}{\ell_T}\right) + o_p(1)$$

uniformly in t, s, where $E^*(\cdot)$ denotes expectation conditional on $\{(y_t, x_t)\}_{t=1}^T$. The taper function $a(\cdot)$ is an even covariance-generating kernel with $a(0) = 1$ and $\int |a(u)|\, du < \infty$, and $\ell_T \to \infty$ is a dependence bandwidth that governs how quickly cross-correlations among multipliers decay. The role of ℓ_T is analogous to the block length in tapered block bootstrap methods (Paparoditis and Politis, 2001, 2002): a larger ℓ_T induces stronger dependence among the ξ_t^*, which is needed when the quantile score $z_t(\tau)$ is highly autocorrelated. For concreteness, we adopt Gaussian innovations together with the Bartlett taper, so the multiplier draw is generated from a sequence of i.i.d. mean zero, variance one innovations $\{\zeta_s\}_{s=1}^{T+\ell_T-1}$ by

$$\xi_t^* = \frac{1}{\sqrt{\ell_T}} \sum_{j=0}^{\ell_T-1} \zeta_{t+j} - \frac{1}{T}\sum_{s=1}^{T} \frac{1}{\sqrt{\ell_T}} \sum_{j=0}^{\ell_T-1} \zeta_{s+j}, \qquad t = 1, \ldots, T,$$

which produces a dependent multiplier sequence with Bartlett taper $a(x) = \max(1 - |x|, 0)$. In principle, the same moving-average filter can be applied to other innovation families, such as Rademacher or Mammen two-point draws (Mammen, 1993), provided the resulting multiplier array satisfies Assumption 5. We keep the Gaussian specification in the formal theory and simulations below because it yields convenient covariance calculations and matches the dependent multiplier construction in Shao (2010). The choice of ℓ_T is discussed at the end of this section. A centered symmetric Bartlett construction would remove the need for explicit boundary adjustments. We adopt the simpler moving-average form, since the resulting edge distortion is of order $\ell_T/T = o(1)$ under Assumption 5.

Define the restricted projected score $\hat{g}_{t,0}(\tau) = R\hat{D}_0(\tau)^{-1}\hat{z}_{t,0}(\tau)$, where 0 denotes quantities formed under the null restriction, and define its centered version by

$$\tilde{g}_{t,0}(\tau) = R\hat{D}_0(\tau)^{-1}\tilde{z}_{t,0}(\tau) = \hat{g}_{t,0}(\tau) - T^{-1}\sum_{s=1}^{T} \hat{g}_{s,0}(\tau)$$

and the Bartlett special case of the lag-window functional introduced in Section 2.2

$$\mathcal{L}_T^{\mathrm{B}}(\{v_t\}; \ell) := \mathcal{L}_T(\{v_t\}; k_{\mathrm{B}}, \ell) = \frac{1}{T}\sum_{t=1}^{T}(v_t^\circ)^2 + 2\sum_{j=1}^{\ell}\left(1 - \frac{j}{\ell}\right)\frac{1}{T}\sum_{t=j+1}^{T} v_t^\circ v_{t-j}^\circ,$$

where $k_{\mathrm{B}}(x) = \max(1 - |x|, 0)$, $v_t^{\circ} = v_t - \bar{v}_T$, and $\bar{v}_T = T^{-1}\sum_{t=1}^{T} v_t$. The corresponding observed restricted statistic is

$$T_{T,0}^{r}(\tau) := \frac{T^{-1/2}\sum_{t=1}^{T}\hat{g}_{t,0}(\tau)}{\sqrt{\mathcal{L}_T^{\mathrm{B}}(\{\hat{g}_{t,0}(\tau)\}_{t=1}^{T};\ell_T)}}. \tag{2}$$

This is the observed statistic used in the bootstrap test developed below, and we will show that it is first-order equivalent to the feasible HAC statistic (1) under $\mathbb{H}_0$ in Section 4.

To avoid solving the quantile regression problem in each bootstrap draw, we define the bootstrap statistic directly as the one-step score update suggested by the restricted Bahadur expansion in Proposition 2 in Section 4 (Horowitz, 1998). The bootstrap perturbation and its scalar numerator are

$$\Delta_T^*(\tau) := \hat{D}_0(\tau)^{-1}\frac{1}{T}\sum_{t=1}^{T}\xi_t^*\tilde{z}_{t,0}(\tau), \qquad N_T^*(\tau) = \sqrt{T}\,R\Delta_T^*(\tau).$$

The bootstrap-side projected score $g_t^*(\tau) = \xi_t^* R\hat{D}_0(\tau)^{-1}\tilde{z}_{t,0}(\tau)$ is then used to construct a draw-specific HAC studentizer $\hat{s}_T^*(\tau) = \sqrt{\mathcal{L}_T^{\mathrm{B}}(\{g_t^*(\tau)\}_{t=1}^{T};\ell_T)}$, and the implemented bootstrap statistic is

$$T_T^*(\tau) = N_T^*(\tau)/\hat{s}_T^*(\tau).$$

Thus, the implemented bootstrap is a one-step bootstrap-t statistic: it uses the same centering rule and dependence bandwidth ℓ_T as the restricted observed statistic, with Bartlett weights in both the multiplier covariance structure and the bootstrap studentizer, and avoids auxiliary exact bootstrap re-estimation.

For B independent draws of the multiplier sequence, the bootstrap test is implemented by comparing the bootstrap draws with the observed restricted statistic (2), and the bootstrap p-value is computed by

$$\hat{p}_T(\tau) := P^*\left(|T_T^*(\tau)| \geq \left|T_{T,0}^{r}(\tau)\right|\right) \approx \frac{1}{B}\sum_{b=1}^{B}\mathbf{1}\left(\left|T_{T,b}^*(\tau)\right| \geq \left|T_{T,0}^{r}(\tau)\right|\right),$$

where $T_{T,b}^*(\tau)$ is the bootstrap statistic from the b-th multiplier draw. The null hypothesis is rejected at level α when $\hat{p}_T(\tau) < \alpha$, or equivalently when $|T_{T,0}^{r}(\tau)|$ exceeds the empirical $(1-\alpha)$ quantile of $\{|T_{T,b}^*(\tau)|\}_{b=1}^{B}$. The complete workflow of estimation procedure is summarized in Algorithm 1.

Remark 1. *The adoption of Gaussian innovations and Bartlett taper is a benchmark choice*

Algorithm 1 Dependent Wild Bootstrap

Input: data $\{(y_t, x_t)\}_{t=1}^{T}$, quantile index τ, restriction (R, r_0), dependence bandwidth ℓ_T, number of bootstrap draws B.
Compute the restricted estimator $\hat{\beta}_0^r(\tau)$ subject to $Rb = r_0$, and estimate $\hat{D}_0(\tau)$ from the restricted residuals.
Form the restricted scores $\hat{z}_{t,0}(\tau) = x_t^* \psi_\tau(\hat{u}_{t,0}(\tau))$ and the recentered scores $\tilde{z}_{t,0}(\tau) = \hat{z}_{t,0}(\tau) - \bar{z}_0(\tau)$ for $t = 1, \dots, T$.
Compute the observed restricted projected scores $\hat{g}_{t,0}(\tau) = R\hat{D}_0(\tau)^{-1}\hat{z}_{t,0}(\tau)$ and the observed statistic $T_{T,0}^r(\tau)$.
for $b = 1, \dots, B$ **do**
    Draw $\zeta_{1,b}, \dots, \zeta_{T+\ell_T-1,b}$ as i.i.d. mean-zero, unit-variance innovations (Gaussian in the baseline implementation) and set $\xi_{t,b}^* = \ell_T^{-1/2} \sum_{j=0}^{\ell_T-1} \zeta_{t+j,b} - \bar{\xi}_b^*$ for $t = 1, \dots, T$, where $\bar{\xi}_b^* = T^{-1} \sum_{t=1}^{T} \ell_T^{-1/2} \sum_{j=0}^{\ell_T-1} \zeta_{t+j,b}$.
    Compute $N_{T,b}^*(\tau) = \sqrt{T}\, R\hat{D}_0(\tau)^{-1} T^{-1} \sum_{t=1}^{T} \xi_{t,b}^* \tilde{z}_{t,0}(\tau)$.
    Form $g_{t,b}^*(\tau) = \xi_{t,b}^* R\hat{D}_0(\tau)^{-1}\tilde{z}_{t,0}(\tau)$ and compute $\hat{s}_{T,b}^*(\tau) = \sqrt{\mathcal{L}_T^{\mathrm{B}}(\{g_{t,b}^*(\tau)\}; \ell_T)}$.
    Set $T_{T,b}^*(\tau) = N_{T,b}^*(\tau)/\hat{s}_{T,b}^*(\tau)$.
end for
Output: $\hat{p}_T(\tau) = B^{-1} \sum_{b=1}^{B} \mathbf{1}(|T_{T,b}^*(\tau)| \geq |T_{T,0}^r(\tau)|)$; reject $\mathbb{H}_0$ at level α if $\hat{p}_T(\tau) < \alpha$.

for implementation rather than a theoretical restriction. On the choice of multiplier draw, what matters for first-order validity is the covariance and moment behavior summarized in Assumption 5. In practice, one can also choose filtered two-point innovations, such as Rademacher or Mammen draws, although extending the appendix proofs beyond the Gaussian construction would require replacing the Gaussian moment calculation used in Lemma 1 in the next section. On the choice of taper, Bartlett is natural because it is non-negative, compactly supported, easy to generate through the moving-average filter above, and aligned with the Bartlett studentizer used in the bootstrap. Other taper options such as Parzen and quadratic spectrum are also admissible if they generate a valid covariance array satisfying Assumption 5, and their effect should be primarily finite-sample: stronger induced dependence can improve size control when the score is indeed persistent, whereas in weak-dependence or nearly i.i.d. settings it can change finite-sample size relative to the independent-multiplier benchmark in either direction. In the special i.i.d. case, the natural practical benchmark is the independent multiplier choice $a(x) = \mathbf{1}(x = 0)$*, which corresponds to setting* $\ell_T = 1$ *and reduces the procedure to the familiar i.i.d. wild bootstrap. That choice falls outside our diverging bandwidth asymptotic scheme, so throughout the paper we keep the Bartlett rule as a unified default and view the independent multiplier case as a useful finite-sample robustness check when the sample projected score shows little serial dependence. We conduct robustness checks on these* *choices through simulations in Section 5.*

Choice of ℓ_T. The dependent wild bootstrap introduces its own tuning parameter, namely the dependence bandwidth ℓ_T. It is conceptually different from the HAC bandwidth S_T and the Jacobian bandwidth m_T: while S_T controls long-run variance estimation for the observed statistic, ℓ_T determines how much serial dependence is covered by the multiplier sequence and, through the Bartlett studentizer, how dependence is represented by the bootstrap sampling. In this sense, the bootstrap requires bandwidth selection just as HAC inference does. However, the relevant objectives are different: unlike S_T, which is selected to optimize the long-run variance estimator, ℓ_T is not claimed to be directly optimal for the bootstrap distribution itself. Instead, we choose it using an MSE-motivated proxy based on the scalar projected score, which aligns the dependence window with the serial dependence entering both the multiplier covariance and the bootstrap studentizer. For the first-order theory developed here, the main requirements are that $\ell_T \to \infty$, $\ell_T/T \to 0$, and $\ell_T = o(T^{1/2})$. When the dependence bandwidth is selected from the data, we denote the resulting feasible selector by $\hat{\ell}_T$, which is required to satisfy Assumption 5 in the next section.

Using the scalar projected-score representation, we propose a data-driven, MSE-motivated choice of ℓ_T. Let

$$g_t(\tau) = RD(\tau)^{-1}x_t^*\psi_\tau(u_t(\tau)), \quad \gamma_g(j;\tau) = E[g_t(\tau)g_{t-j}(\tau)], \quad \text{and} \quad \Omega_g(\tau) = \sum_{j=-\infty}^{\infty} \gamma_g(j;\tau).$$

For the Bartlett lag-window estimator applied to $\{g_t(\tau)\}$, standard HAC calculations under the summability conditions in Assumption 7 imply the leading expansion

$$\text{AMSE}_g(\ell) = \frac{V_g(\tau)}{T}\ell + \frac{B_g(\tau)^2}{\ell^2} + o\left(\frac{\ell}{T} + \ell^{-2}\right),$$

where $B_g(\tau) = 2\sum_{j=1}^{\infty} j\,\gamma_g(j;\tau)$ and $V_g(\tau) = 2\sum_{j=-\infty}^{\infty}\gamma_g(j;\tau)^2$. Minimizing the leading term yields

$$\ell_{g,\text{opt}}(\tau) = c_g(\tau)T^{1/3}, \quad \text{and} \quad c_g(\tau) = \left\{\frac{2B_g(\tau)^2}{V_g(\tau)}\right\}^{1/3}.$$

Hence, the $T^{1/3}$ dependence window is the Bartlett-AMSE-optimal rate for the scalar projected score, and the remaining task is to estimate the scale constant $c_g(\tau)$. To construct a feasible rule, let

$$\hat{g}^{\circ}_{t,0}(\tau) = \hat{g}_{t,0}(\tau) - \frac{1}{T}\sum_{s=1}^{T}\hat{g}_{s,0}(\tau), \quad \text{and} \quad \hat{\gamma}_g(j;\tau) = \frac{1}{T}\sum_{t=j+1}^{T}\hat{g}^{\circ}_{t,0}(\tau)\hat{g}^{\circ}_{t-j,0}(\tau),$$

and choose a preliminary truncation sequence $L_T \to \infty$. Define $\hat{B}_T(\tau) = 2\sum_{j=1}^{L_T} j\,\hat{\gamma}_g(j;\tau)$, and let $\hat{V}_T(\tau)$ be any strictly positive pilot estimator of $V_g(\tau)$ built from the same preliminary covariance estimates. We then set

$$\hat{c}_T(\tau) = \Pi_{[\underline{c},\bar{c}]}\left(\left\{\frac{2\hat{B}_T(\tau)^2}{\hat{V}_T(\tau)}\right\}^{1/3}\right), \quad \text{and} \quad \hat{\ell}_T(\tau) = \left\lceil \hat{c}_T(\tau)T^{1/3}\right\rceil, \tag{3}$$

where $\Pi_{[\underline{c},\bar{c}]}(x) = \min\{\bar{c}, \max\{\underline{c}, x\}\}$ clips the pilot constant to a fixed compact interval $0 < \underline{c} < \bar{c} < \infty$. The clipping step keeps the selector inside the admissible class used in the bootstrap theory, while the un-clipped pilot preserves the AMSE interpretation. The formal admissibility result for this plug-in selector is stated in Proposition 1 below.

In practice, a fixed benchmark Bartlett-type rule $\ell_T = \max\{2, \lceil cT^{1/3}\rceil\}$ with $c = 2$ and $\lceil\cdot\rceil$ denoting the ceiling function provides a useful default. Monte Carlo results in Section 5 show that the plug-in and fixed bandwidths deliver broadly similar performance in size control and local power.

4 Asymptotic Theory

In this section, we provide the asymptotic theory for the proposed dependent wild bootstrap-based procedure. For brevity, the technical proofs of all main results are deferred to Appendix A. The assumptions below are stated for a fixed $\tau \in \mathcal{T}^{\circ}$. Extension to a compact range of quantiles requires additional uniform empirical process arguments but follows the same structure.

Assumption 1. *The underlying process $\{(y_t, x_t)\}_{t\in\mathbb{Z}}$ is strictly stationary and strong-mixing with coefficients $\alpha(m)$ satisfying $\sum_{m=1}^{\infty} m^{2+\eta}\alpha(m)^{\eta/(4+\eta)} < \infty$ for some $\eta > 0$. Moreover $E\|x_t^*\|^{4+\eta} < \infty$.*

Assumption 2. *For each τ, $Q_{u_t(\tau)|x_t}(\tau \mid x_t) = 0$ almost surely. The conditional density $f_{u_t(\tau)|x_t}(u \mid x_t)$ exists in a neighborhood of zero, is twice continuously differentiable in u, and its first two derivatives are uniformly bounded in that neighborhood. There exist constants $0 < \underline{f} < \bar{f} < \infty$ such that $\underline{f} \le f_{u_t(\tau)|x_t}(0 \mid x_t) \le \bar{f}$ almost surely, uniformly in $\tau \in \mathcal{T}$.*

Assumption 3. *The matrices $D(\tau)$ and $J(\tau)$ are positive definite for each $\tau \in \mathcal{T}$. In addition, the long-run variance of the projected score $RD(\tau)^{-1}x_t^*\psi_\tau(u_t(\tau))$ is strictly positive.*

Assumption 4. *The HAC kernel $k(\cdot)$ belongs to the class $\mathcal{K}$ defined in Section 2.2: $k : \mathbb{R} \to [-1, 1]$ is even and continuous with $k(0) = 1$, $\int_{-\infty}^{\infty} k^2(z)\, dz < \infty$, and $\int_0^{\infty} \bar{k}(z)\, dz < \infty$, where $\bar{k}(z) = \sup_{x \geq z} |k(x)|$. Here S_T denotes a deterministic target rate, while $\hat{S}_T$ denotes the feasible selector used in practice. The bandwidth sequence S_T satisfies $S_T \to \infty$, $S_T = o(T^{1/2})$, and the feasible selector satisfies $\hat{S}_T / S_T \xrightarrow{p} 1$.*

Assumption 5. *Conditional on the sample, the bootstrap multipliers $\{\xi_t^*\}_{t=1}^T$ have mean zero, variance one, uniformly bounded $(2 + \eta)$ moments, and covariance function $E^*(\xi_t^*\xi_s^*) = a\left(\frac{t-s}{\ell_T}\right) + o_p(1)$ uniformly in t, s, where $a(\cdot)$ is an even taper kernel satisfying $a(0) = 1$ and $\int |a(u)|du < \infty$. The $o_p(1)$ term absorbs the boundary correction in the Bartlett moving-average construction from Section 3, whose finite-sample covariance matches $a((t - s)/\ell_T)$ only up to edge effects. The dependence bandwidth satisfies $\ell_T \to \infty$, $\ell_T /T \to 0$, and the feasible selector is of the form $\hat{\ell}_T = \lceil \hat{c}_T T^\kappa \rceil$ for some deterministic $\kappa \in (0, 1/2)$ and random scale factor $\hat{c}_T$ satisfying $0 < \underline{c} \leq \hat{c}_T \leq \bar{c} < \infty$ with probability approaching one.*

Assumption 6. *The Powell density kernel $K(\cdot)$ used in $\hat{D}(\tau)$ and $\hat{D}_0(\tau)$ is bounded, symmetric, Lipschitz continuous, compactly supported, satisfies $\int K(u)\, du = 1$, $\int uK(u)\, du = 0$, and $0 < \int u^2K(u)\, du < \infty$. The density-estimation bandwidth is of the form $m_T = T^{-\gamma_m}$ for some $\gamma_m \in (1/8, 1/2)$.*

Assumption 7. *For the scalar projected score $g_t(\tau) = RD(\tau)^{-1}x_t^*\psi_\tau(u_t(\tau))$, assume $E|g_t(\tau)|^{4+\eta} < \infty$, $\sum_{j=-\infty}^{\infty}(1 + |j|) \left|E[g_t(\tau)g_{t-j}(\tau)]\right| < \infty$, and the fourth-order cumulants of $\{g_t(\tau)\}$ are absolutely summable. The HAC and multiplier bandwidths satisfy $\max\{S_T, \ell_T\} = o(T^{1/2})$.*

Assumption 1 imposes the weak-dependence and moment conditions used throughout the paper, and is standard in dependent HAC theory for quantile regression. Assumption 2 places the target quantile at zero and requires a locally smooth conditional density that is uniformly bounded and bounded away from zero, which provides local identification and supports Powell-type Jacobian estimation. Assumption 3 rules out singular Jacobian and long-run variance matrices, so both the unrestricted and restricted asymptotic variances are well defined. Assumption 4 collects the kernel and bandwidth conditions needed for consistent feasible HAC studentization. Assumption 5 is the bootstrap analogue: it controls the moments and covariance structure of the dependent multipliers and formalizes the admissible growth of the dependence bandwidth ℓ_T. Assumption 6 contains the kernel and bandwidth restrictions for the Jacobian estimators, which imply

$$\|\hat{D}(\tau) - D(\tau)\| + \|\hat{D}_0(\tau) - D(\tau)\| = O_p\big(m_T^2 + (Tm_T)^{-1/2}\big) = o_p(T^{-1/4}).$$

The change from once to twice continuously differentiable densities in Assumption 2 is used here: it delivers the standard $O(m_T^2)$ bias term for the Powell Jacobian estimator, which together with the variance term $(Tm_T)^{-1/2}$ yields the displayed $o_p(T^{-1/4})$ rate. Assumption 7 then supplies the summability and cumulant conditions needed to control projected-score autocovariances uniformly over the lags entering the observed and bootstrap HAC studentizers. Its $(4+\eta)$ moment restriction is in fact implied by Assumption 1, since $|\psi_\tau| \le 1$, but is stated separately to keep the HAC argument self-contained.

