A consistent specification test for dynamic quantile models

Peter Horvath
Department of Economics, Duke University

Jia Li
Department of Economics, Duke University

Zhipeng Liao
Department of Economics, UCLA

Andrew J. Patton
Department of Economics, Duke University

Correct specification of a conditional quantile model implies that a particular conditional moment is equal to zero. We nonparametrically estimate the conditional moment function via series regression and test whether it is identically zero using uniform functional inference. Our approach is theoretically justified via a strong Gaussian approximation for statistics of growing dimensions in a general time series setting. We propose a novel bootstrap method in this nonstandard context and show that it significantly outperforms the benchmark asymptotic approximation in finite samples, especially for tail quantiles such as Value-at-Risk (VaR). We use the proposed new test to study the VaR and CoVaR (Adrian and Brunnermeier (2016)) of a collection of US financial institutions.

Keywords. Bootstrap, VaR, series regression, strong approximation.

JEL classification. C14, C22, C52.

1. Introduction

Quantile models allow the researcher to learn about the location of a given variable, when the probability level of the quantile is near one-half, or about the tails of the variable, when the probability level is near zero or one. In the former case, quantiles represent a robust alternative to the mean, reducing the sensitivity to a few large observations. In the latter case, quantiles are used to measure the risk of a variable. Given their many uses, the literature on quantile models, which began with Koenker and Basset (1978), is now voluminous; see Koenker (2005) and Komunjer (2013) for recent reviews. Quantiles

that lie in the tails of profit and loss distributions are usually given the moniker "Value-at-Risk" (VaR). This risk measure is at the center of the Basel accords (Basel Committee (2010)) on banking supervision, which guide regulatory policies in 28 countries or jurisdictions around the world. Moreover, trading desks and regulators monitor daily risk exposures using VaR, and it has become a mainstay of the risk management industry; see McNeil, Frey, and Embrechts (2015) for example. The results presented in this paper apply to all dynamic quantile models, but they are particularly relevant for dynamic VaR models.

The use of parametric models for conditional quantiles naturally leads to a need for specification tests for these models, as a misspecified model can lead to erroneous policy decisions or suboptimal predictions. Denoting the variable of interest Y_{t+1} and the information set $\mathcal{F}_t$, the conditional q-quantile of Y_{t+1} given $\mathcal{F}_t$, henceforth denoted as f_t, satisfies $\mathbb{P}(Y_{t+1} \le f_t|\mathcal{F}_t) = q$, which is equivalent to the conditional moment restriction:

$$\mathbb{E}[1_{\{Y_{t+1} \le f_t\}} - q|\mathcal{F}_t] = 0.$$

In applications, researchers are often interested in some specific conditioning variable X_t, which may be a vector, in the information set. For any X_t adapted to $\mathcal{F}_t$, the above equation implies

$$\mathbb{E}[1_{\{Y_{t+1} \le f_t\}} - q|X_t = x] = 0, \quad \text{for all } x \in \mathcal{X}, \tag{1.1}$$

where $\mathcal{X}$ is the support of X_t. A consistent specification test can be obtained by estimating the conditional expectation function on the left-hand side of equation (1.1) and testing whether it is identically zero. This inference problem is nontrivial because it concerns the global, instead of local, behavior of the conditional expectation function. In a recent paper, Li and Liao (2020) proposed a uniform nonparametric inference method based on series regression for general time-series data. Under this approach, one can estimate the conditional moment function by regressing $1_{\{Y_{t+1} \le f_t\}} - q$ on an asymptotically growing number of approximating functions of X_t. Because of the growing dimension, the asymptotic problem is non-Donsker, that is, the functional estimator does not admit a functional central limit theorem. The approach of Li and Liao (2020) instead relies on a strong Gaussian approximation theory, which can be used to characterize the asymptotic properties of the test statistic; also see Chernozhukov, Lee, and Rosen (2013) and Belloni, Chernozhukov, Chetverikov, and Kato (2015) for applications of the strong approximation technique in microeconometric contexts.

Two theoretical extensions of existing work are required for the analysis in this paper. First, the inference procedure proposed by Li and Liao (2020) is directly based on an asymptotic Gaussian approximation. However, in a realistically calibrated Monte Carlo experiment (see Section 3) we find that such an asymptotic approximation works well only for quantiles near the middle of the distribution, while it suffers from substantial size distortion for quantiles in the tails. As VaR-type applications often involve probability levels of 95% or above, this is problematic. We overcome this issue by proposing a novel bootstrap method for computing the critical values for our test statistic, and we

establish its asymptotic validity. This bootstrap theory is nonstandard because it concerns uniform inference on the nonparametric series estimator, which appears to be new in time-series analysis; in particular, the non-Donsker issue in the original problem also manifests in the "bootstrap world." The proposed bootstrap procedure is easy to implement, and our Monte Carlo analysis shows that the bootstrap has satisfactory size control in realistic scenarios.

The second theoretical extension required for our analysis pertains to the presence of generated variables. First, note that the variable $1_{\{Y_{t+1} \le f_t\}} - q$ in the conditional moment restriction (1.1) depends in a nondifferentiable way on estimated parameters through the conditional quantile f_t. To address this, we provide sufficient conditions under which the preliminary estimation error is asymptotically negligible for our nonparametric testing problem. Intuitively, the estimated parameters typically converge at a parametric rate, and hence, the resulting error is negligible compared with the statistical noise in our uniform nonparametric inference. Although the intuition is straightforward, the formal theoretical justification is nontrivial due to the technical interaction between the nonsmoothness of the indicator function and the growing dimension of the series regressors. We carry out the theoretical analysis by developing a bracketing-based chaining argument in the growing-dimensional time-series setting, which is new to the literature and notably more complicated than the theory presented in Li and Liao (2020) for smooth transformations.[1] In addition, we allow the conditioning variable X_t to be a generated variable (e.g., volatility estimates from GARCH models, or the quantile estimate itself), which is not considered in Li and Liao (2020), either. This latter problem is distinct from the presence of estimation error in f_t because the generated X_t variable enters into a growing number of regressors in the series estimation.

The third contribution of this paper, beyond the theoretical contributions described above, is our empirical analysis of the well-known CoVaR originally proposed by Adrian and Brunnermeier (2016). CoVaR is a measure of the systemic risk of a firm, obtained via quantile regressions of returns on a firm and a market index. We apply the proposed new test to the models presented in Adrian and Brunnermeier (2016) and draw two main conclusions. First, we find that the conditional quantile specification for the market loss is broadly supported by our tests, while the specification for individual firms' losses appears to have room for improvement. Second, we find that the critical covariate in the CoVaR specification is the market volatility measure; the remaining six covariates considered by Adrian and Brunnermeier (2016) do not appear to affect the results of our model specification tests.

Given their widespread use, numerous methods for testing conditional quantile restrictions in dynamic models have been proposed in the literature; see Komunjer (2013) for a summary. Specifically, Christoffersen (1998) noted that if the quantile model is correct, then the indicator variable $1_{\{Y_{t+1} \le f_t\}}$ is i.i.d. Bernoulli with success probability q. He proposes a test of this implication against the alternative that the indicator follows a first-order Markov process. Engle and Manganelli (2004) instead used a linear regression

---

[1]Our analysis is also distinct from prior work in microeconometric settings, because the latter often relies on symmetrization-based empirical process theory that is specific to the i.i.d. setting.

to test whether the indicator variable is predictable using $\mathcal{F}_t$-measurable instruments. Both of these approaches have power against specific parametric alternatives, and can be thought of as testing a fixed number of unconditional moments implied by the conditional moment restriction (1.1). Our nonparametric method complements these parametric approaches by permitting the detection of a priori unknown forms of model misspecification. It is also possible to carry out a nonparametric test using Bierens's test; see, for example, Bierens (1982, 1990), Bierens and Ploberger (1997), Bierens and Ginther (2001), and Escanciano and Velasco (2010). Our test is distinct from the Bierens test because we directly examine the conditional expectation function estimated by series regression, whereas the Bierens test examines the (Donsker-type) empirical process of a continuum of instrumented unconditional moments. These two approaches are generally deemed complementary to each other; see, for example, the discussion in Chernozhukov, Lee, and Rosen (2013).

The rest of the paper is organized as follows. Section 2 describes our specification test and establishes its theoretical properties. Section 3 reports the finite-sample properties of the test via Monte Carlo experiments. Section 4 provides an empirical illustration of the proposed method. Section 5 concludes. The Appendix contains all proofs, with technical lemmas collected in the Online Supplementary Material (Horvath, Li, Liao, Patton (2022)).

2. A consistent specification test for dynamic quantile models

2.1 *The setting and motivating examples*

Consider a univariate time series $(Y_t)_{t\geq 0}$ adapted to a filtration $(\mathcal{F}_t)_{t\geq 0}$. We focus on the conditional q-quantile of Y_{t+1} given $\mathcal{F}_t$-information for some $q \in (0, 1)$, denoted as $Q_q(Y_{t+1}|\mathcal{F}_t)$. Our econometric goal is to nonparametrically test whether a candidate model for $Q_q(Y_{t+1}|\mathcal{F}_t)$ is correctly specified. Arguably the most prominent application of conditional quantile models is estimating VaR and related quantities (e.g., expected shortfall). In that context, Y_t is the loss of an asset portfolio and the conditional quantile $Q_q(Y_{t+1}|\mathcal{F}_t)$ is the one-period-ahead VaR at confidence level q. Below, we focus our discussion on the VaR for ease of exposition, while noting that our theory is applicable for any generic conditional quantile model.

The candidate VaR model involves a process $f_t(\theta)$, where θ is a finite-dimensional parameter taking values in $\mathbb{R}^{d_\theta}$, such that for a given θ the model reports $f_t(\theta)$ as the $\mathcal{F}_t$-conditional q-quantile. The VaR model is completed by the empirical researcher's choice of an estimator $\hat{\theta}_n$, yielding $f_t(\hat{\theta}_n)$ as the estimated VaR. We call the candidate VaR model, summarized by the pair $(f_t(\cdot), \hat{\theta}_n)$ correctly specified if

$$Q_q(Y_{t+1}|\mathcal{F}_t) = f_t(\theta^\star), \quad \text{for all } t \geq 1, \tag{2.1}$$

where $\theta^\star$ is the probability limit of $\hat{\theta}_n$.[2] To anticipate results below, we note that we are agnostic about how $\hat{\theta}_n$ is constructed: for example, it may be computed via quantile regressions, GMM, maximum likelihood, factor models, or even Bayesian methods. Instead, we only assume that $\hat{\theta}_n$ approaches the limit $\theta^\star$ in large samples with $n^{1/2}$-rate of convergence. In our asymptotic analysis, this formalizes the notion that the estimation of this finite-dimensional parameter is "relatively easy" in that it converges at the parametric rate. The following two examples further clarify the empirical context.