The next proposition records the formal connection between the AMSE-motivated selector in (3) and the feasible-selector conditions used throughout the bootstrap theory.

Proposition 1. *Suppose Assumptions 1, 2, 3, 6, and 7 hold for a fixed τ. If the pilot quantities defined in (3) satisfy $\hat{B}_T(\tau) \xrightarrow{p} B_g(\tau)$ and $\hat{V}_T(\tau) \xrightarrow{p} V_g(\tau)$ with $0 < B_g(\tau) < \infty$ and $0 < V_g(\tau) < \infty$, then the clipped plug-in selector $\hat{\ell}_T(\tau) = \lceil \hat{c}_T(\tau) T^{1/3} \rceil$ satisfies the bandwidth-selector part of Assumption 5 with $\kappa = 1/3$. In particular,*

$$\hat{\ell}_T(\tau) \to \infty, \qquad \hat{\ell}_T(\tau)/T \to 0, \quad \text{and} \quad \hat{\ell}_T(\tau)/T^{1/3} = O_p(1).$$

With these assumptions in place, the first step in the asymptotic analysis is to derive the restricted Bahadur expansion for the sample estimator. Under the strong-mixing framework of Galvao and Yoon (2024), Assumptions 1, 2, and 6 give the unrestricted expansion, and the next proposition shows that the restricted counterpart under both the null and local alternatives follows by combining that expansion with a first-order linearization of the equality-constrained Karush-Kuhn-Tucker (KKT) system.

Proposition 2. *Suppose Assumptions 1–6 hold and write $\mathbb{G}_T(\tau) = \frac{1}{\sqrt{T}} \sum_{t=1}^T x_t^* \psi_\tau(u_t(\tau))$. Then the unrestricted estimator satisfies $\sqrt{T}\big(\hat{\beta}(\tau) - \beta_0(\tau)\big) = D(\tau)^{-1}\mathbb{G}_T(\tau) + o_p(1)$. Under $\mathbb{H}_0 : R\beta_0(\tau) = r_0$, the restricted estimator satisfies $\sqrt{T}\big(\hat{\beta}_0^r(\tau) - \beta_0(\tau)\big) = A_R(\tau)\mathbb{G}_T(\tau) + o_p(1)$, where $A_R(\tau) = D(\tau)^{-1} - D(\tau)^{-1}R'\big(RD(\tau)^{-1}R'\big)^{-1}RD(\tau)^{-1}$. Under a contiguous local alternative sequence $\mathbb{H}_{1,T} : \; R\beta_{0,T}(\tau) = r_0 + \delta/\sqrt{T}$ with fixed drift $\delta \in \mathbb{R}$ and $\|\beta_{0,T}(\tau) - \beta_0(\tau)\| = O(T^{-1/2})$, let $u_{t,T}(\tau) = y_t - x_t^{*\prime}\beta_{0,T}(\tau)$ and $\mathbb{G}_{T,1}(\tau) = \frac{1}{\sqrt{T}} \sum_{t=1}^T x_t^* \psi_\tau(u_{t,T}(\tau))$. Then,* *the unrestricted estimator satisfies*

$$\sqrt{T}\big(\hat{\beta}(\tau) - \beta_{0,T}(\tau)\big) = D(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + o_p(1),$$

and hence,

$$\sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big) = \delta + RD(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + o_p(1).$$

Moreover, the restricted estimator satisfies

$$\sqrt{T}\big(\hat{\beta}_0^r(\tau) - \beta_{0,T}(\tau)\big) = A_R(\tau)\mathbb{G}_{T,1}(\tau) - D(\tau)^{-1}R'\big(RD(\tau)^{-1}R'\big)^{-1}\delta + o_p(1).$$

Proposition 2 provides the linear expansion for the statistic's numerator. The next result is a scalar-restriction corollary of the point wise HAC theory in Galvao and Yoon (2024): it combines that expansion with the HAC covariance approximation to establish the null limit of the feasible observed statistic. Write $P^*(\cdot) = P(\cdot \mid \{(y_t, x_t)\}_{t=1}^T)$ for the conditional bootstrap law, and let $\sigma_R^2(\tau) = R\Sigma(\tau)R'$.

Theorem 1. *Suppose Assumptions 1, 2, 3, 4, 6, and 7 hold. Then, under* $\mathbb{H}_0 : R\beta_0(\tau) = r_0$,

$$T_T^r(\tau) \xrightarrow{d} N(0,1).$$

Theorem 1 establishes the first-order null distribution of the observed feasible HAC statistic (1). To validate the bootstrap-t procedure, however, it remains to show that the observed and bootstrap studentizers estimate the same asymptotic scale.

Lemma 1. *Suppose Assumptions 1, 2, 3, 4, 6, and 7 hold. Then,*

$$\tilde{s}_T^r(\tau) \xrightarrow{p} \sigma_R(\tau).$$

If the observed statistic is reported with the matrix-based studentizer $\hat{s}_T^r(\tau) = \sqrt{R\hat{\Sigma}_r(\tau)R'}$, *then,*

$$|\hat{s}_T^r(\tau) - \tilde{s}_T^r(\tau)| \xrightarrow{p} 0$$

as well, so either observed studentizer can be used in the theorem statements. In addition, if Assumption 5 holds and $\hat{s}_T^*(\tau)$ *is constructed from the Bartlett functional* $\mathcal{L}_T^{\mathrm{B}}$ *applied to* $\{g_t^*(\tau)\}_{t=1}^T$ *with the same centering, normalization, and bandwidth* ℓ_T *as in the restricted* *observed statistic $T^r_{T,0}(\tau)$, then, for every $\epsilon > 0$,*

$$P^*(|\hat{s}^*_T(\tau) - \sigma_R(\tau)| > \epsilon) \xrightarrow{p} 0.$$

Consequently,

$$P^*(|\hat{s}^*_T(\tau) - \tilde{s}^r_T(\tau)| > \epsilon) \xrightarrow{p} 0.$$

By Assumption 6 and the triangle inequality, the same conclusion holds with $\tilde{s}^r_T(\tau)$ replaced by $\hat{s}^r_T(\tau)$.

Lemma 1 is the key bridge between the projected-score covariance approximation implied by Assumption 7 and the bootstrap validity result below. With the numerator expansion from Proposition 2 and the common-scale conclusion of the lemma, we can now compare the studentized bootstrap statistic directly with the observed statistic under the null hypothesis.

Theorem 2. *Suppose Assumptions 1– 7 hold. Then, under $\mathbb{H}_0$,*

$$\sup_{z\in\mathbb{R}} \left| P^*(T^*_T(\tau) \le z) - P\big(T^r_{T,0}(\tau) \le z\big) \right| \xrightarrow{p} 0.$$

Remark 2. *Theorem 2 shows that the conditional bootstrap law consistently estimates the sampling law of the finite-sample statistic actually used for testing, and therefore directly justifies bootstrap p-values and bootstrap critical values formed from $T^r_{T,0}(\tau)$. The result is first-order: the proof works by showing that the bootstrap and observed laws converge to the same Gaussian limit, without claiming any higher-order refinement relative to normal critical values.*

Corollary 1. *Under the conditions of Theorem 2,*

$$\sup_{z\in\mathbb{R}} |P^*(T^*_T(\tau) \le z) - \Phi(z)| \xrightarrow{p} 0.$$

Corollary 1 is only the asymptotic normal consequence of Theorem 2; the theorem itself is the result that directly validates the bootstrap test. An immediate implication of Theorem 2 is that bootstrap critical values deliver asymptotically correct size.

Corollary 2. *Let $c^*_{1-\alpha}(\tau)$ denote the conditional $(1-\alpha)$ bootstrap quantile of $|T^*_T(\tau)|$. Under the conditions of Theorem 2,*

$$P\big(\left|T^r_{T,0}(\tau)\right| > c^*_{1-\alpha}(\tau)\big) \to \alpha.$$

After establishing validity under the null hypothesis, we turn to the behavior of the procedure under Pitman local alternatives.

Lemma 2. *Suppose Assumptions 1– 7 hold and let* $\mathbb{H}_{1,T} : R\beta_{0,T}(\tau) = r_0 + \delta/\sqrt{T}$ *be a contiguous Pitman sequence with* $\delta \in \mathbb{R}$ *and* $\|\beta_{0,T}(\tau) - \beta_0(\tau)\| = O(T^{-1/2})$. *Define*

$$N_T(\tau) = \sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big), \quad N_{T,0}(\tau) = \frac{1}{\sqrt{T}}\sum_{t=1}^{T}\hat{g}_{t,0}(\tau), \quad s_{T,0}(\tau) = \sqrt{\mathcal{L}_T^{\mathrm{B}}(\{\hat{g}_{t,0}(\tau)\}_{t=1}^{T}; \ell_T)}.$$

Then,

$$\|\hat{D}_0(\tau) - D(\tau)\| = o_p(T^{-1/4}), \quad N_{T,0}(\tau) = N_T(\tau) + o_p(1), \quad \hat{s}_T^r(\tau) \xrightarrow{p} \sigma_R(\tau), \quad s_{T,0}(\tau) \xrightarrow{p} \sigma_R(\tau).$$

In addition,

$$\sup_{z\in\mathbb{R}} |P^*(T_T^*(\tau) \le z) - \Phi(z)| \xrightarrow{p} 0, \qquad c_{1-\alpha}^*(\tau) \xrightarrow{p} z_{1-\alpha/2}.$$

The next theorem then combines the unrestricted local drift with Lemma 2 to obtain the bootstrap test's asymptotic power function.

Theorem 3. *Suppose Assumptions 1– 7 hold and consider a contiguous Pitman sequence of local alternatives* $\mathbb{H}_{1,T} : R\beta_{0,T}(\tau) = r_0 + \delta/\sqrt{T}$ *with* $\delta \in \mathbb{R}$ *and* $\|\beta_{0,T}(\tau) - \beta_0(\tau)\| = O(T^{-1/2})$. *Then,*

$$T_T^r(\tau) \xrightarrow{d} N(\delta/\sigma_R(\tau), 1)\,, \qquad T_{T,0}^r(\tau) - T_T^r(\tau) = o_p(1).$$

Moreover, if $c_{1-\alpha}^*(\tau)$ *denotes the conditional* $(1-\alpha)$ *bootstrap quantile of* $|T_T^*(\tau)|$, *then* $c_{1-\alpha}^*(\tau) \xrightarrow{p} z_{1-\alpha/2}$. *Hence the bootstrap test based on* $c_{1-\alpha}^*(\tau)$ *has asymptotic power*

$$1 - \Phi\big(z_{1-\alpha/2} - \delta/\sigma_R(\tau)\big) \; + \; \Phi\big(-z_{1-\alpha/2} - \delta/\sigma_R(\tau)\big)\,.$$

In particular, this limit equals α *at* $\delta = 0$ *and is strictly larger than* α *for every* $\delta \neq 0$.

Remark 3. *The common local drift behind Theorem 3 follows from the same linearization used under the null. Let* $u_{t,T}(\tau) = y_t - x_t^{*\prime}\beta_{0,T}(\tau)$ *and* $\mathbb{G}_{T,1}(\tau) = \frac{1}{\sqrt{T}}\sum_{t=1}^{T} x_t^*\psi_\tau(u_{t,T}(\tau))$. *Then*

$$\sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big) = \delta + RD(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + o_p(1),$$

while Proposition 2 gives

$$\sqrt{T}\big(\hat{\beta}_0^r(\tau) - \beta_{0,T}(\tau)\big) = A_R(\tau)\mathbb{G}_{T,1}(\tau) - D(\tau)^{-1}R'\big(RD(\tau)^{-1}R'\big)^{-1}\delta + o_p(1).$$

Consequently,

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\hat{g}_{t,0}(\tau) = RD(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + \delta + o_p(1) = \sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big) + o_p(1),$$

so the restricted projected-score statistic and the feasible HAC statistic continue to share the same deterministic shift under Pitman local alternatives.

Finally, under the assumptions maintained in the corresponding theorem, the preceding results continue to hold when the tuning parameters are selected from the data, provided the feasible selectors satisfy the same rate conditions as their deterministic targets.

Corollary 3. *Under the assumptions of Theorems 1–3, if the selectors $\hat{S}_T$ and $\hat{\ell}_T$ are measurable functions of the sample satisfying Assumptions 4 and 5, then the conclusions of Theorems 1–3 continue to hold with these tuning parameters in place of deterministic sequences.*

5 Numerical Studies

In this section, we conduct Monte Carlo simulations to compare the finite-sample performance of the proposed bootstrap procedure in both univariate and multivariate models.

Univariate model. For the univariate scenario, we consider two data-generating processes (DGPs):

Model 1. The first design is a homoscedastic location-shift model formulated as $y_t = \beta_0 + \beta_1 x_t + u_t$, where the regressor x_t follows $x_t = \rho_x x_{t-1} + \varepsilon_t$ with $\varepsilon_t \sim N(0, 1-\rho_x^2)$, and the latent disturbance u_t follows $u_t = \rho u_{t-1} + v_t$ with $v_t \sim N(0, 1-\rho^2)$. To initialize the disturbance, set $u_1 = \rho u_0 + v_1$ with $u_0 \sim N(0,1)$. Throughout, let $\rho_x = 0.8$ and $\rho \in \{0, 0.5, 0.9\}$, so that the disturbance is either serially uncorrelated or mildly or strongly correlated. The baseline coefficients are $(\beta_0, \beta_1) = (0, 1)$, so the null hypothesis is $\mathbb{H}_0 : \beta_1(\tau) = 1$. We study empirical size at the 5% nominal level for $\rho \in \{0, 0.5, 0.9\}$ and sample sizes $T \in \{100, 200, 500, 1000\}$.

Model 2. The second design keeps the same linear quantile regression specification $y_t = \beta_0 + \beta_1 x_t + u_t$ and the same AR(1) regressor process $x_t = \rho_x x_{t-1} + \varepsilon_t$ with $\rho_x = 0.8$, but replaces the homoscedastic disturbance by a conditionally heteroscedastic AR(1)-GARCH(1,1) error. Specifically,

$$u_t = \rho u_{t-1} + \eta_t, \qquad \eta_t = \sigma_t z_t, \qquad \sigma_t^2 = \omega + \alpha\eta_{t-1}^2 + \beta\sigma_{t-1}^2,$$

where $z_t \sim N(0,1)$. We let the GARCH parameters $(\alpha, \beta, \omega) = (0.05, 0.90, 0.05)$ and adopt a burn-in of 500 observations before retaining the final sample. As in Model 1, we set $(\beta_0, \beta_1) = (0, 1)$, consider $\rho \in \{0, 0.5, 0.9\}$ and sample sizes $T \in \{100, 200, 500, 1000\}$, and evaluate size under the null hypothesis $\mathbb{H}_0 : \beta_1(\tau) = 1$ for each $\tau \in \{0.5, 0.9\}$.

For each combination of (T, ρ, τ), we report the rejection frequencies of the conventional QR test based on the `nid` covariance estimator from the R package `quantreg`[1], the feasible HAC estimator using the quadratic spectrum kernel together with the AR(1)-based plug-in bandwidth rule of Galvao and Yoon (2024) with estimated dependence parameters (feasible HAC), the corresponding HAC estimator using the true autocorrelation coefficient implied by the DGP in bandwidth selection (infeasible HAC), the smooth extended tapered block bootstrap estimator (SETBB), and the proposed dependent wild bootstrap procedure (DWB). Each design is simulated with 1000 Monte Carlo replications. The two HAC estimators employ the quadratic spectrum kernel. For SETBB, we use 499 bootstrap replications, the Sheather-Jones selector for the smoothing bandwidth as in Gregory et al. (2018), and taper parameter $c = 0.43$ as in Paparoditis and Politis (2001) and Shao (2010). We follow Gregory et al. (2018) and choose the block length using the nonparametric plug-in method developed by Lahiri et al. (2007). This block size, once selected, remains fixed for all bootstrap resamples.[2] For DWB, the null hypothesis is imposed by fixing the slope at its benchmark value and re-estimating the intercept under the restricted quantile regression, and the recentered restricted score is then perturbed by a Bartlett-type dependent multiplier sequence with 499 dependent multiplier draws. To studentize the bootstrap statistic with a Bartlett long-run variance estimator using bandwidth $\ell_T = \max\left\{2, \lceil cT^{1/3} \rceil\right\}$, we consider both the fixed scale constant $c = 2$ (fixed DWB) and the plug-in selector (3), which estimates c from the data and bounds it between 0.5 and 5 (plug-in DWB).

Figure 1 compares size control across the six methods and shows similar patterns in both panels. When serial dependence is absent ($\rho = 0$), all methods are close to the nominal 5% level, with SETBB slightly undersized. As dependence strengthens ($\rho = 0.5$ and 0.9), the conventional test becomes increasingly oversized. The feasible and infeasible HAC procedures substantially reduce the size distortion, but they still over-reject under stronger serial dependence. Their behavior is nearly indistinguishable throughout, indicating that the data-driven

[1] `nid` presumes local (in τ) linearity (in x) of the conditional quantile function and computes a Huber sandwich estimate using a local estimate of the sparsity. See `http://www.econ.uiuc.edu/~roger/research/rq/FAQ` for more details about this covariance estimator.

[2] SETBB is implemented by the R package `QregBB`.

optimal bandwidth tracks the oracle benchmark closely, consistent with Galvao and Yoon (2024). By contrast, SETBB and DWB are often among the methods closest to the nominal level when $\rho = 0.5$, but SETBB becomes more noticeably oversized once dependence is very strong, especially in Model 2, whereas DWB is the least oversized method in the strong-dependence, upper-tail case ($\rho = 0.9$ and $\tau = 0.9$). The comparison between DWB with fixed $c = 2$ and DWB with the plug-in selector yields very similar rejection frequencies overall, with neither implementation dominating uniformly. Appendix Table A1 reports the average block length across designs and shows that the plug-in selector chooses larger dependence windows as persistence increases.

For the power comparison, we move the true slope coefficient away from the null by setting $\beta_1 = 1 + \delta$ and compare *size-adjusted* rejection frequencies as $|\delta|$ increases. Size adjustment is essential, because without first equalizing null rejection rates, a test that over-rejects under the null would mechanically appear more powerful even if it were not genuinely more effective at detecting alternatives. Figure 2 presents the resulting power curves with sample size $T = 1000$. As expected, power rises toward one for all six methods as $|\delta|$ increases. After size adjustment, the six curves are usually quite close: the two HAC procedures and the bootstrap are often nearly indistinguishable, while SETBB is sometimes more conservative in the upper-tail cases. As serial dependence strengthens, however, power declines in the sense that a larger deviation from the null is required to attain a given rejection probability. This effect is particularly pronounced when $\rho = 0.9$, since strong serial dependence and volatility clustering jointly amplify persistence in the score process and further reduce the effective information in the sample. Power is lower in Model 2 than in Model 1 when $\rho = 0.9$, where conditional heteroscedasticity raises the long-run variance of the quantile score and thereby reduces the signal-to-noise ratio of the studentized statistic under local alternatives. Both fixed and plug-in DWB remain competitive in strongly dependent upper-tail scenarios and, taken together with the size results, perform well relative to the HAC and SETBB alternatives.

In Appendices B.2–B.4, we report additional checks on the robustness of DWB when (1) the Gaussian disturbance in Models 1 and 2 is replaced by standardized left-skewed heavy-tailed t draws, (2) the data $\{(y_t, x_t)\}_{t=1}^{T}$ are truly i.i.d., and (3) the multipliers in the dependent wild bootstrap are drawn from the Rademacher distribution with either the Parzen or quadratic spectrum taper. The results suggest that both fixed and plug-in DWB continue to show similar patterns across these scenarios.