EXAMPLE 1 (GARCH VaR). VaR estimation is often based on volatility models such as GARCH. For example, in a Gaussian GARCH(1,1) model, asset return Y_t and its volatility v_t follow

$$Y_{t+1} = v_{t+1}\epsilon_{t+1}, \qquad v_{t+1}^2 = \omega + \beta v_t^2 + \gamma Y_t^2, \qquad \epsilon_t \sim \text{i.i.d.}\, \mathcal{N}(0,1).$$

We collect the model parameters by setting $\theta = (\omega, \beta, \gamma)$. Given θ, the volatility can be computed subsequently, which we denote by $v_{t+1}(\theta)$. The conditional q-quantile is given by $f_t(\theta) = z_q v_{t+1}(\theta)$, where z_q is the q-quantile of the standard normal distribution.

EXAMPLE 2 (CoVaR). Adrian and Brunnermeier (2016) proposed a CoVaR model to measure systematic risk of financial institutions. The key component is a linear conditional quantile model of the market portfolio's loss given the loss of a financial firm and other macroeconomic states, which is given by

$$Q_q\big(Y_{t+1}^{\text{market}}|C_t, Y_{t+1}^{\text{firm}}\big) = \alpha + \beta^\top C_t + \gamma Y_{t+1}^{\text{firm}},$$

where Y_{t+1}^{market} and Y_{t+1}^{firm} are the losses of the market portfolio and the firm, respectively, and C_t collects risk-relevant covariates such as the TED spread and market volatility. It is useful to note that, in this example, the conditioning information set $\mathcal{F}_t$ contains not only the predetermined macroeconomic states (i.e., C_t) but also the contemporaneous firm asset loss (i.e., Y_{t+1}^{firm}), although the information set $\mathcal{F}_t$ is indexed by t under our notational convention.[3]

2.2 *Testing VaR implied conditional moment restrictions*

The conditional quantile restriction (2.1) can be equivalently written as a conditional moment restriction as follows:

$$\mathbb{E}\big[Z_{t+1}^\star|\mathcal{F}_t\big] = 0, \quad \text{where}\, Z_{t+1}^\star \equiv 1_{\{Y_{t+1} \leq f_t(\theta^\star)\}} - q. \tag{2.2}$$

---

[2]Note that the VaR component of a given model can be correctly specified in the sense of equation (2.1) even if the complete model is misspecified. For example, a researcher may construct a fully parametric model that misspecifies some part of the joint distribution of the data, but is correct for the conditional quantile. In this scenario, if $\hat{\theta}_n$ is the maximum likelihood estimator, then $\theta^\star$ is the corresponding pseudo-true parameter from quasi-maximum likelihood estimation (see Komunjer (2005)).

[3]This convention is adopted to emphasize the fact that the dependent variable, say Y_{t+1}^{market}, can enter $\mathcal{F}_t$ only through its lagged values, that is, Y_s^{market} for $s \leq t$. Otherwise, if $\mathcal{F}_t$ were also spanned by Y_{t+1}^{market}, the conditional quantile model would be degenerate.

To make further progress, we consider a finite-dimensional $\mathcal{F}_t$-adapted "instrument" process $X_t^\star$ taking values in a compact set $\mathcal{X}$. Equation (2.2) then implies

$$h(x) = 0, \quad \text{for all } x \in \mathcal{X}, \tag{2.3}$$

where $h(x) = \mathbb{E}[Z_{t+1}^\star | X_t^\star = x]$ denotes the conditional expectation function of $Z_{t+1}^\star$ given $X_t^\star$. We note that $Z_{t+1}^\star$ is not directly observable because it depends on the pseudo-true parameter $\theta^\star$. As suggested by our notation, we also allow $X_t^\star$ to (possibly) depend on $\theta^\star$.

We propose a nonparametric test that takes (2.3) as the null hypothesis. Note that (2.3) is generally an implication of (2.2), because the σ-field spanned by $X_t^\star$ is a subset of $\mathcal{F}_t$. But there is an interesting exception: if the dynamic model is Markovian with state variable $X_t^\star$, as is often assumed in economic models, then these two conditions coincide.

Looking at the condition (2.3), we see clearly that the testing problem is functional in nature, because it concerns the global, instead of local, behavior of the $h(\cdot)$ function. In other words, the inference must be *uniform* across all $x \in \mathcal{X}$. The difficulty of doing so stems from the fact that it is a non-Donsker problem (for which the conventional weak-convergence-based inference is not applicable). In a recent paper, Li and Liao (2020) developed a Yurinskii-type strong approximation to address this issue in a general time series context.[4] Under their approach, we can nonparametrically regress $Z_{t+1}^\star$ on the conditioning variable $X_t^\star$ using the series method, and then invoke the strong approximation theory to show that the nonparametric estimation error function can be approximated, or "coupled," by a diverging sequence of Gaussian processes. The Gaussian approximation then permits feasible uniform inference.

The problem of evaluating conditional quantile models leads to two complications that are not considered in Li and Liao (2020). First, the variables $(Z_{t+1}^\star, X_t^\star)$ depend on $\theta^\star$, which needs to be replaced by $\hat{\theta}_n$ in a feasible procedure. The technical challenge here is that $Z_{t+1}^\star$ depends on $\theta^\star$ in a nonsmooth way because of the presence of the indicator function; this issue is further complicated by the fact that the nonparametric series regression involves a growing number of series approximation functions.

Second, the finite-sample performance of the strong Gaussian coupling is found to be poor for quantiles in the tails (details are presented in the next section). VaR and related quantities, like CoVaR, invariably focus on quantiles with $q \geq 0.95$, and so an alternative inference procedure is needed. We propose a novel bootstrap procedure to compute critical values, and justify its theoretical validity in the current nonstandard (non-Donsker) context for uniform functional inference.

2.3 *The test and its asymptotic properties*

We now provide details on our test and prove its asymptotic validity, namely that it controls size under the null hypothesis and is consistent against fixed alternatives. We first

---

[4]Yurinskii coupling has also been used by Chernozhukov, Lee, and Rosen (2013) and Belloni et al. (2015) for constructing uniform inference in a microeconometric context with i.i.d. data.

define the test statistic. To capture more explicitly how the $(Z_{t+1}^{\star}, X_t^{\star})$ variables are related to the pseudo-true parameter $\theta^{\star}$, we introduce two functions, $Z_{t+1}(\theta)$ and $X_t(\theta)$, such that

$$Z_{t+1}^{\star} = Z_{t+1}(\theta^{\star}), \qquad X_t^{\star} = X_t(\theta^{\star}). \tag{2.4}$$

Note that $Z_{t+1}(\theta) = 1_{\{Y_{t+1} \leq f_t(\theta)\}} - q$, while the form of $X_t(\theta)$ depends on the specific application. As mentioned above, we assume that $\hat{\theta}_n$ is a $n^{1/2}$-consistent estimator for $\theta^{\star}$.[5] The feasible analogues of the quantities in (2.4) are then given by

$$\widehat{Z}_{t+1} \equiv Z_{t+1}(\hat{\theta}_n), \qquad \widehat{X}_t \equiv X_t(\hat{\theta}_n).$$

We nonparametrically regress $\widehat{Z}_{t+1}$ on $\widehat{X}_t$ using the series method. To do so, we consider m_n approximating basis functions $P(x) = (p_1(x), \ldots, p_{m_n}(x))^{\top}$. By convention, we assume that the constant term is always included by setting $p_1(\cdot) = 1$. The nonparametric interpretation of the series estimation relies on taking $m_n \to \infty$ so that the conditional expectation function $h(\cdot)$ can be approximated sufficiently well by a linear combination of the approximating functions. Commonly-used basis functions include polynomials, Fourier series, splines, and wavelets; see Chen (2007) for additional details.

We conduct least-squares regression of $\widehat{Z}_{t+1}$ on $P(\widehat{X}_t)$ and obtain the regression coefficient as

$$\widehat{b}_n \equiv \left(\sum_{t=1}^{n} P(\widehat{X}_t)P(\widehat{X}_t)^{\top}\right)^{-1} \left(\sum_{t=1}^{n} P(\widehat{X}_t)\widehat{Z}_{t+1}\right).$$

The estimator of the conditional expectation function $h(\cdot)$ is $\widehat{h}_n(\cdot) \equiv P(\cdot)^{\top}\widehat{b}_n$. The associated standard error function is estimated by

$$\widehat{\sigma}_n(\cdot) \equiv \left(P(\cdot)^{\top}\widehat{\Sigma}_n P(\cdot)\right)^{1/2},$$

where, with $\widehat{u}_t \equiv \widehat{Z}_{t+1} - \widehat{h}_n(\widehat{X}_t)$, we set

$$\widehat{Q}_n \equiv n^{-1}\sum_{t=1}^{n} P(\widehat{X}_t)P(\widehat{X}_t)^{\top}, \qquad \widehat{A}_n \equiv n^{-1}\sum_{t=1}^{n} \widehat{u}_t^2 P(\widehat{X}_t)P(\widehat{X}_t)^{\top}, \qquad \widehat{\Sigma}_n \equiv \widehat{Q}_n^{-1}\widehat{A}_n\widehat{Q}_n^{-1}.$$

Finally, we define the "sup-t" test statistic, $\widehat{T}_n$, as

$$\widehat{T}_n \equiv \sup_{x \in \mathcal{X}} \frac{n^{1/2}|\widehat{h}_n(x)|}{\widehat{\sigma}_n(x)}.$$

---

[5]For in-sample specification tests, the $\hat{\theta}_n$ estimator is typically estimated using the full sample with size n. This framework also accommodates (pseudo) "out-of-sample"test when $\hat{\theta}_n$ is estimated using a fixed window provided that the size of the estimation sample is nondegenerate relative to n. If the estimation is performed under a recursive or rolling scheme, we would have a more complicated situation with a sequence of estimators, say $(\hat{\theta}_{n,t})_{1 \leq t \leq n}$. In that case, we need to strengthen the rate requirement to a uniform version, that is, $\max_{1 \leq t \leq n} \|\hat{\theta}_{n,t} - \theta^{\star}\| = O_p(n^{-1/2})$. This condition may be justified using a functional central limit theorem for the estimator process $\hat{\theta}_{n,\lfloor n \cdot \rfloor}$ for the recursive setting, as well as the rolling setting provided that the rolling window size is nondegenerate with respect to n. If the rolling window is of fixed size as considered in Giacomini and White (2006), one may follow that prior work and treat the estimates for θ as an observed sequence, and the issue of estimation error becomes irrelevant for the asymptotic inference.

The first part of our analysis is to establish a strong approximation for the sup-t statistic $\widehat{T}_n$ using the supremum of a Gaussian process. In the infeasible case with $\theta^\star$ known, such an approximation could be obtained by directly using the theory of Li and Liao (2020). The key complication here is to analyze the effect of replacing $\theta^\star$ with the estimator $\hat{\theta}_n$. We aim to provide sufficient conditions under which such an effect is asymptotically negligible. The theoretical message is easy to understand. Intuitively, the estimation error of the finite-dimensional estimator $\hat{\theta}_n$ is relatively small (as it converges at the parametric $n^{1/2}$-rate) compared with the statistical noise in the second-stage nonparametric inference with slower rate of convergence. This theory justifies a least-squares procedure that is very easy to implement, and sets a useful benchmark for further refinement in future work.

The technical formalization of this simple intuition turns out to be nontrivial. The main source of complication is that $Z_{t+1}^\star$ depends on $\theta^\star$ in a nonsmooth manner. In prior literature on M-estimation and GMM, the standard approach for addressing the lack of smoothness relies on the concept of stochastic equicontinuity from empirical process theory. Under stochastic equicontinuity, the nonsmooth sample moment function can be effectively replaced with its limiting version, which is typically smooth (i.e., twice differentiable); see, for example, Andrews (1994) and van der Vaart (1998).

This "off-the-shelf" empirical-process approach, however, is insufficient for our analysis. The reason is that our nonparametric series estimation involves a growing number of regressors (i.e., $m_n \to \infty$). Consequently, the behavior of the regression coefficient $\widehat{b}_n$ relies on a growing number of moments that are nonsmooth in the θ parameter. While the usual stochastic equicontinuity argument can be used to deal with a fixed number of nonsmooth sample moments, it does not guarantee a uniform approximation when the dimension grows to infinity, which is exactly the technical challenge here. To address this issue, we start from first principles and use a bracketing-based chaining argument to characterize the local modulus of continuity of a growing-dimensional empirical process for time series data.[6] This analysis appears to be new to the literature, and should be generally useful in future work involving nonsmoothness and growing dimensions in time series econometrics.

We now introduce the regularity conditions for our asymptotic theory. Since $\hat{\theta}_n$ is assumed to be a $n^{1/2}$-consistent estimator of $\theta^\star$, our analysis regarding the plug-in effect concentrates on $n^{-1/2}$-neighborhoods of $\theta^\star$ of the form

$$B_n(R) \equiv \left\{ \theta \in \Theta : n^{1/2} \| \theta - \theta^\star \| \leq R \right\}, \quad \text{for } R > 0.$$

That is, $B_n(R)$ is a closed ball centered at $\theta^\star$ with radius $Rn^{-1/2}$. We denote by $F_{t+1|t}(\cdot)$ the $\mathcal{F}_t$-conditional distribution function of Y_{t+1}, and write the $\mathcal{F}_t$-conditional expectation of $Z_{t+1}(\theta)$ as

$$\bar{Z}_{t+1}(\theta) \equiv \mathbb{E}[1_{\{Y_{t+1} \leq f_t(\theta)\}} - q | \mathcal{F}_t] = F_{t+1|t}\big(f_t(\theta)\big) - q.$$

---

[6]In microeconometric settings with i.i.d. data, empirical-process arguments are mostly based on symmetrization. Symmetrization relies crucially on the independence assumption, and is not applicable in our time-series setting.

In addition, let $\partial_\theta \bar{Z}_{t+1}(\cdot)$ denote the gradient of $\bar{Z}_{t+1}(\cdot)$ with respect to θ, and further set

$$g(x) \equiv \mathbb{E}\big[\partial_\theta \bar{Z}_{t+1}(\theta^\star)|X_t^\star = x\big], \qquad \eta_t \equiv \partial_\theta \bar{Z}_{t+1}(\theta^\star) - g(X_t^\star).$$

ASSUMPTION 1. (i) $\hat{\theta}_n - \theta^\star = O_p(n^{-1/2})$; (ii) there exist some $\mathcal{F}_t$-measurable variables $(L_t, L_{X,t})$ such that, for all $y_1, y_2 \in \mathbb{R}$,

$$\big|F_{t+1|t}(y_1) - F_{t+1|t}(y_2)\big| \leq L_t|y_1 - y_2|,$$

and, for any $R > 0$ and all $\theta_1, \theta_2 \in B_n(R)$,

$$\big|f_t(\theta_1) - f_t(\theta_2)\big| \leq L_t\|\theta_1 - \theta_2\|, \qquad \big\|X_t(\theta_1) - X_t(\theta_2)\big\| \leq L_{X,t}\|\theta_1 - \theta_2\|.$$

(iii) $\bar{Z}_t(\cdot)$ is continuously differentiable at $\theta^\star$ and, for some $\mathcal{F}_t$-measurable variable $\bar{L}_t$,

$$\big|\bar{Z}_t(\theta) - \bar{Z}_t(\theta^\star) - \partial_\theta \bar{Z}_t(\theta^\star)^\top (\theta - \theta^\star)\big| \leq \bar{L}_t\|\theta - \theta^\star\|^2$$

for any $\theta \in B_n(R)$; (iv) L_t, $L_{X,t}$, and $L_t L_{X,t}$ are L_p-bounded for some $p > 2d_\theta$, and $\partial_\theta \bar{Z}_t(\theta^\star)$ and $\bar{L}_t$ are L_2-bounded; (v) $g(\cdot)$ is a continuously differentiable function and, for each $j \in \{1, \ldots, d_\theta\}$, there exists $\gamma_{j,n} \in \mathbb{R}^{m_n}$ such that

$$\sup_{x \in \mathcal{X}}\big|g_j(x) - P(x)^\top \gamma_{j,n}\big| = o\big((\log n)^{-1/2}\big);$$

(vi) for each $j \in \{1, \ldots, d_\theta\}$, the largest eigenvalue of the matrix $\mathrm{Var}(n^{-1/2}\sum_{t=1}^n P(X_t^\star)\eta_{j,t})$ is bounded; (vii) $\sup_{x \in \mathcal{X}}\|P(x)\|^{-1} = o((\log n)^{-1/2})$.

　　Assumption 1 collects regularity conditions that we use to control the effect of plugging in the $\hat{\theta}_n$ estimator. We make a few remarks on the roles that these conditions play in our theory. Condition (i) imposes, as mentioned above, the $n^{1/2}$-consistency of $\hat{\theta}_n$; this is quite natural for parametric (e.g., the maximum likelihood) or semiparametric (e.g., least-squares, quantile regression, GMM, etc.) estimators typically used in practice. Conditions (ii)–(iv) impose smoothness conditions on $F_{t+1|t}(\cdot)$, $f_t(\cdot)$, and $X_t(\cdot)$ over local neighborhoods around $\theta^\star$. Based on an empirical-process type argument, we can exploit the smoothness in $F_{t+1|t}(\cdot)$ and $\bar{Z}_{t+1}(\theta)$ to bypass the issue of nonsmoothness in $Z_{t+1}(\cdot)$. Also note that in some applications the instrument $X_t^\star$ does not depend on $\theta^\star$, so we can take $L_{X,t} = 0$. Conditions (v) and (vi) impose additional regularities on $\partial_\theta \bar{Z}_{t+1}(\theta^\star)$, which in turn measures the sensitivity of the second-stage regression with respect to the estimation error in $\hat{\theta}_n$. These conditions are relatively mild. Finally, condition (vii) states that the growing-dimensional vector $\|P(x)\|$ diverges no slower than $(\log n)^{1/2}$, which partly reflects the slower (nonparametric) rate of convergence of the series estimator.

　　We also need regularity conditions for making uniform series inference in the (simpler) infeasible setting, which requires some additional notation. The regression residual in the infeasible case is given by

$$u_t^\star \equiv Z_{t+1}^\star - \mathbb{E}\big[Z_{t+1}^\star|X_t^\star\big] = Z_{t+1}^\star - h(X_t^\star).$$

We then set

$$Q_n \equiv n^{-1} \sum_{t=1}^{n} \mathbb{E}[P(X_t^\star)P(X_t^\star)^\top], \qquad \widehat{Q}_n^\star \equiv n^{-1} \sum_{t=1}^{n} P(X_t^\star)P(X_t^\star)^\top,$$

$$A_n \equiv n^{-1} \sum_{t} \mathbb{E}[(u_t^\star)^2 P(X_t^\star)P(X_t^\star)^\top], \qquad \widehat{A}_n^\star \equiv n^{-1} \sum_{t=1}^{n} (u_t^\star)^2 P(X_t^\star)P(X_t^\star)^\top.$$

Note that $\widehat{Q}_n^\star$ and $\widehat{A}_n^\star$ are the infeasible sample analogues of Q_n and A_n, respectively. Below, for any $\varepsilon > 0$, we denote $\mathcal{X} \oplus \varepsilon = \{x + u : x \in \mathcal{X}, \|u\| \leq \varepsilon\}$, that is, an ε-enlargement of $\mathcal{X}$.[7] We use $\|\cdot\|_S$ to denote the matrix spectral norm.

ASSUMPTION 2. *We have* (i) *the function* $h(x) \equiv \mathbb{E}[Z_{t+1}^\star | X_t^\star = x]$ *is continuously differentiable;* (ii) *there exist a sequence* $b_n^\star$ *of* m_n-*dimensional vectors and real sequences* $\zeta_{0,n}$ *and* $\zeta_{1,n}$, *such that* (*recalling the constant* $p > 2$ *from Assumption* 1)

$$\begin{cases} \sup_{x \in \mathcal{X} \oplus \varepsilon_n} |h(x) - P(x)^\top b_n^\star| = O(n^{-1/2}), \\ \max_{1 \leq l \leq m_n} \sup_{x \in \mathcal{X} \oplus \varepsilon_n} |\partial^j p_l(x)| \leq \zeta_{j,n}, \quad j \in \{0, 1\}, \end{cases} \quad \text{for all } \varepsilon_n \asymp n^{1/p - 1/2},$$

and $\zeta_{0,n} m_n n^{-1/4} + \zeta_{0,n}\zeta_{1,n} m_n n^{-1/2} = o(1/\log n)$; (iii) *the eigenvalues of* Q_n *and* A_n *are bounded from above and away from zero and* $\overline{A}_n = \mathrm{Var}(n^{-1/2} \sum_{t=1}^{n} P(X_t^\star) u_t^\star)$ *has bounded eigenvalues;* (iv) $\|\widehat{Q}_n^\star - Q_n\|_S = O_p(\delta_{Q,n})$, $\|\widehat{A}_n^\star - A_n\|_S = O_p(\delta_{A,n})$, *and* $m_n^{1/2}(\delta_{Q,n} + \delta_{A,n}) = o_p(1/\log n)$; (v) $\log(\zeta_n^L) = O(\log n)$, *where* $\zeta_n^L \equiv \sup_{x_1, x_2 \in \mathcal{X}} \|P(x_1) - P(x_2)\| / \|x_1 - x_2\|$.