Multivariate model. Next, we consider a multivariate scenario that extends Model 1 to

$$y_t = \beta_0 + \beta_1 x_{1t} + \beta_2 x_{2t} + \beta_3 x_{3t} + u_t,$$

where the disturbance u_t still follows the same AR(1) process as in the univariate model. To make the setup empirically relevant, we let the three regressors follow independent AR(1) processes with autoregressive coefficients 0.95, 0.85, and 0.53, respectively. All remaining simulation settings are unchanged, except that we test each slope coefficient separately under the null hypotheses $\mathbb{H}_0 : \beta_j(\tau) = 1$ for $j = 1, 2, 3$.

Figure 3 compares size control and reveals patterns similar to those in the univariate model. When $\rho = 0$, all methods remain reasonably close to nominal size, although SETBB is slightly undersized while the other methods' size distortions are slightly larger at $\tau = 0.9$ than at the median when sample size is small. When $\rho = 0.5$ and 0.9, the conventional test becomes substantially oversized, the feasible and infeasible HAC estimators reduce but do not eliminate that distortion, and SETBB and DWB are the methods closest to nominal size across the three coefficient tests. The feasible and infeasible HAC estimators are again nearly indistinguishable, while the relative advantage of SETBB and DWB is most visible when regressor's persistence is high (panels (a) and (b)). The size distortions are largest when testing coefficients associated with the most persistent regressors, especially β_1, and noticeably smaller for β_3, whose regressor is less persistent. In sum, the multivariate results suggest that DWB's size-control gains are not an artifact of the univariate model. Appendix Figure A3 presents the size-adjusted rejection frequencies for testing $\mathbb{H}_0 : \beta_j(\tau) = 1$ for $j = 1, 2, 3$ as the true coefficient $1 + \delta$ moves away from the null. As in the univariate model, all six methods' power curves track each other closely and exhibit similar local-power performance.

Many covariates. Galvao and Yoon (2024) briefly examine the performance of the HAC estimators in models with additional regressors and find that, for a fixed sample size T, size control deteriorates as the number of covariates increases. They attribute this pattern to a downward bias in the HAC standard errors when the regressor dimension becomes larger, and find the commonly used finite-sample correction method, multiplying the HAC variance estimator by $T/(T - p - 2)$ where p denotes the number of covariates, does not fix the over-rejection issue. Similar finite-sample deterioration has also been documented in conditional mean regression models with many controls (see, e.g., Cattaneo et al., 2018; Jochmans, 2022).

Following Galvao and Yoon (2024), we extend Models 1 and 2 by augmenting the specification with z_t, a $p \times 1$ vector of additional covariates. Each component of z_t follows an AR(1) process analogous to that of x_t, with autoregressive coefficients drawn independently from the uniform distribution on $[0.5, 0.9]$, while all other settings of the design remain the same. As in their setup, we fix the sample size at $T \in \{500, 1000\}$ and consider $p \in \{5, 10, 20, 40\}$. Figure 4 reports the rejection frequencies of the competing methods as p varies for each combination of (T, ρ, τ). In both models, the empirical sizes of the conventional test and the two HAC procedures generally increase with p, and this deterioration becomes more pronounced as serial dependence strengthens, particularly when $\rho = 0.9$. For a fixed p, size control improves slightly as T increases. These findings are broadly consistent with Galvao and Yoon (2024); see Section S.2.3.4 of their Online Appendix. SETBB exhibits more stable size control around the median and under weak or moderate dependence. When dependence becomes strong, however, its rejection frequencies rise more noticeably in the upper tail, especially in Model 2, even though they remain less sensitive to p than those of the conventional and HAC procedures. By contrast, both fixed and plug-in DWB maintain relatively stable rejection frequencies across p and often stays closer to the nominal 5% level, suggesting that DWB's finite-sample advantages may extend to empirically relevant settings with many covariates. A plausible explanation for this difference in size control, as p grows relative to T, is that DWB calibrates the critical value from the finite-sample distribution of the restricted, self-studentized statistic itself, using the scalar projected score rather than relying solely on the asymptotic normal approximation together with a HAC variance estimator constructed from the full high-dimensional score vector. Imposing the null restriction and recentering the score may further stabilize the bootstrap by removing finite-sample drift in the restricted score before resampling. However, we emphasize that the theory developed in this paper is stated for fixed p, and a formal treatment of high-dimensional quantile regression merits a separate study in the future.

6 Application: U.S. Housing Prices

In an empirical application, we study the determinants of U.S. national house price appreciation using six monthly series obtained from the Federal Reserve Bank of St. Louis (https://fred.stlouisfed.org/). The dependent variable is the year-on-year percent change in the Case-Shiller National Home Price Index (CSUSHPISA), which measures broad housing market appreciation across the country. We consider five regressors: real disposable personal income growth (`DSPIC96`, year-on-year percent change), the 30-year fixed mortgage rate (`MORTGAGE30US`, monthly average, in percent), the months' supply of houses (`MSACSR`), building permits growth (`PERMIT`, year-on-year percent change), and the civilian unemployment rate (`UNRATE`, in percent). Their role in housing markets has been well documented in the literature (see, e.g., Himmelberg et al., 2005; Knoll et al., 2017; Glaeser and Gyourko, 2018). The sample spans January 1988 through December 2025, yielding $T = 456$ monthly observations. The year-on-year percent change construction for house prices, income, and permits involves twelve-month overlapping windows and therefore induces mechanical MA(11) serial correlation in both the dependent variable and the error process. This overlapping-window structure directly motivates the use of robust inference rather than conventional standard errors.

We estimate the quantile regression

$$\begin{aligned} Q_{y_t}(\tau \mid x_t) = \beta_0(\tau) + \beta_1(\tau)\,\Delta^{12}\log(\mathrm{DPI}_t) + \beta_2(\tau)\,\mathrm{MORTGAGE}_t \\ + \beta_3(\tau)\,\mathrm{SUPPLY}_t + \beta_4(\tau)\,\Delta^{12}\log(\mathrm{PERMIT}_t) + \beta_5(\tau)\,\mathrm{UNRATE}_t \end{aligned}$$

at equally spaced quantile levels $\tau \in \{0.05, 0.10, 0.15, \ldots, 0.95\}$. The coefficient $\beta_j(\tau)$ gauges how the j-th regressor shifts the τ-th quantile of housing price growth, allowing the effect to differ across the distribution of housing market outcomes. Figure 5 plots the coefficient curve $\hat{\beta}_j(\tau)$ together with 95% pointwise confidence bands. The conventional bands use the usual `nid` standard error (shaded area), the feasible HAC bands use the QS kernel, the fixed and plug-in DWB bands use 999 multiplier draws with the rule $\ell_T = \lceil 2T^{1/3} \rceil$ and the data-driven bandwidth selector, and the SETBB bands are percentile bands based on 999 replications.

Comparing the confidence bands in Figure 5 yields four main observations. First, the `nid` bands are generally much narrower than the four robust alternatives, which is consistent with conventional standard errors understating uncertainty in a persistent monthly time series. For real disposable income growth, for example, the `nid` bands suggest a negative effect over much of the distribution, but that conclusion weakens once the robust bands are used, especially at several middle quantiles. Second, the feasible HAC and the two DWB procedures are much closer to one another than any of them is to SETBB, and they usually support the same broad qualitative conclusions. In particular, the DWB bands are often slightly wider at quantiles away from the tails. Third, fixed DWB and plug-in DWB perform very similarly: their bands nearly overlap and they draw the same conclusion throughout. Fourth, SETBB often produces the widest bands, especially below the median. Overall, the four robust procedures all point to more cautious conclusions than the conventional method.

In Appendix C, we revisit the application in Galvao and Yoon (2024), who study the relationship between electricity spot prices and fundamental supply and demand variables such as natural gas prices and solar and wind production forecasts. The overall comparison of the five procedures is broadly similar.

7 Conclusion

This study develops a dependent wild bootstrap procedure for inference in linear time series quantile regression with serially correlated disturbances. By combining restricted quantile scores, dependent multipliers, a one-step bootstrap update, and HAC-based studentization, the procedure targets the same inferential object as feasible HAC testing while avoiding repeated solution of a non-smooth quantile regression problem within each bootstrap draw. We establish asymptotic validity under strong mixing and standard smoothness and bandwidth conditions, and we show that the bootstrap test retains nontrivial local power under Pitman alternatives. The scalar projected-score formulation also yields a data-driven plug-in selector for the bootstrap dependence bandwidth, and the clipped version of that selector remains within the same first-order validity framework. The Monte Carlo evidence indicates that, relative to the conventional and robust HAC- and smooth extended tapered block bootstrap-based competitors, the proposed dependent wild bootstrap tends to provide better size control when dependence is strong, with only modest power differences. The empirical application to U.S. housing prices shows that accounting for serial dependence can materially affect inference, with the bootstrap often yielding wider confidence bands than conventional quantile regression methods. Extending the method to more challenging settings such as nonstationary or high-dimensional time series would be a natural direction for future work.

References

Adrian, T. and Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7):1705–1741.

Andrews, D. W. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, 59(3):817–858.

Cai, Z., Chen, H., and Liao, X. (2023). A new robust inference for predictive quantile regression. *Journal of Econometrics*, 234(1):227–250.

Cai, Z., Chen, Y., Hong, S., and Tsvetano, D. (2026). Unified inference for predictive mean and quantile regressions via empirical likelihood. *Working Paper*, Department of Economics, University of Kansas.

Cattaneo, M. D., Jansson, M., and Newey, W. K. (2018). Inference in linear regression models with many covariates and heteroscedasticity. *Journal of the American Statistical Association*, 113(523):1350–1361.

Davidson, R. and MacKinnon, J. G. (1999). The size distortion of bootstrap tests. *Econometric Theory*, 15(3):361–376.

Engle, R. F. and Manganelli, S. (2004). CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics*, 22(4):367–381.

Fan, R. and Lee, J. H. (2019). Predictive quantile regressions under persistence and conditional heteroskedasticity. *Journal of Econometrics*, 213(1):261–280.

Feng, X., He, X., and Hu, J. (2011). Wild bootstrap for quantile regression. *Biometrika*, 98(4):995–999.

Fitzenberger, B. (1998). The moving blocks bootstrap and robust inference for linear least squares and quantile regressions. *Journal of Econometrics*, 82(2):235–287.

Galvao, A. F. and Yoon, J. (2024). Hac covariance matrix estimation in quantile regression. *Journal of the American Statistical Association*, 119(547):2305–2316.

Glaeser, E. and Gyourko, J. (2018). The economic implications of housing supply. *Journal of Economic Perspectives*, 32(1):3–30.

Gregory, K. B., Lahiri, S. N., and Nordman, D. J. (2015). A smooth block bootstrap for statistical functionals and time series. *Journal of Time Series Analysis*, 36(3):442–461.

Gregory, K. B., Lahiri, S. N., and Nordman, D. J. (2018). A smooth block bootstrap for quantile regression with time series. *Annals of Statistics*, 46(3):1138–1166.

Hagemann, A. (2017). Cluster-robust bootstrap inference in quantile regression models. *Journal of the American Statistical Association*, 112(517):446–456.

Himmelberg, C., Mayer, C., and Sinai, T. (2005). Assessing high house prices: Bubbles, fundamentals and misperceptions. *Journal of Economic Perspectives*, 19(4):67–92.

Horowitz, J. L. (1998). Bootstrap methods for median regression models. *Econometrica*, pages 1327–1351.

Jochmans, K. (2022). Heteroscedasticity-robust inference in linear regression models with many covariates. *Journal of the American Statistical Association*, 117(538):887–896.

Kato, K. (2012). Asymptotic normality of powell's kernel estimator. *Annals of the Institute of Statistical Mathematics*, 64(2):255–273.

Knoll, K., Schularick, M., and Steger, T. (2017). No price like home: Global house prices, 1870–2012. *American Economic Review*, 107(2):331–353.

Koenker, R. (2005). *Quantile Regression*. Econometric Society Monographs. Cambridge University Press, Cambridge, UK.

Koenker, R. and Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1):33–50.

Koenker, R. and Xiao, Z. (2004). Unit root quantile autoregression inference. *Journal of the American Statistical Association*, 99(467):775–787.

Koenker, R. and Xiao, Z. (2006). Quantile autoregression. *Journal of the American Statistical Association*, 101(475):980–990.

Künsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics*, 17(3):1217–1241.

Lahiri, S., Furukawa, K., and Lee, Y.-D. (2007). A nonparametric plug-in rule for selecting optimal block lengths for block bootstrap methods. *Statistical Methodology*, 4(3):292–321.

Lee, J. H. (2016). Predictive quantile regression with persistent covariates: Ivx-qr approach. *Journal of Econometrics*, 192(1):105–118.

Linton, O. and Whang, Y.-J. (2007). The quantilogram: With an application to evaluating directional predictability. *Journal of Econometrics*, 141(1):250–282.

Mammen, E. (1993). Bootstrap and wild bootstrap for high dimensional linear models. *Annals of Statistics*, 21(1):255–285.

Newey, W. and West, K. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.

Oka, T. and Qu, Z. (2011). Estimating structural changes in regression quantiles. *Journal of Econometrics*, 162(2):248–267.

Paparoditis, E. and Politis, D. (2002). The tapered block bootstrap for general statistics from stationary sequences. *Econometrics Journal*, 5(1):131–148.

Paparoditis, E. and Politis, D. N. (2001). Tapered block bootstrap. *Biometrika*, 88(4):1105–1119.

Powell, J. (1991). Estimation of monotonic regression models under quantile restrictions. In Barnett, W., Powell, J., and Tauchen, G., editors, *Nonparametric and Semiparametric Methods in Econometrics and Statistics*. Cambridge University Press, Cambridge.

Shao, X. (2010). The dependent wild bootstrap. *Journal of the American Statistical Association*, 105(489):218–235.

White, H., Kim, T.-H., and Manganelli, S. (2015). Var for var: Measuring tail dependence using multivariate regression quantiles. *Journal of Econometrics*, 187(1):169–188.

Xiao, Z. (2009). Quantile cointegrating regression. *Journal of Econometrics*, 150(2):248–260.

Zhao, Q. (2000). Restricted regression quantiles. *Journal of Multivariate Analysis*, 72(1):78–99.

Figure 1: Empirical size comparison: Models 1 & 2

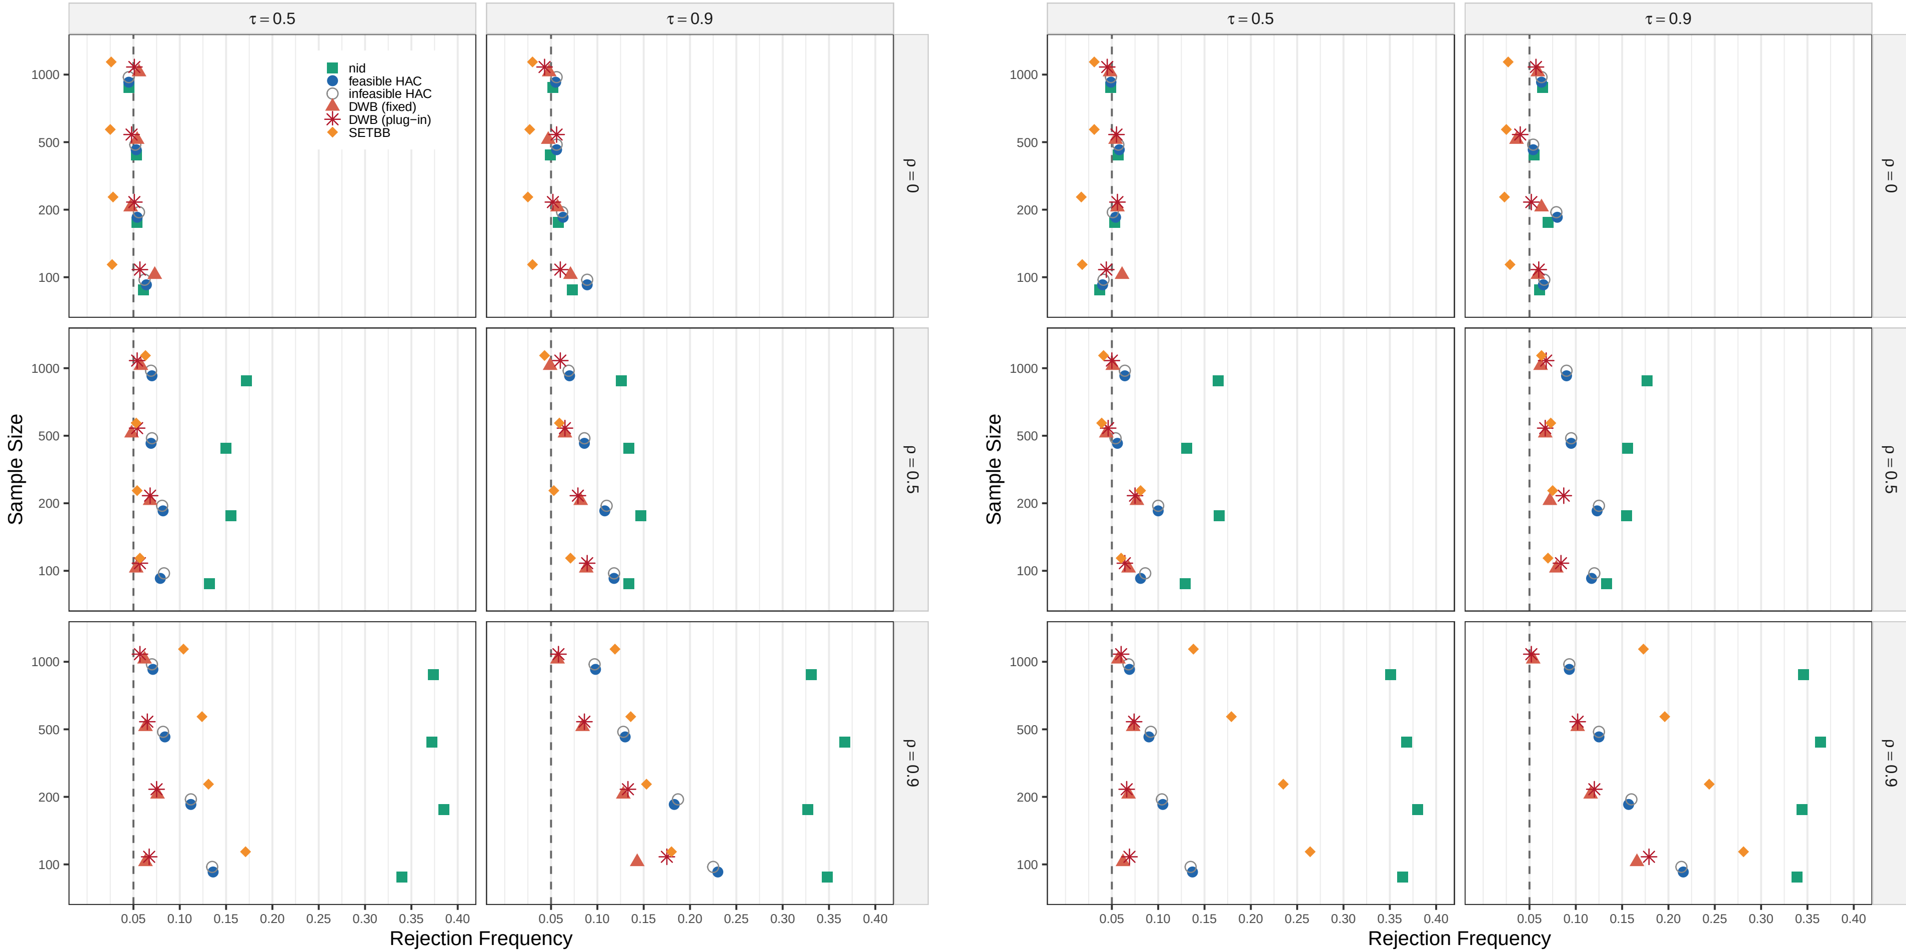

(a) Model 1 **(b)** Model 2

Notes. This figure reports empirical rejection rates of testing the null hypothesis $\mathbb{H}_0 : \beta_1(\tau) = 1$ by the conventional QR test based on the `nid` covariance estimator from the R package `quantreg`, the HAC estimator using the feasible bandwidth selector of Galvao and Yoon (2024) (feasible-HAC), the corresponding HAC estimator using the true autocorrelation coefficient in bandwidth selection (infeasible-HAC), the smooth extended tapered block bootstrap (SETBB) method of Gregory et al. (2018), and the proposed dependent wild bootstrap procedure with fixed scale constant $c = 2$ (fixed DWB) and plug-in selector in (3) with $c \in [0.5, 5]$ (plug-in DWB). The data generating processes Model 1 and Model 2 are defined in Section 5. The degree of autocorrelation is determined by $\rho \in \{0, 0.5, 0.9\}$, and the slope coefficient is tested at the quantiles $\tau \in \{0.5, 0.9\}$. The sample size $T = 1000$. The empirical rejection rate is computed at the 5% nominal level based on 1000 replications and 499 bootstrap draws.