The conditions in Assumption 2 are fairly standard for analyzing the series estimation. A few remarks are in order. Condition (ii) introduces the population analogue $b_n^\star$ for the regression coefficient $\widehat{b}_n$. The existence of $b_n^\star$ is ensured by well-known approximation theory (see, e.g., Chen (2007) and many references therein), and the precision of the approximation may be stated explicitly in terms of the smoothness of the $h(\cdot)$ function and the dimensionality of $X_t^\star$. We note that the uniform bound conditions are stated over ε_n-enlargements of $\mathcal{X}$, which is slightly stronger than a more conventional condition stated over $\mathcal{X}$. This modification is needed here because the generated variables $\widehat{X}_t$ may take values outside $\mathcal{X}$ (but they still fall in $\mathcal{X} \oplus \varepsilon_n$ with high probability under maintained assumptions). If the instrument $X_t^\star$ does not depend on $\theta^\star$, or the generated variable $\widehat{X}_t$ is restricted to take values in $\mathcal{X}$, we no longer need the enlargement, which amounts to setting $\varepsilon_n = 0$. The only high-level requirement is condition (iv), which imposes rates of convergence for the infeasible estimators $\widehat{Q}_n^\star$ and $\widehat{A}_n^\star$ under matrix spectral norm. These conditions can be verified using, for example, Lemma 2.1 of Chen and Christensen (2015).

---

[7]We adopt the direct sum notation because $\mathcal{X} \oplus \varepsilon$ is in fact the direct sum of $\mathcal{X}$ and a closed ball centered at zero with radius ε.

ASSUMPTION 3. *Under the null hypothesis, there exists a sequence ξ_n of m_n-dimensional standard Gaussian random vectors such that*

$$
\sup_{x \in \mathcal{X}} \left| \frac{n^{-1/2} P(x)^\top Q_n^{-1} \sum_{t=1}^{n} P(X_t^\star) Z_{t+1}^\star}{\sigma_n(x)} \right| = \sup_{x \in \mathcal{X}} \left| \frac{P(x)^\top \Sigma_n^{1/2} \xi_n}{\sigma_n(x)} \right| + o_p\big((\log n)^{-1/2}\big), \quad (2.5)
$$

where $\Sigma_n \equiv Q_n^{-1} A_n Q_n^{-1}$ and $\sigma_n(x) \equiv (P(x)^\top \Sigma_n P(x))^{1/2}$.

Assumption 3 is the key element for uniform inference and is clearly high-level in nature. This condition essentially states that the infeasible sup-t statistic on the left-hand side of (2.5) can be strongly approximated by the supremum of the centered Gaussian process on the right-hand side of that equation. For our test, we need only that this condition holds under the null hypothesis. This high-level condition may be verified as a special case of the strong approximation theory developed by Li and Liao (2020) for general time-series data. More specifically, we observe that under Assumption 2(iii), Assumption 3 holds if

$$
\left\| \sum_{t=1}^{n} n^{-1/2} P(X_t^\star) Z_{t+1}^\star - A_n^{1/2} \xi_n \right\| = o_p\big((\log n)^{-1/2}\big). \quad (2.6)
$$

Since $n^{-1/2} P(X_t^\star) Z_{t+1}^\star$ forms a martingale difference array under the null hypothesis, we can use the Yurinskii-type coupling result in Li and Liao (2020) (see their Theorem 1) to verify (2.6) under the primitive conditions provided in that paper.[8]

We are now ready to state the asymptotic property of the sup-t statistic $\widehat{T}_n$ under the null hypothesis.

THEOREM 1. *Suppose that Assumptions 1, 2, and 3 hold. Then under the null hypothesis there exists a sequence ξ_n of m_n-dimensional standard normal random variables such that*

$$
\widehat{T}_n - \widetilde{T}_n = o_p\big((\log n)^{-1/2}\big),
$$

where

$$
\widetilde{T}_n = \sup_{x \in \mathcal{X}} \frac{\big| P(x)^\top \Sigma_n^{1/2} \xi_n \big|}{\sigma_n(x)}.
$$

Theorem 1 shows that the sup-t statistic $\widehat{T}_n$ can be strongly approximated by the supremum of a Gaussian process $P(\cdot)^\top \Sigma_n^{1/2} \xi_n / \sigma_n(\cdot)$, for which the sampling variability

---

[8]See Assumption 1 in Li and Liao (2020), as well as Assumption B2 and Lemma B1 in the Appendix of the Online Supplemental Material of that paper for additional technical discussions.

is fully captured by the Gaussian vector ξ_n with growing dimension.[9] Based on this result, a natural way of computing the critical value for the sup-t statistic is to estimate the quantile of the approximating variable $\widetilde{T}_n$ by simulating the Gaussian process, with Σ_n and $\sigma_n(\cdot)$ replaced by their estimators $\widehat{\Sigma}_n$ and $\widehat{\sigma}_n(\cdot)$, respectively. This approach is shown to be asymptotically valid in Li and Liao (2020). However, as we will show in our Monte Carlo experiments in the next section, this approach suffers from nontrivial size distortion at relatively high quantiles (e.g., $q \geq 0.95$), which makes it effectively inapplicable for VaR applications. We thus propose an alternative approach using bootstrap, described in Algorithm 1 below.

ALGORITHM 1 (Bootstrap critical value).

Step 1. *Resample* $(\widehat{Z}^*_{t+1}, \widehat{X}^*_t)_{1 \leq t \leq n}$ *as an i.i.d. sample with replacement from* $(\widehat{Z}_{t+1}, \widehat{X}_t)_{1 \leq t \leq n}$.

Step 2. *Compute* $(\widehat{h}^*_n(\cdot), \widehat{\sigma}^*_n(\cdot))$ *in the same way as* $(\widehat{h}_n(\cdot), \widehat{\sigma}_n(\cdot))$ *but with* $(\widehat{Z}_{t+1}, \widehat{X}_t)_{1 \leq t \leq n}$ *replaced by* $(\widehat{Z}^*_{t+1}, \widehat{X}^*_t)_{1 \leq t \leq n}$, *and then set* $\widehat{T}^*_n = \sup_{x \in \mathcal{X}} n^{1/2} |\widehat{h}^*_n(x) - \widehat{h}_n(x)| / \widehat{\sigma}^*_n(x)$.

Step 3. *Repeat steps* 1–2 *for a large number of times. At significance level* α, *set the critical value* $cv_{n,\alpha}$ *as the* $1 - \alpha$ *quantile of* $\widehat{T}^*_n$ *in the Monte Carlo sample. Reject the null hypothesis (i.e.,* $h(x) = 0$ *for all* $x \in \mathcal{X}$) *if* $\widehat{T}_n > cv_{n,\alpha}$.

Algorithm 1 resembles a "textbook" i.i.d. bootstrap. To compute the critical value, one performs i.i.d. resampling and then repeatedly computes the test statistic. It is useful to note that the $\hat{\theta}_n$ estimator does not need to be recomputed for the bootstrap samples because its plug-in error is asymptotically negligible for the nonparametric test. Since the sup-t statistic is not asymptotically pivotal, we do not expect the bootstrap to deliver a formal theoretical refinement. Instead, we only advocate the bootstrap as a practical and theoretically justified way to conduct feasible inference, which turns out to outperform the asymptotic Gaussian-based method in the applications of interest in this paper.

The validity of the i.i.d. bootstrap in the time series setting of this paper follows from the fact that, under the null hypothesis, the sampling variability in the test statistic is driven by a martingale difference sequence (namely $P(X^\star_t)u^\star_t$). In this case the i.i.d. bootstrap is sufficient to approximate its finite-sample distribution under the null. To construct a uniform confidence band for $h(\cdot)$ under the alternative, one would have to capture autocovariances by using, for example, a block bootstrap, but in the testing context of this paper this is not necessary, because to prove the test's consistency it suffices to simply control the asymptotic magnitude of the critical value.

Theorem 2 establishes the asymptotic property of the bootstrapped test statistic, and further shows the size and power properties of the resulting test.

---

[9]The strong approximation error $\widehat{T}_n - \widetilde{T}_n$ is shown to be $o_p((\log n)^{-1/2})$, which is slightly stronger than a "usual" $o_p(1)$ statement. This stronger statement is needed in the present setting because the probability density of the sup-Gaussian approximating variable $\widetilde{T}_n$ is divergent, but it can be bounded at rate $(\log(m_n))^{1/2} \leq (\log n)^{1/2}$. As a result, the $o_p((\log n)^{-1/2})$ approximation error only leads to an $o(1)$ error in coverage probability, which is needed in our size analysis.

Theorem 2. *Suppose that Assumptions 1, 2, and 3 hold. Then (a) under the null hypothesis there exists a sequence ξ_n^* of m_n-dimensional standard normal random variables such that*

$$\widehat{T}_n^* - \widetilde{T}_n^* = o_p\big((\log n)^{-1/2}\big),$$

where

$$\widetilde{T}_n^* \equiv \sup_{x \in \mathcal{X}} \frac{\big|P(x)^\top \Sigma_n^{1/2} \xi_n^*\big|}{\sigma_n(x)}.$$

(b) for $\alpha \in (0, 1/2)$, the test described in Algorithm 1 has asymptotic level α under the null hypothesis (i.e., $h(x) = 0$ for all $x \in \mathcal{X}$) and has asymptotic power 1 under the alternative hypothesis (i.e., $h(x) \neq 0$ for some $x \in \mathcal{X}$).