Figure 2: Power comparison: Models 1 & 2

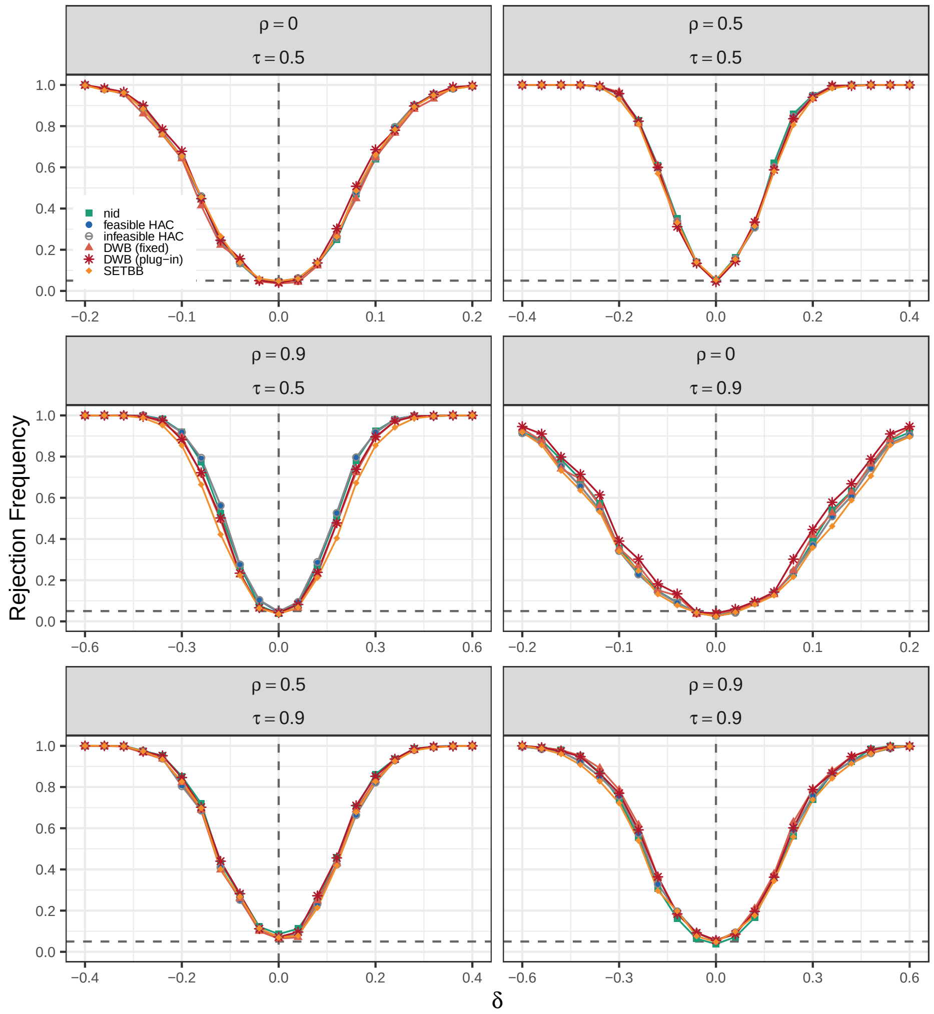

(a) Model 1

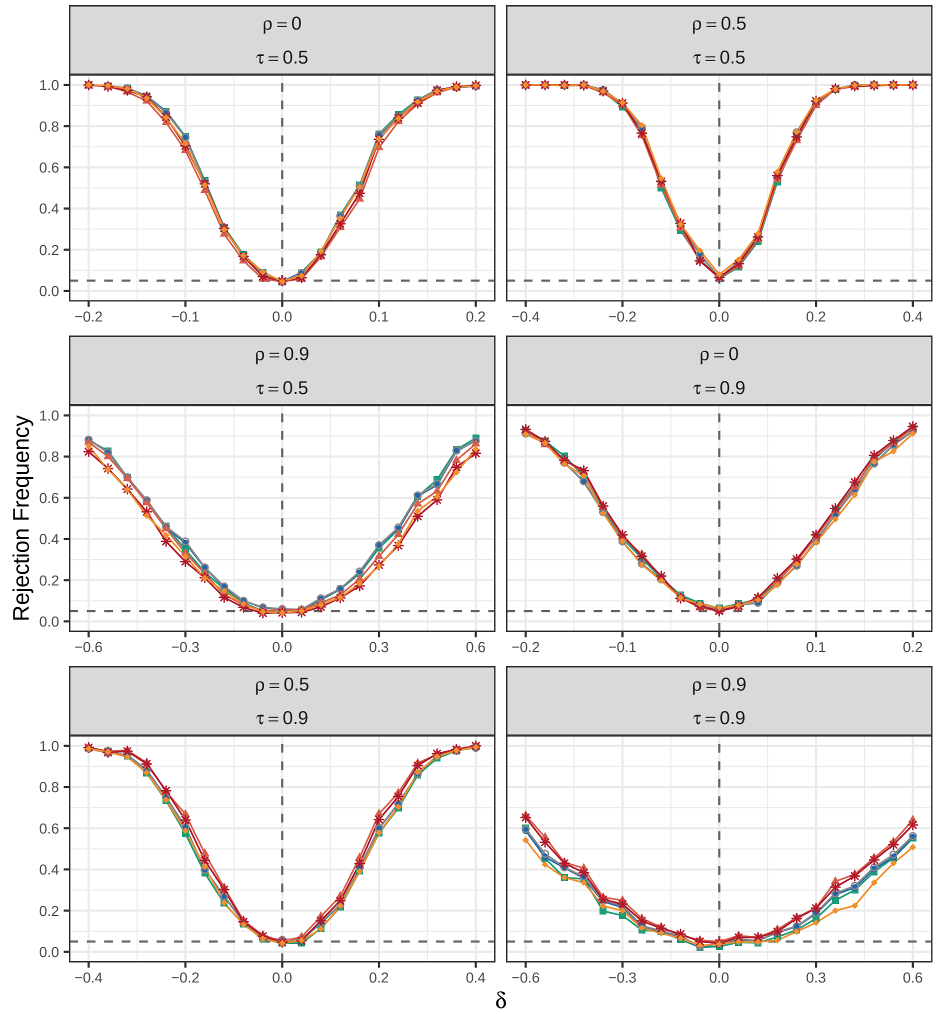

(b) Model 2

Notes. This figure reports empirical rejection rates of testing the null hypothesis $\mathbb{H}_0 : \beta_1(\tau) = 1+\delta$ by the conventional QR test based on the `nid` covariance estimator from the R package `quantreg`, the HAC estimator using the feasible plug-in bandwidth selector of Galvao and Yoon (2024) (feasible-HAC), the corresponding HAC estimator using the true autocorrelation coefficient in bandwidth selection (infeasible-HAC), the smooth extended tapered block bootstrap (SETBB) method of Gregory et al. (2018), and the proposed dependent wild bootstrap procedure with fixed scale constant $c = 2$ (fixed DWB) and plug-in selector in (3) with $c \in [0.5, 5]$ (plug-in DWB). The data generating processes Model 1 and Model 2 are defined in Section 5. δ determines the deviation of the slope coefficient from 1. The degree of autocorrelation is determined by $\rho \in \{0, 0.5, 0.9\}$, and the slope coefficient is tested at the quantiles $\tau \in \{0.5, 0.9\}$. The sample size $T = 1000$. The empirical rejection rate is computed at the 5% nominal level based on 1000 replications and 499 bootstrap draws.

Figure 3: Empirical size comparison: Multivariate quantile regression

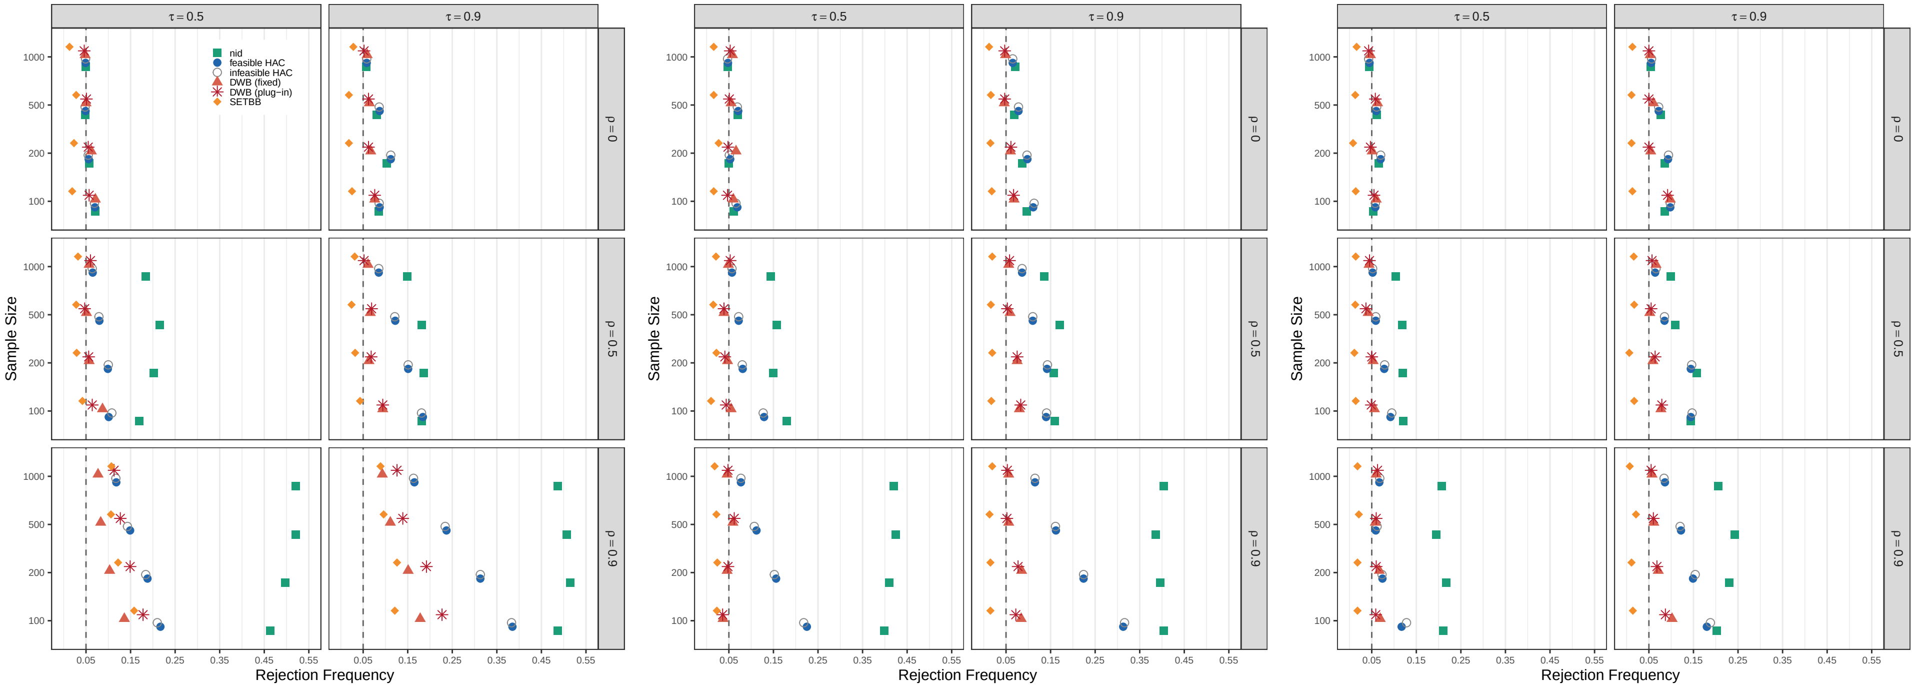

(a) $\mathbb{H}_0 : \beta_1(\tau) = 1$ **(b)** $\mathbb{H}_0 : \beta_2(\tau) = 1$ **(c)** $\mathbb{H}_0 : \beta_3(\tau) = 1$

Notes. This figure reports empirical rejection rates for testing the null hypothesis $\mathbb{H}_0 : \beta_j(\tau) = 1$, where $j \in \{1, 2, 3\}$ indexes the slope coefficient under consideration in the multivariate Model 1 design. The four procedures are the conventional QR test based on the `nid` covariance estimator from the R package `quantreg`, the HAC test using the feasible plug-in bandwidth selector of Galvao and Yoon (2024) (feasible-HAC), the corresponding HAC test using the true autocorrelation coefficients in bandwidth selection (infeasible-HAC), the smooth extended tapered block bootstrap (SETBB) method of Gregory et al. (2018), and the proposed dependent wild bootstrap procedure with fixed scale constant $c = 2$ (fixed DWB) and plug-in selector in (3) with $c \in [0.5, 5]$ (plug-in DWB). The data generating process extends Model 1 in Section 5 to three regressors, with $y_t = \beta_0 + \beta_1 x_{1t} + \beta_2 x_{2t} + \beta_3 x_{3t} + u_t$, where the regressors follow independent AR(1) processes with autoregressive coefficients 0.95, 0.85, and 0.53, respectively, and the disturbance follows the same AR(1) process as in Section 5. The disturbance autocorrelation is determined by $\rho \in \{0, 0.5, 0.9\}$, and inference is conducted at the quantiles $\tau \in \{0.5, 0.9\}$. The empirical rejection rates are computed at the 5% nominal level using 1000 Monte Carlo replications and 499 bootstrap draws.

Figure 4: Empirical size comparison: Models 1 & 2 with p covariates

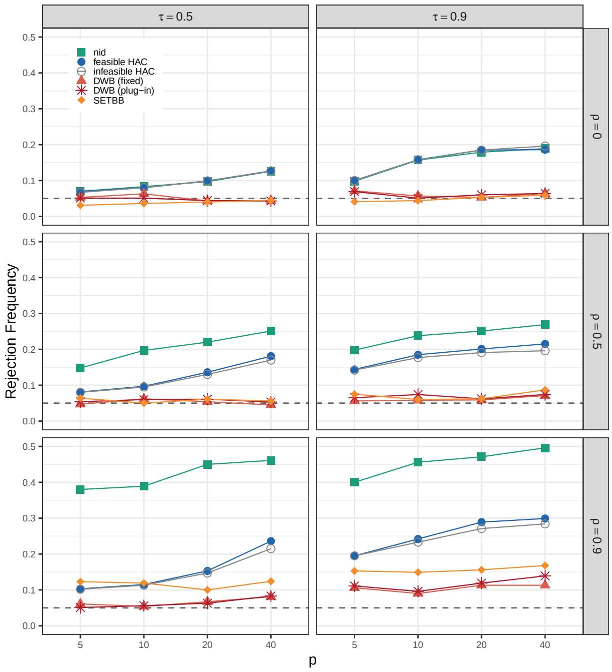

(a) Model 1: $T = 500$

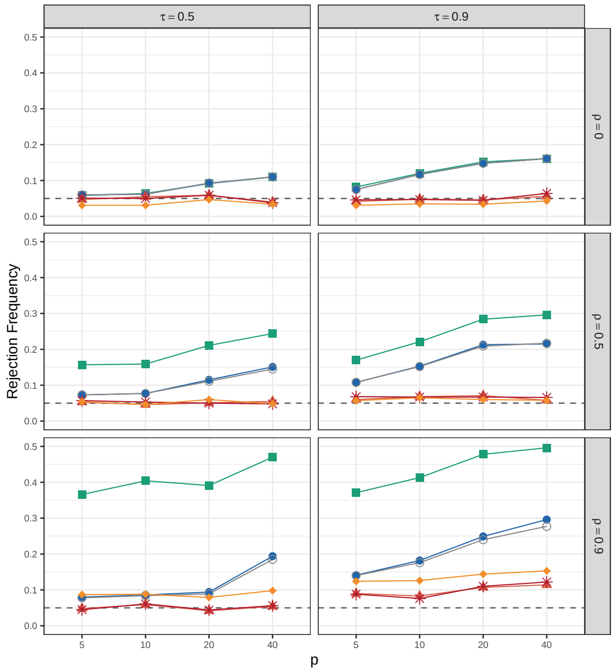

(b) Model 1: $T = 1000$

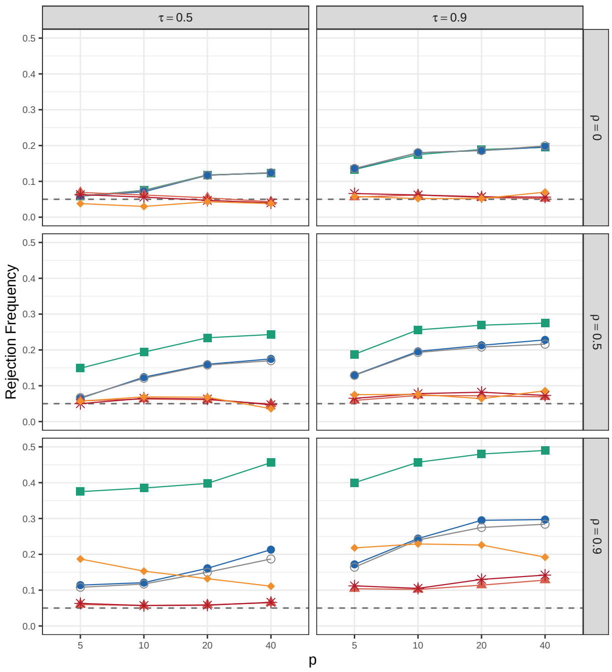

(c) Model 2: $T = 500$

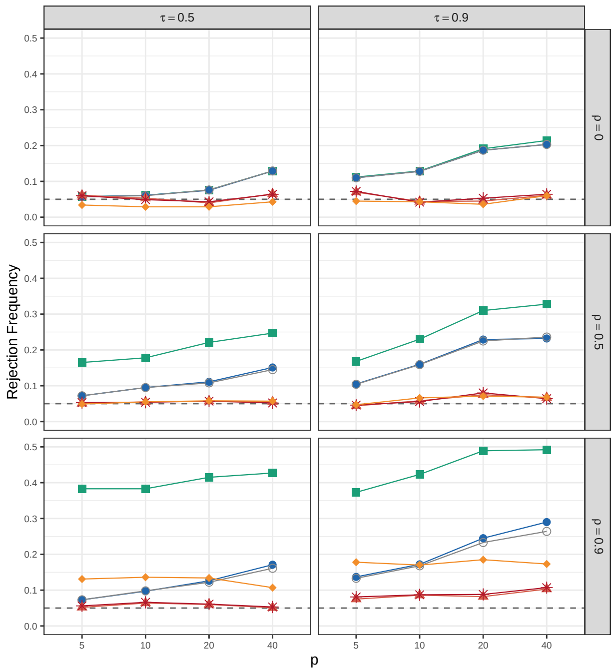

(d) Model 2: $T = 1000$

Notes. This figure reports empirical rejection rates of testing the null hypothesis $\mathbb{H}_0 : \beta_1(\tau) = 1$ by the conventional QR test based on the `nid` covariance estimator from the R package `quantreg`, the HAC estimator using the feasible plug-in bandwidth selector of Galvao and Yoon (2024) (feasible-HAC), the corresponding HAC estimator using the true autocorrelation coefficient in bandwidth selection (infeasible-HAC), the smooth extended tapered block bootstrap (SETBB) method of Gregory et al. (2018), and the proposed dependent wild bootstrap procedure with fixed scale constant $c = 2$ (fixed DWB) and plug-in selector in (3) with $c \in [0.5, 5]$ (plug-in DWB). The data generating processes Model 1 and Model 2 are defined in Section 5 with additional $p \in \{5, 10, 20, 40\}$ covariates. The degree of autocorrelation is determined by $\rho \in \{0, 0.5, 0.9\}$, and the slope coefficient is tested at the quantiles $\tau \in \{0.5, 0.9\}$. The sample size $T \in \{500, 1000\}$. The empirical rejection rate is computed at the 5% nominal level based on 1000 replications and 499 bootstrap draws.