3. Simulations

We now examine the finite-sample properties of the proposed test in a Monte Carlo experiment for GARCH-based VaR models. We consider two data generating processes (DGPs), each with sample size $n = 2000$. Under the first DGP, we generate a time series $(Y_t)_{1 \leq t \leq n}$ of daily losses from a Gaussian GARCH(1,1) process:

$$\text{DGP-N:} \quad \begin{cases} Y_t = v_t \varepsilon_t, & \varepsilon_t \sim \text{i.i.d.} \, \mathcal{N}(0, 1), \\ v_t^2 = \omega + \beta v_{t-1}^2 + \gamma Y_{t-1}^2, \end{cases}$$

with parameters $\omega = 0.05$, $\beta = 0.9$, and $\gamma = 0.05$. Our second DGP is taken from Bontemps (2019), who uses a Student's t EGARCH process:

$$\text{DGP-A:} \quad \begin{cases} Y_t = v_t z_t, & z_t \sim \text{i.i.d.} \, t(0, 1, 4), \\ \log\big(v_t^2\big) = \omega + \gamma\big(|z_{t-1}| - \mathbb{E}\big[|z_{t-1}|\big]\big) + \delta z_{t-1} + \beta \log\big(v_{t-1}^2\big), \end{cases}$$

with parameters $\omega = 0.0001$, $\gamma = 0.3$, $\delta = -0.8$, and $\beta = 0.9$.

In both cases, the quantile model is a Gaussian GARCH(1,1), estimated via (quasi) maximum likelihood. The VaR of Y_{t+1} at confidence level q is obtained as

$$f_t(\hat{\theta}_n) = \Phi^{-1}(q)\sqrt{\hat{\omega}_n + \hat{\beta}_n \hat{v}_t^2 + \hat{\gamma}_n Y_t^2},$$

where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal distribution. Under DGP-N, this model is correct and the null hypothesis is true. Under DGP-A this model is misspecified and the null hypothesis is false. We can thus examine the test's size and power properties using these two DGPs. Below, we consider VaRs for $q \in \{0.75, 0.9, 0.95, 0.99\}$ and fix the significance level of the test at $\alpha = 5\%$.

We use $(Y_t, \hat{v}_t)$ as the (feasible) conditioning variable in the nonparametric regression. To construct the basis functions, we first use a rank-transformation to rescale Y_t and $\hat{v}_t$ onto the $[-1, 1]$ interval, and denote the transformed variables as Y_t' and $\hat{v}_t'$, re-

Table 1. Finite-sample rejection rates for GARCH-VaR models.

	Null Hypothesis				Alternative Hypothesis			
	75%	90%	95%	99%	75%	90%	95%	99%
Gaussian	0.047	0.056	0.082	0.336	1.000	1.000	0.999	0.995
Bootstrap	0.038	0.032	0.031	0.034	0.980	0.974	0.969	0.728

Note: This table reports the rejection frequencies of the specification tests at the 5% significance level for the GARCH-based VaR model with probability levels ranging from 75% to 99%. The critical values are computed either based on the asymptotic Gaussian approximation or the bootstrap. The left and right panels are for the null hypothesis (i.e., DGP-N) and the alternative hypothesis (i.e., DGP-A), respectively.

spectively; we then set $\widehat{X}_t = (Y_t', \hat{v}_t')$.[10] We use $m_n = 6$ series terms, with the form $P(\widehat{X}_t) = (1, Y_t', \hat{v}_t', Y_t'\hat{v}_t', LP_2(Y_t'), LP_2(\hat{v}_t'))^\top$, where $LP_2(\cdot)$ denotes the second-order Legendre polynomial, which is employed to reduce the multicollinearity among the series terms.

Below, we examine the finite-sample performance of the specification test via 10,000 Monte Carlo trials. The critical value of the test is computed in two ways: the first is based on the asymptotic Gaussian approximation (Theorem 1), and the second is based on the bootstrap procedure (Theorem 2) with 1000 resamples.

Table 1 reports the finite-sample rejection frequencies for the two tests. The left panel reports results under the null hypothesis. When $q = 0.75$, we see that both the asymptotic Gaussian approximation and the bootstrap method control size well. However, as q increases, the Gaussian approximation leads to nontrivial overrejections. For example, the test rejects 33.6% of the time when $q = 0.99$. In contrast, the rejection rates of the bootstrap method are generally close to the 5% nominal level across all settings, although the test appears to be somewhat conservative.[11]

The right panel of Table 1 reports the rejection rates of the tests under the alternative hypothesis. Both tests exhibit nontrivial power to detect the model misspecification. As expected, power is lower at higher quantiles, reflecting the reduction in information available to detect model misspecification as we move deeper into the tail. The rejection rates of the test based on the asymptotic Gaussian approximation are greater than those

---

[10]The use of rank-transformation may be formally justified as follows. Let $F_Y(\cdot)$ and $F_v(\cdot)$ denote the cumulative distribution functions of Y_t and v_t, respectively, and then let $X_t^\star$ collect the transformed variables $2F_Y(Y_t) - 1$ and $2F_v(v_t) - 1$, which take values in $[-1, 1]$. The rank-transformed variables Y_t' and $\hat{v}_t'$ can be written analogously as $2\widehat{F}_Y(Y_t) - 1$ and $2\widehat{F}_v(\hat{v}_t) - 1$, where $\widehat{F}_Y(\cdot)$ and $\widehat{F}_v(\cdot)$ are the sample analogues of $F_Y(\cdot)$ and $F_v(\cdot)$, respectively. Therefore, the difference between $\widehat{X}_t$ and $X_t^\star$ stems from replacing $\theta^\star \equiv (F_Y(\cdot), F_v(\cdot), \omega, \beta, \gamma)$ with its estimator $\hat{\theta}_n$. Note that a Donsker theorem for weakly dependent data implies that cumulative distribution functions can be estimated at the $n^{1/2}$-rate of convergence under the uniform metric (see, e.g., Theorem 1 of Dehling, Durieu, and Volny (2009), Theorem 3.1 of Dedecker, Rio, and Merlevède (2014), and Chapter 7 of Rio (2017)). Hence, although $\hat{\theta}_n$ contains a functional component, it is still a $n^{1/2}$-consistent estimator for $\theta^\star$. The theory developed above can be easily generalized to accommodate this slightly more general "plug-in," by replacing the Euclidean distance on θ with a generic metric.

[11]The size distortion resulted from the Gaussian-based critical value does not arise from the plug-in estimation error in $\hat{\theta}_n$. In simulation results not presented here, we also considered the infeasible setting with known $\theta^\star$, and found that the Gaussian-based test's rejection rate is at 35.1% (resp., 10.9%) when $q = 99\%$ (resp., 95%). More generally, the rejection rates in the infeasible setting are very similar to those in the feasible setting presented here.

based on the bootstrap, suggesting higher power, but given the former's size distortion at high quantiles, that test is not always reliable.

4. Empirical application

In an influential paper, Adrian and Brunnermeier (2016) proposed a measure of a firm's systemic risk known as "CoVaR" based on how the market portfolio's VaR differs according to whether the firm's return is at a normal level or a "stressed" level. To define this measure, let Y_t^{market} and $Y_t^{(i)}$ be the losses of the market and firm i on week t, respectively (measured simply as the negative of their returns that week). For some probability level q, such as 0.95 or 0.99, the firm's q-VaR is obtained by fitting the following linear quantile regression:

$$Q_q\big(Y_{t+1}^{(i)}|C_t\big) = \alpha_q^{(i)} + \beta_q^{(i)\top} C_t, \tag{4.1}$$

where C_t is a collection of covariates described below. The market portfolio's q-VaR is modeled similarly, but with an additional covariate, namely the contemporaneous loss of firm i:

$$Q_q\big(Y_{t+1}^{\text{market}}|C_t, Y_{t+1}^{(i)}\big) = \tilde{\alpha}_q^{(i)} + \tilde{\beta}_q^{(i)\top} C_t + \tilde{\gamma}_q^{(i)} Y_{t+1}^{(i)}. \tag{4.2}$$

To facilitate the definition of CoVaR, it is helpful to define the following function:

$$\psi_q^{(i)}(c, y) \equiv Q_q\big(Y_{t+1}^{\text{market}}|C_t = c, Y_{t+1}^{(i)} = y\big), \tag{4.3}$$

which is the model-implied q-quantile of the market portfolio when the covariates C_t take value c and the loss of firm i equals y. The CoVaR of firm i is then defined as

$$\text{CoVaR}_q^{(i)} \equiv \psi_q^{(i)}\big(C_t, Q_q\big(Y_{t+1}^{(i)}|C_t\big)\big) - \psi_q^{(i)}\big(C_t, Q_{0.5}\big(Y_{t+1}^{(i)}|C_t\big)\big). \tag{4.4}$$

In words, $\text{CoVaR}_q^{(i)}$ measures the change in the VaR of the market portfolio when firm i's loss moves (hypothetically) from its conditional median to its conditional q-quantile. If the market's VaR changes markedly when the the loss of firm i moves to its q-quantile, then the market VaR is sensitive to the losses of firm i and that firm is said to have high systemic risk. If the market's VaR is insensitive to the losses of firm i, then that firm has little impact on the market and is said to have low systematic risk.

The estimation of CoVaR thus relies on two building blocks: (4.1) exclusively pertains to the conditional quantile of firm i's loss, while (4.2) captures the relationship between the market's VaR and the firm. For simplicity, we refer to (4.1) and (4.2) as the VaR and CoVaR specification, respectively, although the CoVaR risk measure is computed using equation (4.4). Since both components of the CoVaR measure are conditional quantile models, we can apply the proposed test to examine the empirical specifications proposed by Adrian and Brunnermeier (2016).

Our empirical analysis uses the same data as Adrian and Brunnermeier (2016).[12] The dataset contains all publicly traded US commercial banks, broker-dealers, insurance companies, and real estate companies for the period from January 1971 to June

---

[12]This data is available from the American Economic Review website: https://www.aeaweb.org/articles?id=10.1257/aer.20120555.

2013, a total of 1823 firms and 2209 weeks. The covariates C_t are the weekly real estate sector return (Housing), the weekly market return of the S&P 500 index (MktRet), short-term TED spread (TED), change in the credit spread (Credit), change in 3-month yield (Yld3M), change in the slope of the yield curve (TERM), and equity volatility (MktSD). We refer to Adrian and Brunnermeier (2016) for a more detailed description of these variables.