Figure 5: Quantile regression coefficients: U.S. housing prices

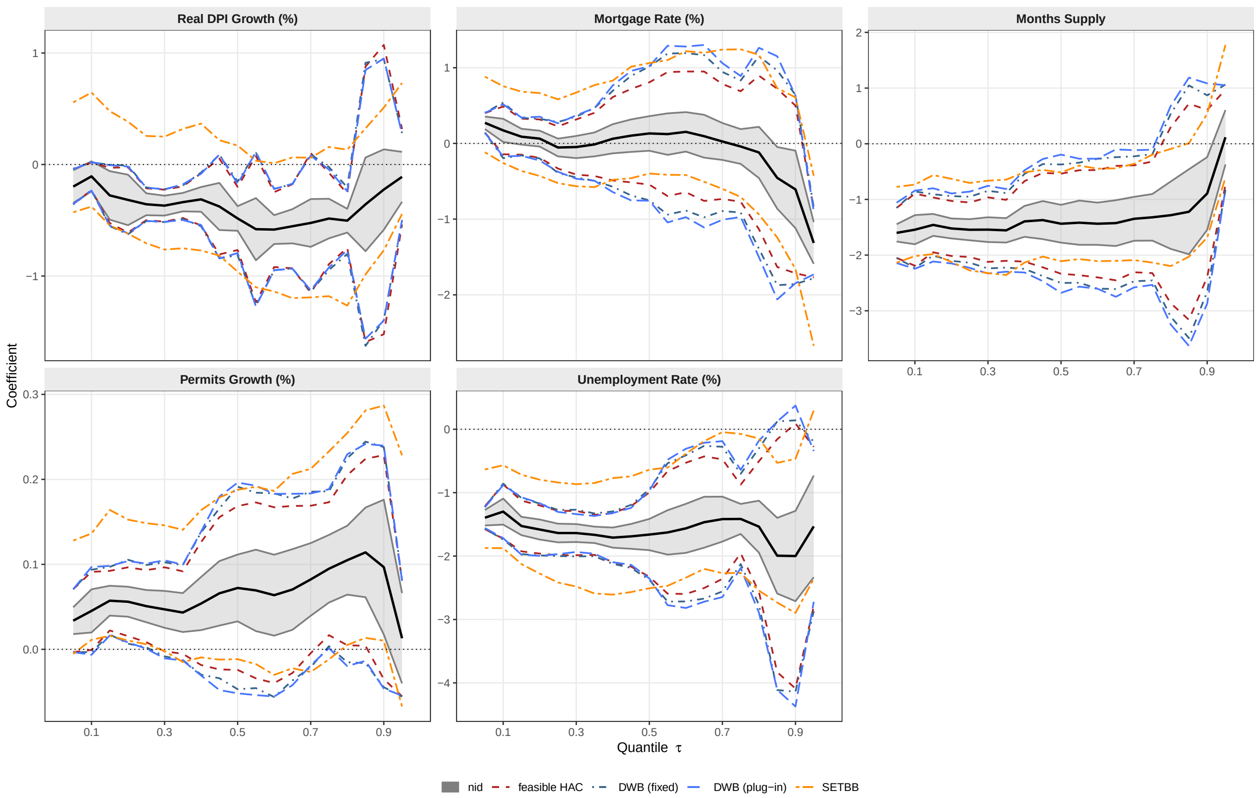

Notes. This figure presents the estimates of quantile regression $Q_{y_t}(\tau \mid x_t) = \beta_0(\tau) + \beta_1(\tau)\,\Delta^{12}\log(\mathrm{DPI}_t) + \beta_2(\tau)\,\mathrm{MORTGAGE}_t + \beta_3(\tau)\,\mathrm{SUPPLY}_t + \beta_4(\tau)\,\Delta^{12}\log(\mathrm{PERMIT}_t) + \beta_5(\tau)\,\mathrm{UNRATE}_t$ at $\tau \in \{0.05, 0.10, \ldots, 0.95\}$. y_t denotes the year-on-year percentage change in the Case-Shiller National Home Price Index, DPI denotes tje real disposable personal income, MORTGAGE denotes the 30-year fixed mortgage rate, SUPPLY denotes the months' supply of houses, PERMIT denotes the number of building permits for housing units, and UNRATE denotes the unemployment rate. The sample period is January 1988–December 2025. Each panel corresponds to one regressor. The solid black curve is the point estimate $\hat{\beta}_j(\tau)$ for regressor j at quantile τ. The grey shaded region is the conventional standard errors-based band, the dashed red lines bound the feasible HAC-based band by Galvao and Yoon (2024), the dashed orange lines bound the SETBB-based band of Gregory et al. (2018) using 999 replications, and the two dashed blue lines respectively bound the fixed and plug-in DWB-based band using 999 multiplier draws. The dotted horizontal line denotes zero.

Online Appendix

A Technical Proofs

A.1 Proof of Proposition 1

Proof. Write the un-clipped pilot constant as

$$\hat{c}_T^0(\tau) := \left\{ \frac{2\hat{B}_T(\tau)^2}{\hat{V}_T(\tau)} \right\}^{1/3},$$

and define its population counterpart by

$$c_g(\tau) := \left\{ \frac{2B_g(\tau)^2}{V_g(\tau)} \right\}^{1/3}.$$

By the assumed consistency of $\hat{B}_T(\tau)$ and $\hat{V}_T(\tau)$, together with $0 < B_g(\tau) < \infty$ and $0 < V_g(\tau) < \infty$, the continuous mapping theorem gives

$$\hat{c}_T^0(\tau) \xrightarrow{p} c_g(\tau) \in (0, \infty).$$

Let $\Pi_{[\underline{c},\bar{c}]}(x) = \min\{\bar{c}, \max\{\underline{c}, x\}\}$. Since this clipping map is continuous,

$$\hat{c}_T(\tau) = \Pi_{[\underline{c},\bar{c}]}\left(\hat{c}_T^0(\tau)\right) \xrightarrow{p} \Pi_{[\underline{c},\bar{c}]}\left(c_g(\tau)\right).$$

Moreover, by construction,

$$0 < \underline{c} \leq \hat{c}_T(\tau) \leq \bar{c} < \infty$$

for every T. Therefore

$$\hat{\ell}_T(\tau) = \left\lceil \hat{c}_T(\tau) T^{1/3} \right\rceil = \left\lceil \hat{c}_T(\tau) T^{\kappa} \right\rceil \qquad \text{with } \kappa = 1/3 \in (0, 1/2).$$

Since $\hat{c}_T(\tau) \geq \underline{c} > 0$, we have

$$\hat{\ell}_T(\tau) \geq \underline{c}\, T^{1/3} \to \infty.$$

Since $\hat{c}_T(\tau) \le \bar{c} < \infty$,

$$\frac{\hat{\ell}_T(\tau)}{T} \le \bar{c}\, T^{-2/3} + T^{-1} \to 0.$$

Finally,

$$\frac{\hat{\ell}_T(\tau)}{T^{1/3}} \le \hat{c}_T(\tau) + T^{-1/3} \le \bar{c} + 1$$

for all sufficiently large T, so $\hat{\ell}_T(\tau)/T^{1/3} = O_p(1)$. These are exactly the selector conditions required in Assumption 5 for the special case $\kappa = 1/3$, which proves the proposition. ∎

A.2 Proof of Proposition 2

Proof. The unrestricted Bahadur expansion is exactly the dependent data quantile regression expansion proved in Galvao and Yoon (2024). We therefore only need to derive the restricted formulas.

Under $\mathbb{H}_0$, let

$$Q_T(b) = \sum_{t=1}^{T} \rho_\tau(y_t - x_t^{*\prime} b)\,, \qquad \Psi_T(b) = \frac{1}{\sqrt{T}} \sum_{t=1}^{T} x_t^* \psi_\tau(y_t - x_t^{*\prime} b)\,,$$

and write $\Delta_T^r = \sqrt{T}\big(\hat{\beta}_0^r(\tau) - \beta_0(\tau)\big)$. Because the unrestricted estimator is $T^{1/2}$-consistent and the restriction set $\{b : Rb = r_0\}$ is a closed affine subspace, the standard convexity argument for restricted quantile regression implies $\Delta_T^r = O_p(1)$. A subgradient form of the Karush-Kuhn-Tucker (KKT) conditions for the equality-constrained problem then yields a multiplier λ_T such that

$$0 = \Psi_T\Big(\hat{\beta}_0^r(\tau)\Big) - R'\lambda_T, \qquad R\Delta_T^r = 0,$$

where the second relation is exact because $R\hat{\beta}_0^r(\tau) = r_0 = R\beta_0(\tau)$ under $\mathbb{H}_0$.

Using the standard local quantile regression expansion underlying the unrestricted Bahadur representation, uniformly for h in bounded sets,

$$\Psi_T\left(\beta_0(\tau) + \frac{h}{\sqrt{T}}\right) = \mathbb{G}_T(\tau) - D(\tau)h + o_p(1).$$

Evaluating this display at $h = \Delta_T^r$ gives the linearized constrained KKT system

$$0 = \mathbb{G}_T(\tau) - D(\tau)\Delta_T^r - R'\lambda_T + o_p(1), \qquad R\Delta_T^r = 0.$$

Premultiplying the first equation by $RD(\tau)^{-1}$ and using $R\Delta_T^r = 0$ yields

$$\lambda_T = \big(RD(\tau)^{-1}R'\big)^{-1} RD(\tau)^{-1}\mathbb{G}_T(\tau) + o_p(1).$$

Substituting back gives

$$\Delta_T^r = \Big[D(\tau)^{-1} - D(\tau)^{-1}R'\big(RD(\tau)^{-1}R'\big)^{-1}RD(\tau)^{-1}\Big]\,\mathbb{G}_T(\tau) + o_p(1),$$

which is the stated formula with $A_R(\tau)$.

Under the local alternative sequence $\mathbb{H}_{1,T}$, define

$$\Delta_{T,1}^r = \sqrt{T}\big(\hat{\beta}_0^r(\tau) - \beta_{0,T}(\tau)\big).$$

Exact feasibility of the restricted estimator now gives

$$R\Delta_{T,1}^r = \sqrt{T}\big(r_0 - R\beta_{0,T}(\tau)\big) = -\delta.$$

Because the local alternative is contiguous, $\|\beta_{0,T}(\tau) - \beta_0(\tau)\| = O(T^{-1/2})$. Assumption 2 therefore implies that the Jacobian evaluated along this local sequence is still $D(\tau) + o(1)$, so the same local expansion around $\beta_{0,T}(\tau)$ yields

$$\sqrt{T}\big(\hat{\beta}(\tau) - \beta_{0,T}(\tau)\big) = D(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + o_p(1).$$

Premultiplying by R and using $R\beta_{0,T}(\tau) = r_0 + \delta/\sqrt{T}$ gives

$$\sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big) = \delta + RD(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + o_p(1).$$

For the restricted estimator, the same local expansion around $\beta_{0,T}(\tau)$ yields

$$0 = \mathbb{G}_{T,1}(\tau) - D(\tau)\Delta^r_{T,1} - R'\lambda_{T,1} + o_p(1), \qquad R\Delta^r_{T,1} = -\delta.$$

Premultiplying by $RD(\tau)^{-1}$ gives

$$\lambda_{T,1} = \big(RD(\tau)^{-1}R'\big)^{-1}\big\{RD(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + \delta\big\} + o_p(1),$$

and substituting back yields

$$\Delta^r_{T,1} = A_R(\tau)\mathbb{G}_{T,1}(\tau) - D(\tau)^{-1}R'\big(RD(\tau)^{-1}R'\big)^{-1}\delta + o_p(1).$$

This proves the proposition. ■

A.3 Proof of Theorem 1

Proof. This result is a scalar-restriction corollary of the pointwise HAC theory in Galvao and Yoon (2024). The only additional step is to pass from their vector-valued quantile regression statistic to the scalar restriction $R\beta_0(\tau) = r_0$. Under $\mathbb{H}_0$, we have $r_0 = R\beta_0(\tau)$. By the unrestricted part of Proposition 2, which is the notation used here for the dependent-data Bahadur representation in equation (3.1) of Galvao and Yoon (2024),

$$\sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big) = RD(\tau)^{-1}\mathbb{G}_T(\tau) + o_p(1), \qquad \mathbb{G}_T(\tau) = \frac{1}{\sqrt{T}}\sum_{t=1}^{T} z_t(\tau),$$

where $z_t(\tau) = x_t^*\psi_\tau(u_t(\tau))$.

We first establish the asymptotic law of the numerator. For any fixed $c \in \mathbb{R}^{p+1}$, the scalar process $\{c'z_t(\tau)\}_{t\in\mathbb{Z}}$ is strictly stationary and strong-mixing with the same mixing coefficients as the original data process. Because $|\psi_\tau(u)| \leq 1$ and $E\|x_t^*\|^{4+\eta} < \infty$ by Assumption 1, we have $E|c'z_t(\tau)|^{4+\eta} < \infty$. Therefore the strong-mixing central limit theorem used in Galvao and Yoon (2024) applies to every such linear combination. Hence, by the Cramér-Wold device and Assumption 1,

$$\mathbb{G}_T(\tau) \xrightarrow{d} N\big(0, J(\tau)\big).$$

By Assumption 3, the projected long-run variance $\sigma_R^2(\tau) = RD(\tau)^{-1}J(\tau)D(\tau)^{-1}R'$ is strictly positive. Hence, by the continuous mapping theorem and Slutsky's lemma,

$$\sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big) \xrightarrow{d} N\big(0, RD(\tau)^{-1}J(\tau)D(\tau)^{-1}R'\big) = N\big(0, \sigma_R^2(\tau)\big)\,.$$

For the feasible denominator, the observed-studentizer part of Lemma 1 gives

$$\hat{s}_T^r(\tau) \xrightarrow{p} \sigma_R(\tau)$$

under the same assumption set. Combining the projected numerator expansion with this denominator consistency, Slutsky's lemma gives

$$T_T^r(\tau) = \frac{\sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big)}{\hat{s}_T^r(\tau)} \xrightarrow{d} N(0,1),$$

which proves the theorem. ■

A.4 Proof of Lemma 1

Proof. We begin with the observed projected-score studentizer. Let

$$\hat{z}_t^{\circ}(\tau) = \hat{z}_t(\tau) - \bar{z}_u(\tau), \qquad \bar{z}_u(\tau) = \frac{1}{T}\sum_{t=1}^{T}\hat{z}_t(\tau),$$

and define the centered lag-window matrix

$$\hat{J}_T^{\circ}(\tau) := \frac{1}{T}\sum_{t=1}^{T}\hat{z}_t^{\circ}(\tau)\hat{z}_t^{\circ}(\tau)' + 2\sum_{j=1}^{T-1}k\left(\frac{j}{S_T}\right)\frac{1}{T}\sum_{t=j+1}^{T}\hat{z}_t^{\circ}(\tau)\hat{z}_{t-j}^{\circ}(\tau)'.$$

By construction of the scalar functional,

$$\tilde{s}_T^r(\tau)^2 = R\hat{D}(\tau)^{-1}\hat{J}_T^{\circ}(\tau)\hat{D}(\tau)^{-1}R'.$$

Under Assumptions 1, 2, 3, 4, 6, and 7, the feasible HAC matrix estimator satisfies

$$\hat{\Sigma}_r(\tau) = \hat{D}(\tau)^{-1}\hat{J}_T(\tau)\hat{D}(\tau)^{-1} \xrightarrow{p} \Sigma(\tau)$$

by the fixed-τ pointwise HAC consistency result of Galvao and Yoon (2024). Here Assumptions 1, 2, 3, 4, and 6 match their unrestricted score conditions, while Assumption 7 is used below for the projected-score comparison and the bootstrap studentizer. It therefore suffices to compare $\hat{J}_T^\circ(\tau)$ and $\hat{J}_T(\tau)$. Since

$$\sqrt{T}\,\bar{z}_u(\tau) = \frac{1}{\sqrt{T}} \sum_{t=1}^{T} \hat{z}_t(\tau) = \mathbb{G}_T(\tau) + o_p(1) = O_p(1),$$

we have $\bar{z}_u(\tau) = O_p(T^{-1/2})$. Standard HAC demeaning algebra then implies

$$\hat{J}_T^\circ(\tau) - \hat{J}_T(\tau) = o_p(1),$$

because the effect of replacing $\hat{z}_t(\tau)$ by $\hat{z}_t^\circ(\tau)$ is of order $O_p(S_T/T)$ and $S_T = o(T^{1/2})$ by Assumption 4. Consequently,

$$\tilde{s}_T^r(\tau)^2 - R\hat{\Sigma}_r(\tau)R' = o_p(1).$$

Since $R\hat{\Sigma}_r(\tau)R' \xrightarrow{p} \sigma_R^2(\tau)$ and Assumption 3 gives $\sigma_R^2(\tau) > 0$, the continuous mapping theorem yields

$$\tilde{s}_T^r(\tau) \xrightarrow{p} \sigma_R(\tau) \qquad \text{and} \qquad |\hat{s}_T^r(\tau) - \tilde{s}_T^r(\tau)| \xrightarrow{p} 0.$$

We next turn to the bootstrap studentizer. Write

$$\tilde{g}_{t,0}(\tau) = R\hat{D}_0(\tau)^{-1}\tilde{z}_{t,0}(\tau), \qquad g_t^*(\tau) = \xi_t^* \tilde{g}_{t,0}(\tau).$$

Let

$$\bar{g}_T^*(\tau) = \frac{1}{T} \sum_{t=1}^{T} g_t^*(\tau)$$

and define the uncentered bootstrap Bartlett functional

$$\bar{\mathcal{L}}_T^*(\tau) := \frac{1}{T} \sum_{t=1}^{T} g_t^*(\tau)^2 + 2 \sum_{j=1}^{\ell_T} \left(1 - \frac{j}{\ell_T}\right) \frac{1}{T} \sum_{t=j+1}^{T} g_t^*(\tau) g_{t-j}^*(\tau).$$

Because $\sum_{t=1}^{T} \tilde{g}_{t,0}(\tau) = 0$ by construction, Assumption 5 implies

$$E^*[\bar{g}_T^*(\tau)] = 0$$

and

$$E^*\left[\bar{g}_T^*(\tau)^2\right] = \frac{1}{T^2}\sum_{t=1}^{T}\sum_{s=1}^{T} \tilde{g}_{t,0}(\tau)\tilde{g}_{s,0}(\tau)E^*(\xi_t^*\xi_s^*) = O_p\left(\frac{\ell_T}{T}\right) = o_p(1).$$

Hence the difference between the centered and uncentered Bartlett functionals is $o_{P^*}(1)$ in probability, so it is enough to study $\bar{\mathcal{L}}_T^*(\tau)$.

Conditional on the sample,

$$E^*\left[\bar{\mathcal{L}}_T^*(\tau)\right] = \frac{1}{T}\sum_{t=1}^{T} \tilde{g}_{t,0}(\tau)^2 + 2\sum_{j=1}^{\ell_T}\left(1 - \frac{j}{\ell_T}\right)\frac{1}{T}\sum_{t=j+1}^{T} \tilde{g}_{t,0}(\tau)\tilde{g}_{t-j,0}(\tau)E^*(\xi_t^*\xi_{t-j}^*).$$

For the Bartlett moving-average multipliers,

$$E^*(\xi_t^*\xi_{t-j}^*) = 1 - \frac{j}{\ell_T} + o_p(1)$$

uniformly for $1 \le j \le \ell_T$, so

$$E^*\left[\bar{\mathcal{L}}_T^*(\tau)\right] = \mathcal{L}_T^{\mathrm{B}}\big(\{\tilde{g}_{t,0}(\tau)\}_{t=1}^{T}; \ell_T\big) + o_p(1).$$

Under $\mathbb{H}_0$, the restricted Jacobian estimator and the unrestricted one are asymptotically equivalent, and the restricted projected score obeys the same lag-window covariance approximation as the unrestricted projected score. Therefore

$$\mathcal{L}_T^{\mathrm{B}}\big(\{\tilde{g}_{t,0}(\tau)\}_{t=1}^{T}; \ell_T\big) \xrightarrow{p} \sigma_R^2(\tau).$$

More explicitly, Proposition 2 implies that under $\mathbb{H}_0$ the restricted and unrestricted estimators differ by $O_p(T^{-1/2})$, while Assumption 6 gives $\|\hat{D}_0(\tau) - \hat{D}(\tau)\| = o_p(1)$. Therefore the restricted and unrestricted projected scores differ by an amount whose contribution to each sample projected autocovariance is $o_p(1)$ uniformly over the lags entering the Bartlett sum. Since the Bartlett weights are uniformly bounded and $\ell_T = o(T^{1/2})$, summing these lagwise differences yields

$$\mathcal{L}_T^{\mathrm{B}}\big(\{\tilde{g}_{t,0}(\tau)\}_{t=1}^T;\ell_T\big)-\mathcal{L}_T^{\mathrm{B}}\big(\{\hat{g}_{t,u}(\tau)\}_{t=1}^T;\ell_T\big)=o_p(1),$$

where $\hat{g}_{t,u}(\tau)=R\hat{D}(\tau)^{-1}\hat{z}_t(\tau)$. The unrestricted Bartlett functional is the Bartlett-kernel special case of the observed projected-score HAC functional, so the observed-studentizer argument above implies its convergence to $\sigma_R^2(\tau)$. The displayed convergence follows.