To conduct the model specification test, we need to select a conditioning variable X_t known at time t. Motivated by the recent literature on the impact of economic uncertainty, we use a variety of uncertainty measures including the economic policy uncertainty index proposed by Baker, Bloom, and Davis (2016) and the financial and macro uncertainty indexes proposed by Jurado, Ludvigson, and Ng (2015).[13] These uncertainty measures are moderately correlated: the correlation between economic policy uncertainty and financial (resp., macro) uncertainty is 0.32 (resp., 0.20), and the correlation between financial and macro uncertainty indexes is 0.59. Below, we use each of these indexes separately as a conditioning variable in our specification test. As in our simulation study, we use the rank-transformation to rescale X_t to have support $[-1, 1]$ and set the basis functions as mth-order Legendre polynomials (resulting in $m + 1$ series terms).[14] Critical values are computed from 1000 bootstrap replications using Algorithm 1 from Section 2, at significance level $\alpha = 5\%$.

Table 2 presents the results of the bootstrap-based specification tests for $q = 95\%$ or 99%, for the three different economic uncertainty measures. To gauge the sensitivity of the test outcome to the choice of the number of series terms in the nonparametric estimation, we present results for $m \in \{6, \dots, 10\}$. We conduct specification tests for each firm separately, and summarize the results by reporting the rejection rates averaged over the cross-section.

From Table 2, we see that the 95%-VaR specification (equation (4.1)) is rejected substantially more frequently than the 5% nominal level: the rejection rate ranges from 15.7% to 20.6% across different implementations. On the other hand, the 95%-CoVaR specification (equation (4.2)) is only rejected approximately at the nominal level, suggesting that this specification is satisfactory for a representative firm. At the 99%-quantile, the rejection rates for both VaR and CoVaR models are lower, and are generally close to or below the nominal level. It is noteworthy that the test results are broadly insensitive to the choice of series terms and measure for economic uncertainty.

Given that the CoVaR specification is rarely rejected in the test results reported in Table 2, is it possible to replace it with a more parsimonious model? To shed light on this question, we consider submodels of equation (4.2) obtained by reducing the number of covariates in C_t. We test these more parsimonious specifications and report the rejection rates across firms in Table 3. For brevity, we focus on the economic policy uncertainty measure as the conditioning variable of the test. As a benchmark, the first two columns

---

[13]The economic policy uncertainty index is available at https://www.policyuncertainty.com, and the financial and macro uncertainty indexes are available at https://www.sydneyludvigson.com/macro-and-financial-uncertainty-indexes.

[14]Recall that the rank-transform of a variable X_t is obtained as $2\widehat{F}_X(X_t) - 1$, where $\widehat{F}_X$ is the empirical CDF of X_t. Further recall that the kth Legendre polynomial can be obtained as $\partial^k (x^2 - 1)^k / \partial x^k / (2^k k!)$.

TABLE 2. Empirical rejection rates using different conditioning variables.

	Policy Uncertainty				Financial Uncertainty				Macro Uncertainty			
	95%		99%		95%		99%		95%		99%	
m	VaR	CoVaR	VaR	CoVaR	VaR	CoVaR	VaR	CoVaR	VaR	CoVaR	VaR	CoVaR
6	0.165	0.064	0.068	0.030	0.204	0.072	0.072	0.031	0.206	0.049	0.081	0.017
7	0.166	0.068	0.067	0.017	0.200	0.073	0.058	0.021	0.193	0.046	0.080	0.015
8	0.170	0.074	0.066	0.009	0.205	0.075	0.058	0.020	0.188	0.050	0.071	0.014
9	0.163	0.078	0.067	0.007	0.196	0.065	0.055	0.018	0.183	0.045	0.070	0.010
10	0.157	0.072	0.067	0.007	0.179	0.053	0.058	0.014	0.162	0.035	0.070	0.008

Note: This table reports the cross-sectional empirical rejection frequencies of the specification tests for Adrian and Brunnermeier's (2016) VaR and CoVaR models given by equations (4.1) and (4.2), respectively. The left, middle, and right panels show results with the conditioning variable X_t being the policy, financial, and macro uncertainty measures, respectively. The quantile for the VaR and CoVaR models is fixed at $q \in \{95\%, 99\%\}$. All tests are implemented at the 5% significance level with Legendre polynomial basis with order $m \in \{6, \ldots, 10\}$, and are based on bootstrapped critical values with 1000 resamples.

labeled "None" report the rejection frequencies of the submodel that has no covariates from C_t. This simple specification is rejected for about 30% of the stocks for the 95%-CoVaR. It is rejected less frequently for the 99%-CoVaR, but the rejection frequency is nevertheless still markedly higher than the 5% nominal level. These findings are in contrast to those in Table 2, suggesting that at least *some* of the covariates are crucial in the CoVaR model.

As the next step, we include each of the covariates one at a time in the CoVaR model and report the rejection frequencies. The results reveal that, when we control for the equity volatility (MktSD), the rejection rates fall to around the levels presented in Table 2, which used the full set of covariates. Meanwhile, controlling for any of the other six covariates (e.g., Housing, TED, etc.) contributes little to reducing the rejection rates, suggesting that these additional covariates are individually unimportant. To confirm this conjecture, we finally consider a CoVaR specification with all covariates included *except* for MktSD, with the corresponding rejection rates reported in the last two columns of Table 3, labeled "All\MktSD." The rejection rates are similar to those with no covariates, which suggests that these covariates are not important for the CoVaR specification jointly. Overall, the results in Table 3 point to equity volatility as the most (and perhaps only) important covariate in the CoVaR specification.

5. Conclusion

This paper proposes a new consistent specification test for quantile models in time-series analysis. Our test is based on a nonparametric, series-based, estimate of a conditional moment that is known to equal zero when the model is correctly specified. We extend the uniform nonparametric inference method of Li and Liao (2020) in two directions that are critical for our empirical application. First, to overcome a size distortion we discover when the asymptotic Gaussian approximation is directly used in quantile applications near the tail, we propose a bootstrap method to obtain critical values for our test statistic. Establishing the validity of the proposed bootstrap method is nonstandard because a "non-Donsker" feature present in the original problem also manifests in

TABLE 3. CoVaR specification with different covariates.

	None		Housing		MktRet		TED		Credit	
m	95%	99%	95%	99%	95%	99%	95%	99%	95%	99%
6	0.293	0.167	0.346	0.115	0.194	0.159	0.255	0.089	0.246	0.161
7	0.329	0.158	0.351	0.077	0.279	0.139	0.249	0.057	0.292	0.148
8	0.309	0.144	0.340	0.059	0.289	0.147	0.257	0.031	0.276	0.138
9	0.287	0.121	0.325	0.069	0.266	0.150	0.250	0.030	0.267	0.123
10	0.301	0.083	0.324	0.038	0.274	0.127	0.230	0.013	0.267	0.091

	Yld3m		TERM		MktSD			All\MktSD	
m	95%	99%	95%	99%	95%	99%		95%	99%
6	0.332	0.127	0.335	0.166	0.053	0.013		0.252	0.123
7	0.366	0.114	0.371	0.143	0.052	0.009		0.274	0.109
8	0.375	0.081	0.370	0.125	0.055	0.008		0.278	0.084
9	0.342	0.087	0.354	0.110	0.057	0.004		0.268	0.078
10	0.332	0.059	0.339	0.062	0.052	0.005		0.283	0.049

Note: This table reports the cross-sectional empirical rejection frequencies of the specification tests for the q-CoVaR model, $q \in \{95\%, 99\%\}$, with different subsets of covariates. The column labels indicate the variable(s) included in C_t for each sub-model. The conditioning variable X_t is the economic policy uncertainty measure. All tests are implemented at the 5% significance level with Legendre polynomial basis with order $m \in \{6, \ldots, 10\}$, and are based on bootstrapped critical values with 1000 resamples.

the "bootstrap world." Second, we deal with the issue that both the dependent variable and the (growing number of) independent variables in our series regression contain estimated parameters that enter in a nonsmooth way. We apply the proposed new tests to a detailed analysis of the well-known CoVaR measure of Adrian and Brunnermeier (2016). We find that their specification for individual firm VaR is rejected more often than expected given the significance level of the test, suggesting that this model can be improved either via additional covariates or a more flexible specification. We also find that just one of the seven covariates used in the CoVaR specification is important for this model passing our specification tests. This suggests either that the CoVaR model can be made much more parsimonious, or perhaps more likely, that there are other covariates not considered in Adrian and Brunnermeier (2016) that could improve the explanatory power of the model.

Appendix: Proofs

Throughout the proofs, we use K to denote a generic finite constant that may change from line to line. For $p \geq 1$, let $\|\cdot\|_p$ denote the L_p norm for random variables. For notational simplicity, we write $\sum_t$ in place of $\sum_{t=1}^{n}$. The proofs rely on several technical lemmas, Lemmas S1–S12, which are collected in the Online Supplemental Material.

Proof of Theorem 1. By Lemma S2 and Lemma S7,

$$\|\widehat{\Sigma}_n - \Sigma_n\|_S = o_p(1). \tag{A.1}$$

Since the eigenvalues of Q_n and A_n are bounded from above and away from zero, the eigenvalues of Σ_n satisfy the same property and, by (A.1), we also have

$$\lambda_{\min}^{-1}(\widehat{\Sigma}_n) + \lambda_{\max}(\widehat{\Sigma}_n) = O_p(1). \tag{A.2}$$

These lemmas also show that, under the null hypothesis,

$$\|\widehat{\Sigma}_n - \Sigma_n\|_S = O_p\big(\delta_{Q,n} + \delta_{A,n} + \zeta_{0,n} m_n^{1/2} n^{-1/4} + \zeta_{1,n} m_n^{1/2} n^{-1/2}\big). \tag{A.3}$$

By (A.2), (A.3), and Assumptions 2(ii), (iv),

$$
\begin{aligned}
\sup_{x \in \mathcal{X}} \left| \frac{\sigma_n(x)}{\widehat{\sigma}_n(x)} - 1 \right| &= \sup_{x \in \mathcal{X}} \left| \frac{\widehat{\sigma}_n^2(x) - \sigma_n^2(x)}{\widehat{\sigma}_n(x)\big(\widehat{\sigma}_n(x) + \sigma_n(x)\big)} \right| \\
&\leq \sup_{x \in \mathcal{X}} \left| \frac{P(x)^\top (\widehat{\Sigma}_n - \Sigma_n) P(x)}{P(x)^\top \widehat{\Sigma}_n P(x)} \right| \\
&= O_p\big(\delta_{Q,n} + \delta_{A,n} + \zeta_{0,n} m_n^{1/2} n^{-1/4} + \zeta_{1,n} m_n^{1/2} n^{-1/2}\big) \\
&= o_p\big((\log n)^{-1}\big). \tag{A.4}
\end{aligned}
$$