It remains to control the conditional fluctuation around the conditional mean. Using the uniformly bounded $(2+\eta)$ moments of the multipliers in Assumption 5 together with the absolute summability of the projected-score autocovariances and fourth-order cumulants in Assumption 7, the same cumulant calculations used in standard HAC variance-consistency arguments imply

$$\mathrm{Var}^*\big(\bar{\mathcal{L}}_T^*(\tau)\big)=o_p(1).$$

To see the structure of this bound, write

$$\bar{\mathcal{L}}_T^*(\tau)=\sum_{j=0}^{\ell_T}\omega_{j,T}\hat{\gamma}_{j,T}^*(\tau),$$

where $\omega_{0,T}=1$, $\omega_{j,T}=2(1-j/\ell_T)$ for $j\geq 1$, and

$$\hat{\gamma}_{j,T}^*(\tau)=\frac{1}{T}\sum_{t=j+1}^{T}g_t^*(\tau)g_{t-j}^*(\tau)=\frac{1}{T}\sum_{t=j+1}^{T}\xi_t^*\xi_{t-j}^*\tilde{g}_{t,0}(\tau)\tilde{g}_{t-j,0}(\tau).$$

Then

$$\mathrm{Var}^*\big(\bar{\mathcal{L}}_T^*(\tau)\big)\leq C\sum_{j=0}^{\ell_T}\sum_{k=0}^{\ell_T}\left|\mathrm{Cov}^*\big(\hat{\gamma}_{j,T}^*(\tau),\hat{\gamma}_{k,T}^*(\tau)\big)\right|.$$

For Gaussian Bartlett multipliers, Isserlis' formula expands each conditional covariance into a finite sum of products of multiplier covariances $r_h^*=E^*(\xi_t^*\xi_{t-h}^*)=a(h/\ell_T)+o_p(1)$. Because $a(\cdot)$ is uniformly bounded and effectively supported on lags of order ℓ_T, the same double-sum reduction used in standard HAC variance proofs shows that the full covariance sum is bounded by a constant multiple of ℓ_T/T times sample averages of fourth-order products of $\{\tilde{g}_{t,0}(\tau)\}$. Assumption 7 makes those sample averages $O_p(1)$ through absolute summability of projected-score autocovariances and fourth-order cumulants. Consequently,

$$\mathrm{Var}^*\big(\bar{\mathcal{L}}_T^*(\tau)\big) = O_p\bigg(\frac{\ell_T}{T}\bigg) = o_p(1),$$

which is the variance bound asserted above. Conditional Chebyshev's inequality therefore gives

$$\bar{\mathcal{L}}_T^*(\tau) - E^*\big[\bar{\mathcal{L}}_T^*(\tau)\big] = o_{P^*}(1)$$

in outer probability, and hence

$$P^*\Big(\big|\hat{s}_T^*(\tau)^2 - \sigma_R^2(\tau)\big| > \epsilon\Big) \xrightarrow{p} 0$$

for every $\epsilon > 0$. Because $\sigma_R(\tau) > 0$, another application of the continuous mapping theorem yields

$$P^*(|\hat{s}_T^*(\tau) - \sigma_R(\tau)| > \epsilon) \xrightarrow{p} 0.$$

The remaining conclusions follow from the triangle inequality:

$$|\hat{s}_T^*(\tau) - \tilde{s}_T^r(\tau)| \leq |\hat{s}_T^*(\tau) - \sigma_R(\tau)| + |\tilde{s}_T^r(\tau) - \sigma_R(\tau)|\,,$$

and, using also $|\hat{s}_T^r(\tau) - \tilde{s}_T^r(\tau)| \xrightarrow{p} 0$,

$$|\hat{s}_T^*(\tau) - \hat{s}_T^r(\tau)| \leq |\hat{s}_T^*(\tau) - \sigma_R(\tau)| + |\hat{s}_T^r(\tau) - \sigma_R(\tau)|\,.$$

This proves the lemma. ■

A.5 Proof of Theorem 2

Proof. The proof has three parts. First, under $\mathbb{H}_0$, the restricted projected score $\tilde{g}_{t,0}(\tau) = R\hat{D}_0(\tau)^{-1}\tilde{z}_{t,0}(\tau)$ is asymptotically equivalent to the centered population projected score $g_t^\circ(\tau) = g_t(\tau) - \bar{g}_T(\tau)$, where $g_t(\tau) = RD(\tau)^{-1}x_t^*\psi_\tau(u_t(\tau))$ and $\bar{g}_T(\tau) = T^{-1}\sum_{t=1}^T g_t(\tau)$. Define the difference sequence

$$\delta_{t,T}(\tau) = \tilde{g}_{t,0}(\tau) - g_t^\circ(\tau).$$

Proposition 2, Assumption 6, and the same null-imposition and recentering algebra used in Lemma 1 imply that the lag-window functional of this difference is negligible:

$$\mathcal{L}_T\big(\{\delta_{t,T}(\tau)\}_{t=1}^{T}; a, \ell_T\big) = o_p(1).$$

Hence, if

$$M_{T,\delta}^*(\tau) := \frac{1}{\sqrt{T}} \sum_{t=1}^{T} \xi_t^* \delta_{t,T}(\tau),$$

then the covariance approximation in Assumption 5 implies that $\mathrm{Var}^*(M_{T,\delta}^*(\tau))$ is the lag-window quadratic form generated by $a(\cdot)$ and evaluated on $\{\delta_{t,T}(\tau)\}_{t=1}^{T}$, up to an $o_p(1)$ remainder. The preceding display therefore yields

$$\mathrm{Var}^*\big(M_{T,\delta}^*(\tau)\big) = o_p(1).$$

Therefore

$$M_{T,\delta}^*(\tau) = o_{P^*}(1)$$

in probability, so replacing $\tilde{g}_{t,0}(\tau)$ by $g_t^\circ(\tau)$ inside the multiplier sum is asymptotically negligible.

Second, Assumptions 5 and 7 are the scalar conditions needed for the dependent multiplier central limit theorem. Therefore the standard argument for dependent wild or multiplier bootstraps, as in Shao (2010), yields

$$\sup_{z\in\mathbb{R}} \left| P^*\left(\frac{1}{\sqrt{T}} \sum_{t=1}^{T} \xi_t^* g_t^\circ(\tau) \le z \right) - \Phi(z/\sigma_R(\tau)) \right| \xrightarrow{p} 0.$$

By the preceding step, the same conclusion holds with $N_T^*(\tau)$ in place of the centered population-score multiplier sum:

$$\sup_{z\in\mathbb{R}} |P^*(N_T^*(\tau) \le z) - \Phi(z/\sigma_R(\tau))| \xrightarrow{p} 0.$$

Intuitively, the multiplier covariance $E^*(\xi_t^*\xi_s^*) = a((t-s)/\ell_T) + o_p(1)$ reproduces the dependence pattern needed for the projected-score sum, while the cumulant summability in Assumption 7 controls the higher-order terms.

Third, Lemma 1 gives studentization by bootstrap: for every $\epsilon > 0$,

$$P^*(|\hat{s}_T^*(\tau) - \sigma_R(\tau)| > \epsilon) \xrightarrow{p} 0.$$

Combining the conditional Gaussian approximation for $N_T^*(\tau)$ with the preceding display, the conditional Slutsky theorem yields

$$\sup_{z \in \mathbb{R}} |P^*(T_T^*(\tau) \leq z) - \Phi(z)| \xrightarrow{p} 0.$$

For the observed statistic, let

$$N_{T,0}^{(0)}(\tau) = \frac{1}{\sqrt{T}} \sum_{t=1}^{T} \hat{g}_{t,0}(\tau), \qquad s_{T,0}^{(0)}(\tau) = \sqrt{\mathcal{L}_T^{\mathrm{B}}(\{\hat{g}_{t,0}(\tau)\}_{t=1}^{T}; \ell_T)}.$$

Under $\mathbb{H}_0$, the same restricted-versus-unrestricted numerator comparison used in Proposition 2 gives

$$N_{T,0}^{(0)}(\tau) = \sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big) + o_p(1).$$

Moreover, the restricted-denominator comparison in the proof of Lemma 1 together with the observed-studentizer statement of that lemma gives

$$s_{T,0}^{(0)}(\tau) \xrightarrow{p} \sigma_R(\tau), \qquad \hat{s}_T^r(\tau) \xrightarrow{p} \sigma_R(\tau),$$

and therefore

$$T_{T,0}^r(\tau) - T_T^r(\tau) = o_p(1).$$

Turning to the observed statistic, Theorem 1 shows $T_T^r(\tau) \xrightarrow{d} N(0,1)$. Hence

$$T_{T,0}^r(\tau) \xrightarrow{d} N(0,1).$$

Since $\Phi(\cdot)$ is continuous, Pólya's theorem yields

$$\sup_{z \in \mathbb{R}} \big|P\big(T_{T,0}^r(\tau) \leq z\big) - \Phi(z)\big| \to 0.$$

Finally, by the triangle inequality,

$$\sup_{z\in\mathbb{R}} \left|P^*(T_T^*(\tau) \le z) - P\big(T_{T,0}^r(\tau) \le z\big)\right|$$

is bounded above by

$$\sup_{z\in\mathbb{R}} |P^*(T_T^*(\tau) \le z) - \Phi(z)| + \sup_{z\in\mathbb{R}} \left|P\big(T_{T,0}^r(\tau) \le z\big) - \Phi(z)\right|,$$

and both terms converge to zero, the first in probability and the second deterministically. This proves the theorem. ■

A.6 Proof of Corollary 1

Proof. Let

$$F_T^*(z) = P^*(T_T^*(\tau) \le z), \qquad F_{T,0}(z) = P\big(T_{T,0}^r(\tau) \le z\big).$$

By Theorem 2,

$$\sup_{z\in\mathbb{R}} |F_T^*(z) - F_{T,0}(z)| \xrightarrow{p} 0.$$

Under $\mathbb{H}_0$, Theorem 1 gives $T_T^r(\tau) \xrightarrow{d} N(0,1)$, and the null equivalence established in the proof of Theorem 2 gives $T_{T,0}^r(\tau) - T_T^r(\tau) = o_p(1)$. Hence $T_{T,0}^r(\tau) \xrightarrow{d} N(0,1)$ as well. With $\Phi(\cdot)$, Pólya's theorem yields

$$\sup_{z\in\mathbb{R}} |F_{T,0}(z) - \Phi(z)| \to 0.$$

Therefore, by the triangle inequality,

$$\sup_{z\in\mathbb{R}} |P^*(T_T^*(\tau) \le z) - \Phi(z)| \le \sup_{z\in\mathbb{R}} |F_T^*(z) - F_{T,0}(z)| + \sup_{z\in\mathbb{R}} |F_{T,0}(z) - \Phi(z)|,$$

and the right-hand side converges to zero in probability. ■

A.7 Proof of Corollary 2

Proof. Let $Z \sim N(0,1)$ and define the distribution function of $|Z|$ by

$$F_{|Z|}(x) = P(|Z| \le x) = \begin{cases} 0, & x < 0, \\ 2\Phi(x) - 1, & x \ge 0. \end{cases}$$

By Corollary 1,

$$\sup_{x \in \mathbb{R}} \left| P^*(|T_T^*(\tau)| \le x) - F_{|Z|}(x) \right| \xrightarrow{p} 0.$$

Because $F_{|Z|}$ is continuous and strictly increasing on $[0, \infty)$, its $(1-\alpha)$-quantile is

$$F_{|Z|}^{-1}(1-\alpha) = z_{1-\alpha/2}.$$

Hence the bootstrap quantile satisfies

$$c_{1-\alpha}^*(\tau) \xrightarrow{p} z_{1-\alpha/2}.$$

For the observed statistic, Theorem 1 and the null equivalence established in the proof of Theorem 2 imply

$$\left| T_{T,0}^r(\tau) \right| \xrightarrow{d} |Z|.$$

Therefore, by Slutsky's lemma,

$$\left| T_{T,0}^r(\tau) \right| - c_{1-\alpha}^*(\tau) \xrightarrow{d} |Z| - z_{1-\alpha/2}.$$

Since $P(|Z| = z_{1-\alpha/2}) = 0$, the portmanteau theorem gives

$$P\big(\left| T_{T,0}^r(\tau) \right| > c_{1-\alpha}^*(\tau)\big) \to P\big(|Z| > z_{1-\alpha/2}\big) = \alpha.$$

∎

A.8 Proof of Lemma 2

Proof. Under the contiguous Pitman sequence

$$\mathbb{H}_{1,T} : R\beta_{0,T}(\tau) = r_0 + \delta/\sqrt{T}, \qquad \delta \in \mathbb{R},$$

Proposition 2 gives

$$\Delta^r_{T,1}(\tau) := \sqrt{T}\big(\hat\beta^r_0(\tau) - \beta_{0,T}(\tau)\big) = O_p(1).$$

Hence $\hat\beta^r_0(\tau) - \beta_{0,T}(\tau) = O_p(T^{-1/2})$. The Jacobian estimator is local in the residuals, and Assumptions 2 and 6 are formulated on a shrinking neighborhood of zero. Therefore the same Powell-kernel expansion used under $\mathbb{H}_0$ applies with $\beta_0(\tau)$ replaced by $\beta_{0,T}(\tau)$, which yields

$$\|\hat D_0(\tau) - D(\tau)\| = o_p(T^{-1/4}).$$

By the same argument, the unrestricted Jacobian estimator and the unrestricted HAC studentizer remain consistent under $\mathbb{H}_{1,T}$:

$$\hat s^r_T(\tau) \xrightarrow{p} \sigma_R(\tau).$$

Next consider the restricted projected-score numerator. Let $\lambda_T(\tau)$ denote the Lagrange multiplier in the restricted quantile-regression first-order conditions. The same linearized KKT system used in the proof of Proposition 2 gives

$$0 = \mathbb{G}_{T,1}(\tau) - D(\tau)\Delta^r_{T,1}(\tau) - R'\lambda_T(\tau) + o_p(1), \qquad R\Delta^r_{T,1}(\tau) = -\delta,$$

and therefore

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T} \hat z_{t,0}(\tau) = R'\lambda_T(\tau) + o_p(1) = O_p(1).$$

Premultiplying the linearized KKT equation by $RD(\tau)^{-1}$ yields

$$RD(\tau)^{-1}R'\lambda_T(\tau) = RD(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + \delta + o_p(1).$$

Hence

$$RD(\tau)^{-1}\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\hat{z}_{t,0}(\tau) = RD(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + \delta + o_p(1).$$

The unrestricted Bahadur expansion gives

$$N_T(\tau) = \sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big) = \delta + RD(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + o_p(1),$$

so the last display implies

$$RD(\tau)^{-1}\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\hat{z}_{t,0}(\tau) = N_T(\tau) + o_p(1).$$

Using $\|\hat{D}_0(\tau)^{-1} - D(\tau)^{-1}\| = o_p(T^{-1/4})$ together with $T^{-1/2}\sum_t \hat{z}_{t,0}(\tau) = O_p(1)$, we obtain the explicit Jacobian substitution step

$$\begin{aligned}
N_{T,0}(\tau) &= R\hat{D}_0(\tau)^{-1}\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\hat{z}_{t,0}(\tau) \\
&= RD(\tau)^{-1}\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\hat{z}_{t,0}(\tau) + R\big(\hat{D}_0(\tau)^{-1} - D(\tau)^{-1}\big)\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\hat{z}_{t,0}(\tau) \\
&= N_T(\tau) + o_p(1).
\end{aligned}$$

For the restricted denominator, define the local-alternative projected score

$$g_{t,T}(\tau) = RD(\tau)^{-1}x_t^{*}\psi_\tau(u_{t,T}(\tau)).$$

Because $\hat{\beta}_0^r(\tau) - \beta_{0,T}(\tau) = O_p(T^{-1/2})$, the restricted residuals satisfy

$$\hat{u}_{t,0}(\tau) - u_{t,T}(\tau) = -x_t^{*\prime}\big(\hat{\beta}_0^r(\tau) - \beta_{0,T}(\tau)\big).$$

Thus the score indicator can differ from its local-alternative counterpart only when $u_{t,T}(\tau)$ falls in a shrinking neighborhood of zero with width of order $T^{-1/2}\|x_t^*\|$. The bounded density in Assumption 2, the moment condition in Assumption 7, and the rate $\max\{S_T, \ell_T\} = o(T^{1/2})$ therefore imply that the lag-window covariance approximation from Lemma 1 continues to hold under $\mathbb{H}_{1,T}$:

$$\max_{0\le j\le \ell_T}\left|\frac{1}{T}\sum_{t=j+1}^{T}\hat{g}_{t,0}(\tau)\hat{g}_{t-j,0}(\tau)-E[g_{t,T}(\tau)g_{t-j,T}(\tau)]\right|\xrightarrow{p}0.$$

The local perturbation changes the mean of $g_{t,T}(\tau)$ by $O(T^{-1/2})$ but does not change its long-run variance at first order, so the same HAC consistency argument as in Lemma 1 yields

$$s_{T,0}(\tau)\xrightarrow{p}\sigma_R(\tau),\qquad P^*(|\hat{s}_T^*(\tau)-\sigma_R(\tau)|>\epsilon)\xrightarrow{p}0$$

for every $\epsilon>0$.

Finally, because the bootstrap statistic is defined directly from the restricted projected scores, the bootstrap Gaussian approximation under the Pitman sequence follows from the same three-step argument as in Theorem 2. Let

$$g^{\circ}_{t,T}(\tau)=g_{t,T}(\tau)-\frac{1}{T}\sum_{s=1}^{T}g_{s,T}(\tau),\qquad \delta^{\mathrm{loc}}_{t,T}(\tau)=\tilde{g}_{t,0}(\tau)-g^{\circ}_{t,T}(\tau).$$

The covariance approximation displayed above implies

$$\mathcal{L}_T\big(\{\delta^{\mathrm{loc}}_{t,T}(\tau)\}_{t=1}^{T};a,\ell_T\big)=o_p(1),$$

so the corresponding multiplier remainder is $o_{P^*}(1)$ in probability. Since the local perturbation is of order $T^{-1/2}$, the long-run variance of $g^{\circ}_{t,T}(\tau)$ remains $\sigma_R^2(\tau)$ at first order, and Assumptions 5 and 7 therefore yield

$$\sup_{z\in\mathbb{R}}|P^*(N_T^*(\tau)\le z)-\Phi(z/\sigma_R(\tau))|\xrightarrow{p}0.$$

Combining this conditional CLT with the bootstrap studentizer consistency displayed above gives

$$\sup_{z\in\mathbb{R}}|P^*(T_T^*(\tau)\le z)-\Phi(z)|\xrightarrow{p}0.$$

Equivalently, if $F_T^*(x) = P^*(|T_T^*(\tau)| \le x)$ and $F_{|Z|}(x) = P(|Z| \le x)$ for $Z \sim N(0,1)$, then

$$\sup_{x \in \mathbb{R}} \left| F_T^*(x) - F_{|Z|}(x) \right| \xrightarrow{p} 0.$$

Since $F_{|Z|}$ is continuous and strictly increasing on $[0, \infty)$, its $(1-\alpha)$-quantile is $z_{1-\alpha/2}$, and therefore

$$c_{1-\alpha}^*(\tau) \xrightarrow{p} z_{1-\alpha/2}.$$

■

A.9 Proof of Theorem 3

Proof. Under the contiguous Pitman sequence $\mathbb{H}_{1,T} : R\beta_{0,T}(\tau) = r_0 + \delta/\sqrt{T}$, the unrestricted local-alternative expansion in Proposition 2 gives

$$\sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big) = \delta + RD(\tau)^{-1}\mathbb{G}_{T,1}(\tau) + o_p(1).$$

The same strong-mixing central limit theorem used in the proof of Theorem 1 applies to the local-alternative score process $x_t^* \psi_\tau(u_{t,T}(\tau))$, and the $T^{-1/2}$ perturbation in $\beta_{0,T}(\tau)$ does not alter the first-order long-run variance. Hence

$$RD(\tau)^{-1}\mathbb{G}_{T,1}(\tau) \xrightarrow{d} N\big(0, \sigma_R^2(\tau)\big).$$

Lemma 2 gives $\hat{s}_T^r(\tau) \xrightarrow{p} \sigma_R(\tau)$, so Slutsky's lemma yields

$$T_T^r(\tau) = \frac{\sqrt{T}(R\hat{\beta}(\tau) - r_0)}{\hat{s}_T^r(\tau)} \xrightarrow{d} N(\delta/\sigma_R(\tau), 1).$$

The same lemma also gives

$$N_{T,0}(\tau) = N_T(\tau) + o_p(1), \qquad s_{T,0}(\tau) \xrightarrow{p} \sigma_R(\tau), \qquad \hat{s}_T^r(\tau) \xrightarrow{p} \sigma_R(\tau),$$

where

$$N_T(\tau) = \sqrt{T}\big(R\hat{\beta}(\tau) - r_0\big), \qquad N_{T,0}(\tau) = \frac{1}{\sqrt{T}}\sum_{t=1}^{T}\hat{g}_{t,0}(\tau), \qquad s_{T,0}(\tau) = \sqrt{\mathcal{L}_T^{\mathrm{B}}(\{\hat{g}_{t,0}(\tau)\}_{t=1}^{T}; \ell_T)}.$$

Hence $N_T(\tau) = O_p(1)$ and

$$T_{T,0}^r(\tau) - T_T^r(\tau) = \frac{N_{T,0}(\tau) - N_T(\tau)}{s_{T,0}(\tau)} + N_T(\tau)\frac{\hat{s}_T^r(\tau) - s_{T,0}(\tau)}{s_{T,0}(\tau)\hat{s}_T^r(\tau)} = o_p(1).$$

Lemma 2 further gives

$$c_{1-\alpha}^*(\tau) \xrightarrow{p} z_{1-\alpha/2}.$$

Combining this with $T_{T,0}^r(\tau) \xrightarrow{d} N(\delta/\sigma_R(\tau), 1)$ and applying Slutsky's lemma again,

$$P\big(\big|T_{T,0}^r(\tau)\big| > c_{1-\alpha}^*(\tau)\big) \to P\big(|Z + \delta/\sigma_R(\tau)| > z_{1-\alpha/2}\big), \qquad Z \sim N(0,1).$$

Evaluating the Gaussian tail probability gives

$$P\big(|Z + \delta/\sigma_R(\tau)| > z_{1-\alpha/2}\big) = 1 - \Phi\big(z_{1-\alpha/2} - \delta/\sigma_R(\tau)\big) \; + \; \Phi\big(-z_{1-\alpha/2} - \delta/\sigma_R(\tau)\big).$$

At $\delta = 0$ this equals α, and for every $\delta \neq 0$ it is strictly larger than α, proving the stated non-trivial local power. ■

A.10 Proof of Corollary 3

Proof. For the observed feasible HAC statistic, Theorem 1 already allows the random bandwidth selector $\hat{S}_T$ through Assumption 4, which imposes $\hat{S}_T/S_T \xrightarrow{p} 1$. Thus only the bootstrap dependence bandwidth requires an additional argument.

Let

$$\underline{\ell}_T = \lceil \underline{c}T^{\kappa} \rceil, \qquad \bar{\ell}_T = \lceil \bar{c}T^{\kappa} \rceil,$$

where $\underline{c}$ and $\bar{c}$ are the deterministic constants from Assumption 5. By that assumption,

$$P\Big(\underline{\ell}_T \le \hat{\ell}_T \le \bar{\ell}_T\Big) \to 1,$$

and both bounding sequences satisfy

$$\underline{\ell}_T \to \infty, \qquad \bar{\ell}_T \to \infty, \qquad \bar{\ell}_T/T \to 0, \qquad \bar{\ell}_T = o(T^{1/2}).$$

Hence, with probability approaching one, the selected dependence bandwidth belongs to a deterministic interval of admissible bandwidths having exactly the same rate properties as the fixed sequence used in Theorems 2 and 3.

Conditional on the sample, $\hat{\ell}_T$ is just a fixed integer. Every place where ℓ_T enters Lemma 1, Theorem 2, and Theorem 3 uses only the facts that the bandwidth diverges, remains $o(T^{1/2})$, and generates uniformly bounded Bartlett weights. The covariance and variance bounds in the proof of Lemma 1 are therefore uniform for $\ell \in [\underline{\ell}_T, \bar{\ell}_T]$, because the projected-score autocovariances and fourth-order cumulants are absolutely summable by Assumption 7. Consequently, replacing ℓ_T by $\hat{\ell}_T$ leaves unchanged the conclusions

$$\text{for every } \epsilon > 0, \qquad P^*(|\hat{s}_T^*(\tau) - \sigma_R(\tau)| > \epsilon) \xrightarrow{p} 0, \qquad \tilde{s}_T^r(\tau) \xrightarrow{p} \sigma_R(\tau),$$

and likewise preserves the equivalence of the observed matrix-based and projected-score studentizers.

The same observation applies to the conditional multiplier central limit argument in Theorem 2. Assumption 5 is stated directly for the feasible selector $\hat{\ell}_T$, so once the selected bandwidth is treated as fixed conditional on the sample, the proof of Theorem 2 goes through verbatim with ℓ_T replaced by $\hat{\ell}_T$. Hence

$$\sup_{z\in\mathbb{R}} \left| P^*(T_T^*(\tau) \le z) - P\big(T_{T,0}^r(\tau) \le z\big) \right| \xrightarrow{p} 0$$

still holds, and all consequences that rely on this approximation, including the bootstrap critical-value and local-power conclusions in Theorem 3, remain valid. Since $P(\underline{\ell}_T \le \hat{\ell}_T \le \bar{\ell}_T) \to 1$, these arguments, conditional on the selected bandwidth, yield the same unconditional conclusions.

Therefore the conclusions of Theorems 1–3 continue to hold when the deterministic tuning sequences are replaced by the measurable selectors $\hat{S}_T$ and $\hat{\ell}_T$ satisfying Assumptions 4 and 5.

■

B Additional Simulation Results

B.1 Selection of Plug-In Dependence Selector

This section reports the mean selected block length $\hat{\ell}_T$ for the plug-in DWB selector used in Models 1 and 2 in Section 5. Note that the plug-in selector (3) estimates the AMSE-optimal scale constant c from the data using an AR(1) pilot on the scalar projected score, and the resulting estimate is clipped to the interval $[\underline{c}, \bar{c}] = [0.5, 5]$. That is, the feasible dependence bandwidth lies between $\lceil 0.5\, T^{1/3} \rceil$ and $\lceil 5\, T^{1/3} \rceil$.

Table A1 shows that the selected dependence bandwidth increases sharply with both dependence and sample size. When $\rho = 0$, the average selected block length remains close to the lower clipping bound, ranging from about 3 at $T = 100$ to about 5 at $T = 1000$ in both models. When dependence strengthens, by contrast, the selected bandwidth becomes wider and grows to 23–27 when $\rho = 0.9$ and $T = 1000$. For a given level of dependence, the median quantile ($\tau = 0.5$) tends to produce slightly larger selected bandwidth than the upper tail ($\tau = 0.9$).

Table A1: Average selected plug-in DWB bandwidth $\hat{\ell}_T$: Models 1 & 2

		Model 1			Model 2		
T	τ	$\rho = 0$	$\rho = 0.5$	$\rho = 0.9$	$\rho = 0$	$\rho = 0.5$	$\rho = 0.9$
100	0.5	3.10	5.04	11.77	3.12	4.90	12.17
	0.9	3.16	4.30	8.18	3.21	4.27	8.18
200	0.5	3.12	6.04	15.51	3.10	6.05	15.56
	0.9	3.10	4.76	11.79	3.18	5.15	11.87
500	0.5	4.02	7.94	21.48	4.02	7.95	21.42
	0.9	4.04	5.99	18.19	4.06	6.59	18.09
1000	0.5	5.00	9.98	27.13	5.00	9.89	27.16
	0.9	5.00	7.61	23.21	5.02	7.93	23.36

Notes. This table reports the average selected plug-in DWB bandwidth $\hat{\ell}_T$ using Models 1 and 2 defined in Section 5. The plug-in selector formulated by (3) estimates the scale constant c in the bandwidth rule $\ell_T = \lceil c\, T^{1/3} \rceil$ and clips it to the interval $[\underline{c}, \bar{c}] = [0.5, 5]$.

B.2 Robustness Check: Left-Skewed Heavy-Tailed Innovations

This robustness test adopts the exact univariate Models 1 and 2 from Section 5 but replaces the Gaussian disturbance innovations with standardized left-skewed, heavy-tailed skew-t draws. Specifically, the innovation distribution is normalized to have mean zero and variance one, with degrees of freedom $\nu = 5$ and skewness parameter -4. In Model 1, the AR(1) disturbance becomes $u_t = \rho u_{t-1} + e_t$, where e_t is the standardized skew-t shock scaled by $\sqrt{1 - \rho^2}$. In Model 2, the AR(1)-GARCH(1,1) recursion is left unchanged, but the Gaussian innovation z_t in $\eta_t = \sigma_t z_t$ is replaced by the same standardized skew-t shock. Thus, relative to Section 5, the robustness check preserves the serial dependence and conditional variance structure while introducing pronounced left-skewness and heavy tails into the disturbance process.

The comparison of size control and power across the six methods is presented in Figures A1 and A2. Overall, the skew-t robustness checks yield results similar to those in the main text: the dependent wild bootstrap continues to show good size control under non-Gaussian dependence without evident power losses.

Figure A1: Empirical size comparison: Models 1 & 2 under left-skewed standardized t disturbance

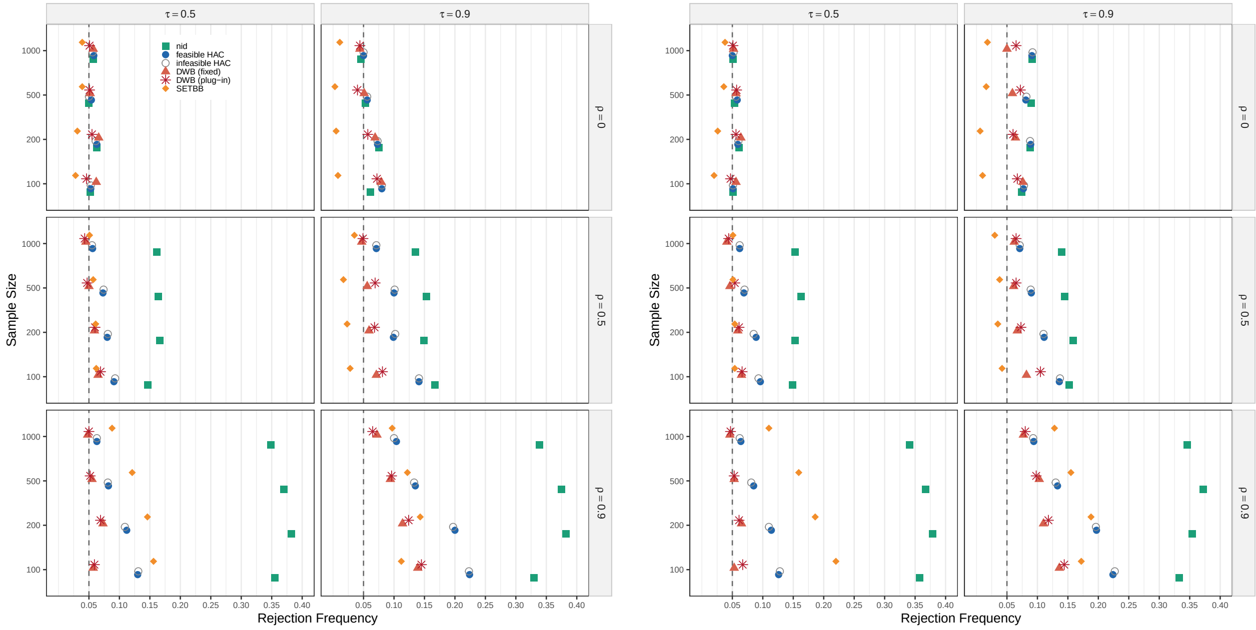

(a) Model 1 **(b)** Model 2

Notes. This figure reports empirical rejection rates of testing the null hypothesis $\mathbb{H}_0 : \beta_1(\tau) = 1$ by the conventional QR test based on the `nid` covariance estimator from the R package `quantreg`, the HAC estimator using the feasible plug-in bandwidth selector of Galvao and Yoon (2024) (feasible-HAC), the corresponding HAC estimator using the true autocorrelation coefficient in bandwidth selection (infeasible-HAC), the smooth extended tapered block bootstrap (SETBB) method of Gregory et al. (2018), and the proposed dependent wild bootstrap procedure with fixed scale constant $c = 2$ (fixed DWB) and plug-in selector in (3) with $c \in [0.5, 5]$ (plug-in DWB). The data generating processes Model 1 and Model 2 are defined in Section 5. The degree of autocorrelation is determined by $\rho \in \{0, 0.5, 0.9\}$, and the slope coefficient is tested at the quantiles $\tau \in \{0.5, 0.9\}$. The sample size $T = 1000$. The empirical rejection rate is computed at the 5% nominal level based on 1000 replications and 499 bootstrap draws.

Figure A2: Power comparison: Models 1 & 2 under left-skewed standardized t disturbance

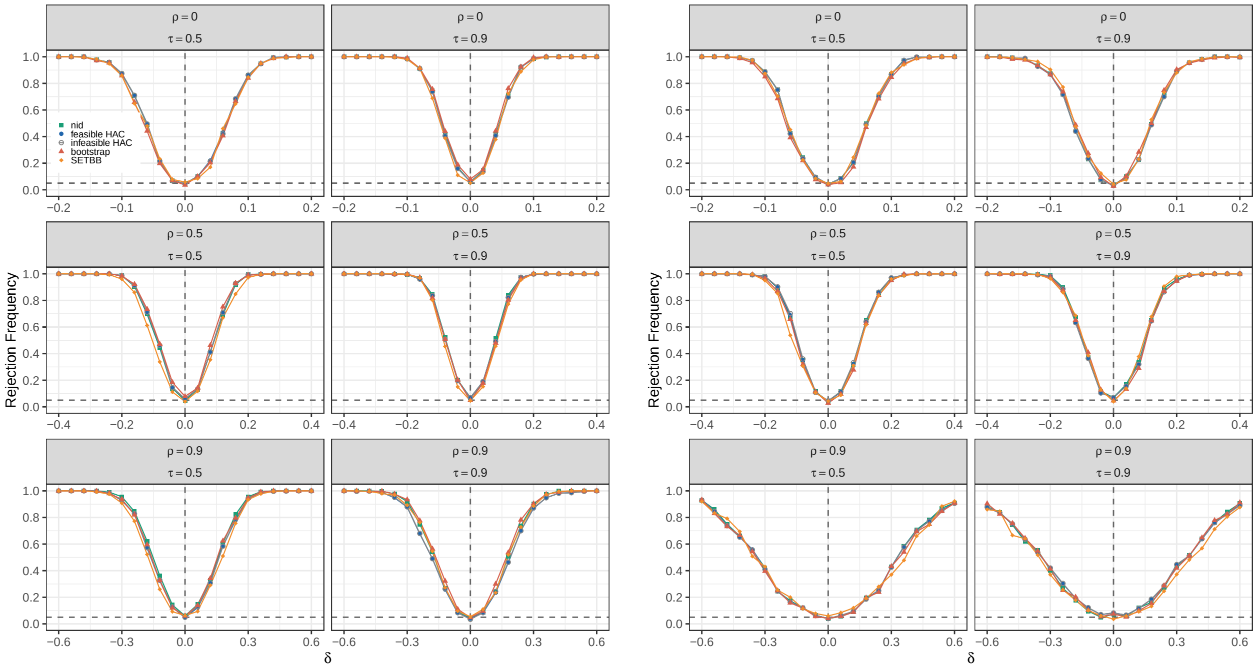

(a) Model 1 **(b)** Model 2

Notes. This figure reports empirical rejection rates of testing the null hypothesis $\mathbb{H}_0 : \beta_1(\tau) = 1 + \delta$ by the conventional QR test based on the `nid` covariance estimator from the R package `quantreg`, the HAC estimator using the feasible plug-in bandwidth selector of Galvao and Yoon (2024) (feasible-HAC), the corresponding HAC estimator using the true autocorrelation coefficient in bandwidth selection (infeasible-HAC), the smooth extended tapered block bootstrap (SETBB) method of Gregory et al. (2018), and the proposed dependent wild bootstrap procedure (bootstrap). The data generating processes Model 1 and Model 2 are defined in Section 5. δ determines the deviation of the slope coefficient from 1. The degree of autocorrelation is determined by $\rho \in \{0, 0.5, 0.9\}$, and the slope coefficient is tested at the quantiles $\tau \in \{0.5, 0.9\}$. The sample size $T = 1000$. The empirical rejection rate is computed at the 5% nominal level based on 1000 replications and 499 bootstrap draws.

B.3 Robustness Check: Multiplier Choice under i.i.d. Data

To address the special case in which the data are independent, we revisit Model 1 under $\rho = 0$ and $\rho_x = 0$, so that both the regressor and the disturbance are serially uncorrelated and the sample is i.i.d. We compare three bootstrap implementations: the benchmark Gaussian innovation and Bartlett taper adopted in the main text, a Bartlett taper driven by two-point Rademacher innovations, and an independent Rademacher multiplier corresponding to $a(x) = \mathbf{1}(x = 0)$ and $\ell_T = 1$. Table A2 reports the empirical rejection frequencies at the 5% nominal level for $\tau \in \{0.5, 0.9\}$ and $T \in \{100, 500, 1000\}$.

Table A2 reveals three observations. First, both fixed and plug-in DWB procedures still produce similar results across all cases. Second, holding the Bartlett taper fixed, replacing Gaussian innovations by Rademacher innovations has only a small effect on rejection frequencies, which suggests that the innovation distribution is a secondary finite-sample choice relative to the dependence structure encoded by $a(\cdot)$. Third, when the data are truly i.i.d., imposing a non-zero Bartlett taper changes finite-sample size relative to the independent-multiplier benchmark, but the direction is not uniform across combinations of (T, τ). In the upper-tail small-sample case $(T, \tau) = (100, 0.9)$, both Bartlett implementations over-reject relative to the independent benchmark, whereas at $(T, \tau) = (1000, 0.5)$ the Gaussian-Bartlett choice is slightly more conservative and the Rademacher-Bartlett version is closer to nominal. Overall, the discrepancies across the three choices are small, and by $T = 1000$ all three implementations exhibit reasonably good size control. This finding supports the discussion in Remark 1 of the main text: the independent multiplier is the natural benchmark in the i.i.d. case, but the Bartlett taper choice remains empirically applicable when used in that setting.

Table A2: Empirical size comparison: i.i.d. data and alternative multiplier choices

		Gaussian-Bartlett		Rademacher-Bartlett		Rademacher-Independent	
T	τ	fixed	plug-in	fixed	plug-in	fixed	plug-in
100	0.5	0.048	0.052	0.046	0.056	0.036	0.048
	0.9	0.080	0.076	0.080	0.074	0.072	0.068
500	0.5	0.050	0.058	0.050	0.058	0.058	0.062
	0.9	0.050	0.056	0.050	0.052	0.048	0.050
1000	0.5	0.062	0.040	0.054	0.044	0.046	0.052
	0.9	0.056	0.052	0.056	0.056	0.048	0.046

Notes. This table reports empirical rejection rates under i.i.d. data ($\rho = 0$, $\rho_x = 0$, Model 1) at the 5% nominal level. Three multiplier variants are compared under both the fixed ($c = 2$) and plug-in bandwidth DWB: Gaussian multiplier and Bartlett taper (Gaussian-Bartlett), two-point Rademacher multiplier and Bartlett taper (Rademacher-Bartlett), and independent Rademacher multiplier (Rademacher-Independent).