Note that $h(\cdot) = 0$ and $b_n^\star = 0$ under the null hypothesis. We can then decompose

$$
\begin{aligned}
n^{1/2}\big(\widehat{b}_n - b_n^\star\big) = \widehat{Q}_n^{-1}\left(n^{-1/2} \sum_t P(X_t^\star) u_t^\star \right) &+ \widehat{Q}_n^{-1}\left(n^{-1/2} \sum_t (P(\widehat{X}_t) - P(X_t^\star)) u_t^\star \right) \\
&+ \widehat{Q}_n^{-1}\left(n^{-1/2} \sum_t P(\widehat{X}_t)\big(\widehat{Z}_{t+1} - Z_{t+1}^\star\big) \right).
\end{aligned}
$$

Observe that $\hat{h}_n(x) - h(x) = P(x)^\top(\widehat{b}_n - b_n^\star) + P(x)^\top b_n^\star - h(x)$. We thus have

$$
\sup_{x \in \mathcal{X}} \left| \frac{n^{1/2}\big(\widehat{h}_n(x) - h(x)\big)}{\widehat{\sigma}_n(x)} - \frac{P(x)^\top \widehat{Q}_n^{-1}\left(n^{-1/2} \sum_t P(X_t^\star) u_t^\star \right)}{\widehat{\sigma}_n(x)} \right| \leq \sum_{j=1}^{4} \sup_{x \in \mathcal{X}} |R_{j,n}(x)|, \tag{A.5}
$$

where

$$R_{1,n}(x) \equiv \frac{P(x)^\top (\widehat{Q}_n^{-1} - Q_n^{-1})\left(n^{-1/2} \sum_t P(X_t^\star) u_t^\star \right)}{\widehat{\sigma}_n(x)},$$

$$R_{2,n}(x) \equiv \frac{P(x)^\top \widehat{Q}_n^{-1}\left(n^{-1/2} \sum_t (P(\widehat{X}_t) - P(X_t^\star)) u_t^\star \right)}{\widehat{\sigma}_n(x)},$$

$$R_{3,n}(x) \equiv \frac{P(x)^\top \widehat{Q}_n^{-1}\left(n^{-1/2} \sum_t P(\widehat{X}_t)\big(\widehat{Z}_{t+1} - Z_{t+1}^\star\big) \right)}{\widehat{\sigma}_n(x)},$$

$$R_{4,n}(x) \equiv \frac{n^{1/2}\big(P(x)^\top b_n^\star - h(x)\big)}{\widehat{\sigma}_n(x)}.$$

Note that $\|n^{-1/2}\sum_t P(X_t^\star)u_t^\star\| = O_p(m_n^{1/2})$. By Lemma S2 and Assumptions 2(ii), (iv),

$$\sup_{x\in\mathcal{X}}\big|R_{1,n}(x)\big| = O_p\big(m_n^{1/2}\big(\delta_{Q,n} + \zeta_{1,n}m_n^{1/2}n^{-1/2}\big)\big) = o_p\big((\log n)^{-1/2}\big).$$

Since $\hat{\theta}_n = \theta^\star + O_p(n^{-1/2})$, by Assumption 2(ii) and Lemma S5, it is easy to see that

$$\sup_{x\in\mathcal{X}}\big|R_{2,n}(x)\big| \le O_p\big(\zeta_{1,n}m_n^{1/2}n^{-1/2}\big) = o_p\big((\log n)^{-1/2}\big).$$

By (A.4) and Lemma S8,

$$\sup_{x\in\mathcal{X}}\big|R_{3,n}(x)\big| = o_p\big((\log n)^{-1/2}\big).$$

Finally, we note that

$$\sup_{x\in\mathcal{X}}\big|R_{4,n}(x)\big| \le O_p(1)\sup_{x\in\mathcal{X}}\frac{\big\|n^{1/2}\big(P(x)^\top b_n^\star - h(x)\big)\big\|}{\|P(x)\|} = o_p\big((\log n)^{-1/2}\big),$$

where the inequality follows from (A.2), and the equality follows from Assumptions 1(vii) and 2. From the estimates above, we see that $\sup_{x\in\mathcal{X}}|R_{j,n}(x)|$ for each $1 \le j \le 4$. Hence, by (A.5),

$$\sup_{x\in\mathcal{X}}\left|\frac{n^{1/2}\big(\widehat{h}_n(x) - h(x)\big)}{\widehat{\sigma}_n(x)} - \frac{P(x)^\top Q_n^{-1}\big(n^{-1/2}\sum_t P(X_t^\star)u_t^\star\big)}{\widehat{\sigma}_n(x)}\right| = o_p\big((\log n)^{-1/2}\big). \quad \text{(A.6)}$$

Let ξ_n be defined as in Assumption 3. By the concentration property of Gaussian processes, we have

$$\sup_{x\in\mathcal{X}}\left|\frac{P(x)^\top \Sigma_n^{1/2}\xi_n}{\sigma_n(x)}\right| = O_p\big((\log n)^{1/2}\big). \quad \text{(A.7)}$$

Then Assumption 3 further implies that

$$\sup_{x\in\mathcal{X}}\left|\frac{n^{-1/2}P(x)^\top Q_n^{-1}\sum_{t=1}^n P(X_t^\star)u_t^\star}{\sigma_n(x)}\right| = O_p\big((\log n)^{1/2}\big). \quad \text{(A.8)}$$

By (A.4) and (A.8),

$$
\left\| \sup_{x \in \mathcal{X}} \left| \frac{n^{-1/2} P(x)^\top Q_n^{-1} \sum_{t=1}^n P(X_t^\star) u_t^\star}{\widehat{\sigma}_n(x)} \right| - \sup_{x \in \mathcal{X}} \left| \frac{n^{-1/2} P(x)^\top Q_n^{-1} \sum_{t=1}^n P(X_t^\star) u_t^\star}{\sigma_n(x)} \right| \right\|
$$

$$
\leq \sup_{x \in \mathcal{X}} \left| \frac{n^{-1/2} P(x)^\top Q_n^{-1} \sum_{t=1}^n P(X_t^\star) u_t^\star}{\sigma_n(x)} \right| \sup_{x \in \mathcal{X}} \left| \frac{\sigma_n(x)}{\widehat{\sigma}_n(x)} - 1 \right|
$$

$$
= o_p\big((\log n)^{-1/2}\big). \tag{A.9}
$$

Combining (A.6) and (A.9), we deduce

$$
\sup_{x \in \mathcal{X}} \left| \frac{n^{1/2} \big(\widehat{h}_n(x) - h(x)\big)}{\widehat{\sigma}_n(x)} \right| - \sup_{x \in \mathcal{X}} \left| \frac{n^{-1/2} P(x)^\top Q_n^{-1} \sum_{t=1}^n P(X_t^\star) u_t^\star}{\sigma_n(x)} \right| = o_p\big((\log n)^{-1/2}\big).
$$

The assertion of the theorem then readily follows from Assumption 3 and the above approximation. □

Next, we prove Theorem 2 in the main text. We need to explicitly introduce some notation for various bootstrap quantities:

$$
\widehat{Q}_n^* = n^{-1} \sum_t P(\widehat{X}_t^*) P(\widehat{X}_t^*)^\top, \qquad \widehat{b}_n^* = (\widehat{Q}_n^*)^{-1} \left(n^{-1} \sum_t P(\widehat{X}_t^*) \widehat{Z}_{t+1}^* \right),
$$

$$
\widehat{u}_t^* = \widehat{Z}_{t+1}^* - P(\widehat{X}_t^*)^\top \widehat{b}_n^*, \qquad \widehat{A}_n^* = n^{-1} \sum_t P(\widehat{X}_t^*) P(\widehat{X}_t^*)^\top (\widehat{u}_t^*)^2,
$$

$$
\widehat{\Sigma}_n^* = (\widehat{Q}_n^*)^{-1} \widehat{A}_n^* (\widehat{Q}_n^*)^{-1}.
$$

PROOF OF THEOREM 2. (a) Step 1. We prove the assertion in part (a) in four steps. In this step, we show that

$$
\sup_{x \in \mathcal{X}} \left| \frac{\sigma_n(x)}{\widehat{\sigma}_n^*(x)} - 1 \right| = o_p\big(m_n^{-1/2} (\log n)^{-1/2}\big). \tag{A.10}
$$

By Lemma S2 and Lemma S9,

$$
\big\| \widehat{Q}_n^* - Q_n \big\|_S = O_p\big(\delta_{Q,n} + \zeta_{1,n} m_n^{1/2} n^{-1/2} + \zeta_{0,n} \log(m_n)^{1/2} m_n^{1/2} n^{-1/2}\big). \tag{A.11}
$$

Under Assumptions 2(ii), (iv),

$$
m_n^{1/2} \big(\delta_{Q,n} + \delta_{A,n} + \zeta_{1,n} m_n^{1/2} n^{-1/2} + \zeta_{0,n} m_n^{1/2} n^{-1/4}\big) = o\big((\log n)^{-1/2}\big). \tag{A.12}
$$

By Lemma S7 and Lemma S11, we see that under the null hypothesis,

$$
\begin{aligned}
\big\| \widehat{A}_n^* - A_n \big\|_S &= O_p\big(\zeta_{0,n} m_n n^{-1/2}\big) + O_p\big(\delta_{A,n} + \zeta_{1,n} m_n^{1/2} n^{-1/2} + \zeta_{0,n} m_n^{1/2} n^{-1/4}\big) \\
&= O_p\big(\delta_{A,n} + \zeta_{1,n} m_n^{1/2} n^{-1/2} + \zeta_{0,n} m_n^{1/2} n^{-1/4}\big).
\end{aligned}
$$

Combining this estimate with (A.11), we obtain (using (A.12))

$$
\begin{aligned}
\big\| \widehat{\Sigma}_n^* - \Sigma_n \big\|_S &= O_p\big(\delta_{Q,n} + \delta_{A,n} + \zeta_{1,n} m_n^{1/2} n^{-1/2} + \zeta_{0,n} m_n^{1/2} n^{-1/4}\big) \\
&= o_p\big(m_n^{-1/2}(\log n)^{-1/2}\big).
\end{aligned}
\tag{A.13}
$$

Consequently,

$$
\begin{aligned}
\sup_{x \in \mathcal{X}} \left| \frac{\sigma_n(x)}{\widehat{\sigma}_n^*(x)} - 1 \right| &= \sup_{x \in \mathcal{X}} \left| \frac{P(x)^\top (\widehat{\Sigma}_n^* - \Sigma_n) P(x)}{\widehat{\sigma}_n^*(x)(\widehat{\sigma}_n^*(x) + \sigma_n(x))} \right| \\
&\leq \sup_{x \in \mathcal{X}} \left| \frac{P(x)^\top (\widehat{\Sigma}_n^* - \Sigma_n) P(x)}{\widehat{\sigma}_n^*(x)^2} \right| \\
&\leq \lambda_{\min}^{-1}\big(\widehat{\Sigma}_n^*\big) \big\| \widehat{\Sigma}_n^* - \Sigma_n \big\|_S \\
&= o_p\big(m_n^{-1/2}(\log n)^{-1/2}\big)
\end{aligned}
\tag{A.14}
$$

which proves (A.10).