B.4 Robustness Check: Alternative Multiplier Tapers

This robustness check examines whether the bootstrap's finite-sample behavior is sensitive to the choice of taper function and distribution family of multiplier draw. We revisit the exact DGPs of Models 1 and 2 in Section 5 and recompute rejection frequencies of dependent wild bootstrap under combinations of Gaussian and Rademacher innovations and Bartlett, Parzen, and quadratic spectrum (QS) tapers: Gaussian-Bartlett (G-B), Gaussian-Parzen (G-P), Gaussian-QS (G-QS), Rademacher-Bartlett (R-B), Rademacher-Parzen (R-P), and Rademacher-QS (R-QS). We conduct both fixed DWB and plug-in DWB as in the main text, and consider $T \in \{100, 500, 1000\}$, $\rho \in \{0, 0.5, 0.9\}$, and $\tau \in \{0.5, 0.9\}$. For the Parzen and QS tapers, we replace the Bartlett moving-average filter with kernel-specific filter weights and studentize each bootstrap draw with the matching scalar lag-window kernel. For the QS filter, which has infinite support, we truncate the weight sequence at $12\ell_T$.

The comparison of size control across the six combinations, as displayed in Table A3, suggests three observations. First, fixed DWB and plug-in DWB exhibit performance quite similar to each other, echoing the main findings from the main text. Second, given a specific taper, replacing Gaussian innovations with Rademacher innovations has only a modest effect on rejection frequencies regardless of the pair (T, ρ, τ). Third, no taper dominates uniformly or exhibits universally excellent size control across scenarios. To summarize, the proposed dependent wild bootstrap procedure appears reasonably robust to the choice of innovations and taper functions that satisfy Assumption 5.

B.5 Power Comparison under Multivariate Scenario

Figure A3 presents the size-adjusted rejection frequencies for testing the null hypotheses $\mathbb{H}_0 : \beta_1(\tau) = 1 + \tau$, $\mathbb{H}_0 : \beta_2(\tau) = 1 + \tau$, and $\mathbb{H}_0 : \beta_3(\tau) = 1 + \tau$ in panels (a), (b), and (c), respectively, for the multivariate quantile regression $y_t = \beta_0 + \beta_1 x_{1t} + \beta_2 x_{2t} + \beta_3 x_{3t} + u_t$ with the model setup defined in Section 5. The sample size is $T = 1000$ in all three panels. As in the univariate scenario, power converges to one for all five methods as $|\delta|$ increases, and the differences are generally modest, with the five curves again tracking each other closely.

Table A3: Empirical size comparison: Alternative choices of multiplier and taper

			G-B		G-P		G-QS		R-B		R-P		R-QS	
T	τ	ρ	fixed	plug-in	fixed	plug-in	fixed	plug-in	fixed	plug-in	fixed	plug-in	fixed	plug-in
Panel A: Model 1														
100	0.5	0	0.060	0.056	0.056	0.060	0.046	0.050	0.054	0.054	0.054	0.050	0.040	0.054
		0.5	0.074	0.096	0.074	0.092	0.068	0.084	0.074	0.090	0.074	0.096	0.072	0.086
		0.9	0.090	0.084	0.088	0.082	0.066	0.066	0.078	0.076	0.084	0.076	0.066	0.060
	0.9	0	0.058	0.044	0.050	0.054	0.050	0.048	0.056	0.048	0.046	0.048	0.056	0.046
		0.5	0.074	0.072	0.070	0.072	0.066	0.068	0.070	0.072	0.072	0.070	0.060	0.074
		0.9	0.120	0.144	0.126	0.140	0.120	0.122	0.124	0.140	0.124	0.144	0.122	0.128
500	0.5	0	0.070	0.062	0.066	0.068	0.068	0.062	0.070	0.058	0.068	0.056	0.072	0.058
		0.5	0.032	0.036	0.032	0.036	0.030	0.038	0.028	0.032	0.032	0.038	0.024	0.036
		0.9	0.046	0.050	0.054	0.058	0.042	0.048	0.046	0.052	0.054	0.050	0.048	0.046
	0.9	0	0.046	0.056	0.042	0.056	0.038	0.056	0.046	0.056	0.040	0.052	0.040	0.058
		0.5	0.058	0.058	0.056	0.056	0.048	0.052	0.056	0.064	0.054	0.062	0.056	0.056
		0.9	0.084	0.086	0.092	0.090	0.082	0.086	0.080	0.086	0.084	0.086	0.078	0.088
1000	0.5	0	0.072	0.064	0.076	0.060	0.058	0.070	0.068	0.062	0.074	0.060	0.072	0.064
		0.5	0.036	0.030	0.028	0.034	0.036	0.032	0.030	0.030	0.032	0.034	0.030	0.036
		0.9	0.046	0.050	0.048	0.044	0.040	0.040	0.044	0.042	0.050	0.042	0.042	0.046
	0.9	0	0.056	0.048	0.050	0.044	0.056	0.044	0.050	0.048	0.044	0.050	0.056	0.042
		0.5	0.066	0.068	0.060	0.072	0.070	0.064	0.066	0.068	0.064	0.068	0.058	0.062
		0.9	0.076	0.082	0.080	0.080	0.072	0.072	0.076	0.082	0.080	0.082	0.074	0.074
Panel B: Model 2														
100	0.5	0	0.058	0.050	0.050	0.044	0.054	0.042	0.058	0.046	0.054	0.046	0.054	0.040
		0.5	0.076	0.070	0.072	0.076	0.070	0.066	0.078	0.076	0.072	0.076	0.068	0.064
		0.9	0.062	0.064	0.078	0.066	0.054	0.052	0.062	0.062	0.066	0.064	0.048	0.052
	0.9	0	0.072	0.068	0.072	0.068	0.062	0.062	0.074	0.070	0.066	0.068	0.062	0.072
		0.5	0.092	0.082	0.084	0.080	0.074	0.080	0.090	0.084	0.084	0.090	0.076	0.088
		0.9	0.158	0.182	0.156	0.180	0.154	0.176	0.156	0.174	0.160	0.182	0.152	0.180
500	0.5	0	0.064	0.058	0.056	0.050	0.052	0.050	0.058	0.054	0.050	0.048	0.052	0.056
		0.5	0.046	0.046	0.044	0.038	0.040	0.042	0.040	0.038	0.042	0.050	0.034	0.040
		0.9	0.058	0.068	0.056	0.066	0.052	0.056	0.064	0.062	0.064	0.058	0.052	0.058
	0.9	0	0.038	0.050	0.042	0.058	0.040	0.050	0.034	0.056	0.040	0.054	0.040	0.050
		0.5	0.046	0.056	0.042	0.060	0.046	0.048	0.046	0.050	0.046	0.052	0.040	0.040
		0.9	0.096	0.106	0.102	0.104	0.100	0.106	0.100	0.104	0.106	0.102	0.096	0.102
1000	0.5	0	0.048	0.056	0.052	0.048	0.050	0.046	0.052	0.054	0.052	0.050	0.054	0.048
		0.5	0.042	0.048	0.038	0.046	0.036	0.046	0.040	0.048	0.048	0.052	0.040	0.052
		0.9	0.058	0.064	0.060	0.056	0.052	0.054	0.062	0.058	0.064	0.056	0.056	0.060
	0.9	0	0.052	0.060	0.048	0.064	0.054	0.060	0.050	0.058	0.050	0.068	0.052	0.060
		0.5	0.056	0.064	0.058	0.060	0.050	0.058	0.054	0.068	0.054	0.062	0.052	0.056
		0.9	0.084	0.090	0.086	0.082	0.084	0.076	0.086	0.080	0.088	0.084	0.084	0.078

Notes. This table reports empirical rejection rates at the 5% nominal level under both the fixed ($c = 2$) and plug-in bandwidth DWB. Six combinations of multipliers and taper functions are considered: Gaussian-Bartlett (G-B), Gaussian-Parzen (G-P), Gaussian-Quadratic Spectrum (G-QS), Rademacher-Bartlett (R-B), Rademacher-Parzen (R-P), and Rademacher-Quadratic Spectrum (R-QS). The upper and lower panels use the exact Models 1 and 2 DGPs from Section 5, respectively.

Figure A3: Power comparison: Multivariate quantile regression

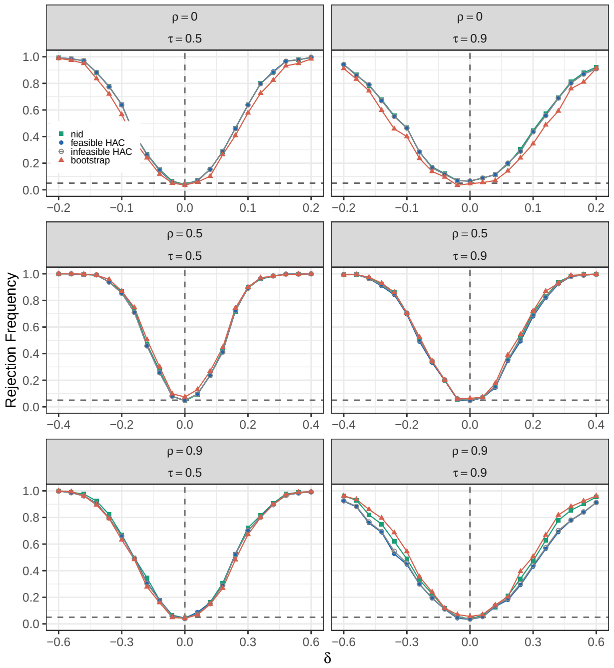

(a) $\mathbb{H}_0 : \beta_1(\tau) = 1$

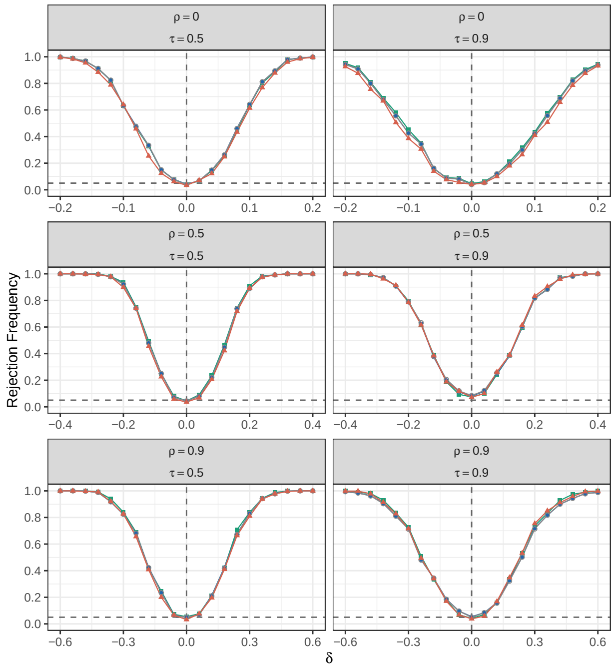

(b) $\mathbb{H}_0 : \beta_2(\tau) = 1$

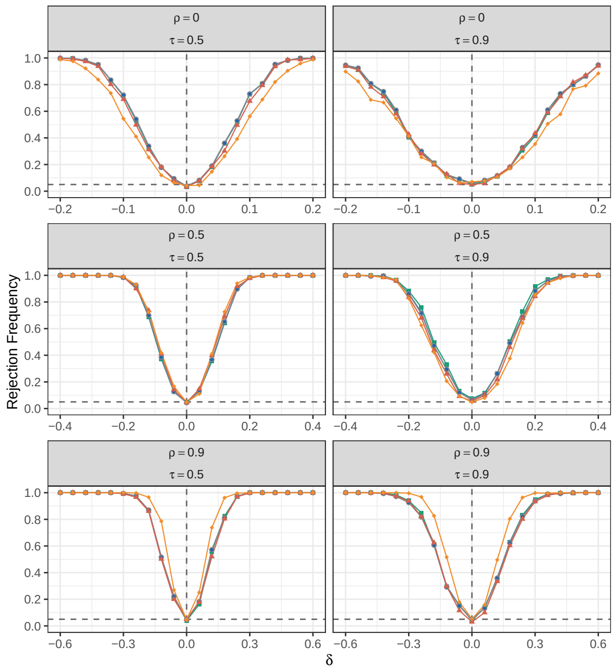

(c) $\mathbb{H}_0 : \beta_3(\tau) = 1$

Notes. This figure reports empirical rejection rates for testing the null hypothesis $\mathbb{H}_0 : \beta_j(\tau) = 1 + \delta$, where $j \in \{1, 2, 3\}$ indexes the slope coefficient under consideration in the multivariate Model 1 design. The four procedures are the conventional QR test based on the `nid` covariance estimator from the R package `quantreg`, the HAC test using the feasible plug-in bandwidth selector of Galvao and Yoon (2024) (feasible-HAC), the corresponding HAC test using the true autocorrelation coefficients in bandwidth selection (infeasible-HAC), the smooth extended tapered block bootstrap (SETBB) method of Gregory et al. (2018), and the proposed dependent wild bootstrap procedure with fixed scale constant $c = 2$ (fixed DWB) and plug-in selector in (3) with $c \in [0.5, 5]$ (plug-in DWB). The data generating process extends Model 1 in Section 5 to three regressors, with $y_t = \beta_0 + \beta_1 x_{1t} + \beta_2 x_{2t} + \beta_3 x_{3t} + u_t$, where the regressors follow independent AR(1) processes with autoregressive coefficients 0.95, 0.85, and 0.53, respectively, and the disturbance follows the same AR(1) process as in Section 5. δ determines the deviation of the slope coefficients from 1. The disturbance autocorrelation is determined by $\rho \in \{0, 0.5, 0.9\}$, and inference is conducted at the quantiles $\tau \in \{0.5, 0.9\}$. The empirical rejection rates are computed at the 5% nominal level using 1000 Monte Carlo replications and 499 bootstrap draws.

C Replication: Electricity Price Application

As a supplementary empirical application, we revisit the empirical application in Galvao and Yoon (2024), who estimate conditional quantile functions for day-ahead electricity prices at peak Hour 17 (5 p.m.) using data from the California electricity market. The dependent variable is the log day-ahead market price (log DAM_{17}), with an unconditional mean of \$44.58 per MWh and a standard deviation of \$15.56. The sample contains $T = 1356$ daily observations. Five regressors are included: the log natural gas spot price (log Gas), the log greenhouse gas emission allowance price (log GHG), the log system-wide load forecast (log Demand), and the solar and wind generation forecasts in GW (Solar, Wind). Solar and wind are scaled by 10^{-3} so that all coefficient estimates are of comparable order. Gas and demand directly raise marginal generation cost and hence prices, while renewable supply displaces thermal units and lowers prices. The GHG allowance price reflects the carbon cost embedded in thermal dispatch.

We estimate the quantile regression

$$\begin{aligned} Q_{\log(\text{DAM}_{17,t})}(\tau \mid x_t) = \beta_0(\tau) + \beta_1(\tau)\log(\text{Gas}_t) + \beta_2(\tau)\log(\text{GHG}_t) \\ + \beta_3(\tau)\log(\text{Demand}_t) + \beta_4(\tau)\text{Solar}_t + \beta_5(\tau)\text{Wind}_t \end{aligned}$$

at $\tau \in \{0.05, 0.10, 0.20, \ldots, 0.90, 0.95\}$. Figure A4 plots the coefficient curves together with 95% confidence bands from five procedures: the conventional nid standard errors (shaded area), the feasible HAC with estimated bandwidth and quadratic spectrum kernel (dashed red), the SETBB method by Gregory et al. (2018) using 999 replications (dashed orange), and the dependent wild bootstrap with 999 bootstrap draws (dashed blue) using either fixed clipping constant $c = 2$ or plug-in selector (3). The replication broadly aligns with the findings of Galvao and Yoon (2024): the feasible HAC bands are uniformly wider than the nid bands, and relying solely on conventional standard errors in time series quantile regression can lead to misleading conclusions. The fixed and plug-in DWB bands closely track each other and are slightly wider than the HAC bands across most of the five panels, particularly in the tails, consistent with the simulation evidence on size control and suggesting more conservative procedures in this application.

Figure A4: Quantile regression coefficients: Electricity prices at peak Hour 17

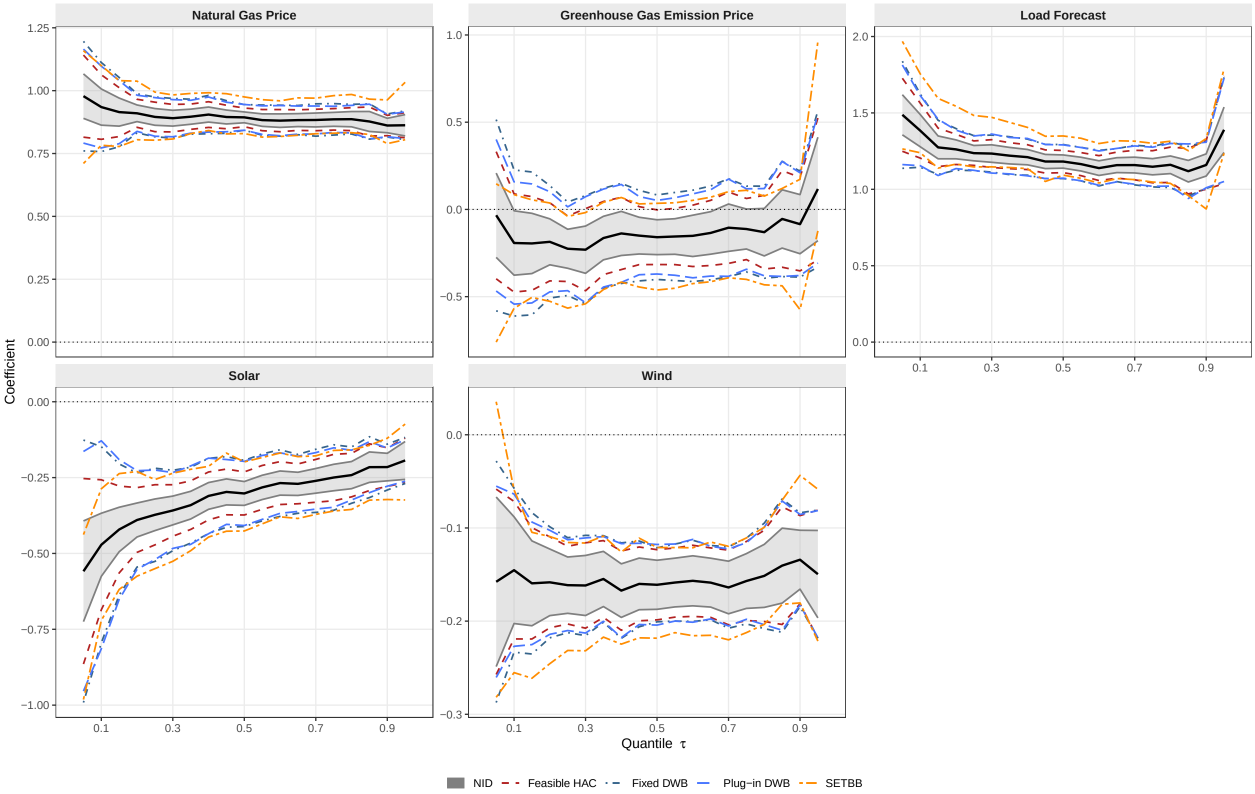

Notes: This figure presents the estimates of quantile regression $Q_{\log(\mathrm{DAM}_{17,t})}(\tau \mid x_t) = \beta_0(\tau) + \beta_1(\tau)\log(\mathrm{Gas}_t) + \beta_2(\tau)\log(\mathrm{GHG}_t) + \beta_3(\tau)\log(\mathrm{Demand}_t) + \beta_4(\tau)\mathrm{Solar}_t + \beta_5(\tau)\mathrm{Wind}_t$ at $\tau \in \{0.05, 0.10, \ldots, 0.90, 0.95\}$. DAM_{17} denotes the day-ahead market electricity price at peak Hour 17, Gas denotes the natural gas spot price, GHG denotes the greenhouse gas emission allowance price, Demand denotes the system-wide load forecast, and Solar and Wind denote the solar and wind generation forecasts, respectively, each scaled by 10^{-3}. The sample contains $T = 1356$ daily observations. Each panel corresponds to one regressor. The solid black curve is the point estimate $\hat{\beta}_j(\tau)$ for regressor j at quantile τ. The grey shaded region is the conventional standard errors-based band, the dashed red lines bound the feasible HAC-based band by Galvao and Yoon (2024), the dashed orange lines bound the SETBB-based band of Gregory et al. (2018) using 999 replications, and the two dashed blue lines respectively bound the fixed and plug-in DWB-based band using 999 multiplier draws. The dotted horizontal line denotes zero.