Step 2. Let $\widehat{\omega}_n^* \equiv n^{-1/2} \sum_t (P(\widehat{X}_t^*)\widehat{Z}_{t+1}^* - \mathbb{E}^*[P(\widehat{X}_t^*)\widehat{Z}_{t+1}^*])$. In this step, we show that

$$
\frac{n^{1/2} P(x)^\top (\widehat{b}_n^* - \widehat{b}_n)}{\widehat{\sigma}_n^*(x)} = \frac{P(x)^\top Q_n^{-1}}{\sigma_n(x)} \widehat{\omega}_n^* + o_p\big((\log n)^{-1/2}\big)
\tag{A.15}
$$

uniformly over $x \in \mathcal{X}$.

First, by Assumption 2(iii),

$$
\frac{\|P(x)\|}{\sigma_n(x)} = O(1).
\tag{A.16}
$$

Then under the null hypothesis, by (A.10) and Lemma S10,

$$
\begin{aligned}
\sup_{x \in \mathcal{X}} &\left| \frac{n^{1/2} P(x)^\top (\widehat{b}_n^* - \widehat{b}_n)}{\widehat{\sigma}_n^*(x)} - \frac{n^{1/2} P(x)^\top (\widehat{b}_n^* - \widehat{b}_n)}{\sigma_n(x)} \right| \\
&\leq \left| \frac{n^{1/2} P(x)^\top (\widehat{b}_n^* - \widehat{b}_n)}{\sigma_n(x)} \right| \sup_{x \in \mathcal{X}} \left| \frac{\sigma_n(x)}{\widehat{\sigma}_n^*(x)} - 1 \right| \\
&\leq \big\| n^{1/2} (\widehat{b}_n^* - \widehat{b}_n) \big\| \frac{\|P(x)\|}{\sigma_n(x)} \sup_{x \in \mathcal{X}} \left| \frac{\sigma_n(x)}{\widehat{\sigma}_n^*(x)} - 1 \right| = o_p\big((\log n)^{-1/2}\big).
\end{aligned}
\tag{A.17}
$$

We can further decompose

$$
\frac{n^{1/2} P(x)^\top (\widehat{b}_n^* - \widehat{b}_n)}{\sigma_n(x)} = \frac{P(x)^\top (\widehat{Q}_n^*)^{-1}}{\sigma_n(x)} \widehat{\omega}_n^* + \frac{P(x)^\top (\widehat{Q}_n^*)^{-1}}{\sigma_n(x)} (\widehat{Q}_n - \widehat{Q}_n^*)(n^{1/2}\widehat{b}_n)
\tag{A.18}
$$

for any $x \in \mathcal{X}$. Under the null hypothesis, Lemma S6 implies that $n^{1/2}\widehat{b}_n = O_p(m_n^{1/2})$. By Assumption 2(iii), Lemma S2, Lemma S9, and (A.16),

$$\frac{\|P(x)^\top (\widehat{Q}_n^*)^{-1}\|}{\sigma_n(x)} = O_p(1). \tag{A.19}$$

Then, by Lemma S9,

$$\left\| \frac{P(x)^\top (\widehat{Q}_n^*)^{-1}}{\sigma_n(x)} (\widehat{Q}_n - \widehat{Q}_n^*)(n^{1/2}\widehat{b}_n) \right\| = O_p\big(\zeta_{0,n}\big(\log(m_n)m_n n^{-1}\big)^{1/2} m_n^{1/2}\big)$$

$$= o_p\big((\log n)^{-1/2}\big), \tag{A.20}$$

where the second line is implied by Assumption 2(ii). By Assumption 2(iii), Lemma S7, and Lemma S12,

$$\|\widehat{\omega}_n^*\| = \left\| n^{-1/2} \sum_t \big(P(\widehat{X}_t^*)\widehat{Z}_{t+1}^* - \mathbb{E}^*\big[P(\widehat{X}_t^*)\widehat{Z}_{t+1}^*\big]\big) \right\| = O_p\big(m_n^{1/2}\big) \tag{A.21}$$

which together with Assumption 2(iii), Lemma S2, Lemma S9, and (A.16) implies that

$$\left\| \frac{P(x)^\top \big((\widehat{Q}_n^*)^{-1} - Q_n^{-1}\big)}{\sigma_n(x)} \widehat{\omega}_n^* \right\| = O_p\big(\zeta_{0,n}\big(\log(m_n)m_n n^{-1}\big)^{1/2} m_n^{1/2}\big)$$

$$= o_p\big((\log n)^{-1/2}\big), \tag{A.22}$$

where the second equality is implied by Assumption 2(ii). The claim in (A.15) follows from (A.17), (A.18), (A.20), and (A.22).

Step 3. In this step, we show that there exists a sequence ξ_n^* of standard m_n-dimensional Gaussian vectors such that

$$\left| \sup_{x \in \mathcal{X}} \frac{|P(x)^\top Q_n^{-1} \widehat{\omega}_n^*|}{\sigma_n(x)} - \sup_{x \in \mathcal{X}} \frac{|P(x)^\top Q_n^{-1} A_n^{1/2} \xi_n^*|}{\sigma_n(x)} \right| = o_p\big((\log n)^{-1/2}\big). \tag{A.23}$$

By Assumption 2(iii), Lemma S7, and Lemma S12,

$$\lambda_{\min}^{-1}\big(\widehat{H}_n^*\big) + \lambda_{\max}\big(\widehat{H}_n^*\big) = O_p(1). \tag{A.24}$$

Denote $\alpha_n(x) = Q_n^{-1}P(x)/\sigma_n(x)$, and note that $\sup_{x \in \mathcal{X}} \|\alpha_n(x)\| \le C_\alpha$ for some finite constant C_α. For given n and $x \in \mathcal{X}$, define $f_{n,x}(z, v) = z\alpha_n(x)^\top P(v)$ for any $z \in \{0, 1\}$ and any $v \in \mathcal{X} \oplus \varepsilon_n$. We then consider the class of functions $\mathcal{F}_n = \mathcal{F}_n^0 \cup (-\mathcal{F}_n^0)$, with $\mathcal{F}_n^0 \equiv \{f_{n,x}(z, t) : x \in \mathcal{X}\}$. To prove (A.23), we shall apply Corollary 2.2 in Chernozhukov, Chetverikov, and Kato (2013) to $\mathcal{F}_n$ under the $\mathcal{D}_n$-conditional probability.

Specifically, by the Cauchy–Schwarz inequality,

$$\|f_{n,x}(z, t)\| \le \zeta_{0,n} m_n^{1/2} \sup_{x \in \mathcal{X}} \|\alpha_n(x)\| \le C_\alpha \zeta_{0,n} m_n^{1/2}.$$

In addition, for any $x_1, x_2 \in \mathcal{X}$,

$$\big|f_{n,x_1}(z, t) - f_{n,x_2}(z, t)\big| \le \zeta_{0,n} m_n^{1/2} \|\alpha_n(x_1) - \alpha_n(x_2)\| \le K\zeta_{0,n} m_n^{1/2} \zeta_n^L \|x_1 - x_2\|, \tag{A.25}$$

where we recall the definition of ζ_n^L from Assumption 2. Therefore, $\mathcal{F}_n$ forms a VC-type class with (constant) envelope $F \equiv C_\alpha \zeta_{0,n} m_n^{1/2}$, and it satisfies the following uniform entropy condition for some constant A:

$$\sup_{\mathbb{Q}} N\big(\mathcal{F}_n, \|\cdot\|_{\mathbb{Q},2}, \epsilon\|F\|_{\mathbb{Q},2}\big) \le \big(A\xi_n^L/\epsilon\big)^{d_x},$$

where the supremum is taken over all finitely discrete probability measures, and we denote by $N(\mathcal{F}_n, \|\cdot\|_{\mathbb{Q},2}, \epsilon\|F\|_{\mathbb{Q},2})$ the covering number for $\mathcal{F}_n$ under the $L^2(\mathbb{Q})$ norm. By (A.24), we have for any $n \ge 1$,

$$\mathbb{E}^*\big[\big|\alpha_n(x)^\top\big(P\big(\widehat{X}_t^*\big)\widehat{Z}_{t+1}^* - \mathbb{E}^*\big[P\big(\widehat{X}_t^*\big)\widehat{Z}_{t+1}^*\big]\big)\big|^2\big] \le K\lambda_{\max}\big(\widehat{H}_n^*\big) = O_p(1),$$

and

$$\mathbb{E}^*\big[\big|\alpha_n(x)^\top\big(P\big(\widehat{X}_t^*\big)\widehat{Z}_{t+1}^* - \mathbb{E}^*\big[P\big(\widehat{X}_t^*\big)\widehat{Z}_{t+1}^*\big]\big)\big|^3\big]$$
$$\le K\zeta_{0,n}m_n^{1/2}\mathbb{E}^*\big[\big|\alpha_n(x)^\top\big(P\big(\widehat{X}_t^*\big)\widehat{Z}_{t+1}^* - \mathbb{E}^*\big[P\big(\widehat{X}_t^*\big)\widehat{Z}_{t+1}^*\big]\big)\big|^2\big]$$
$$= O_p\big(\zeta_{0,n}m_n^{1/2}\big).$$

Thus, applying Corollary 2.2 in Chernozhukov, Chetverikov, and Kato (2013) under the $\mathcal{D}_n$-conditional probability (with $q = \infty$, $\gamma = 1/\log n$, $b = C_\alpha \zeta_{0,n} m_n^{1/2}$, $\sigma = O(1)$, and $K_n = O(\log n)$ in their notation) shows that there exists a sequence ξ_n^* of m_n-dimensional standard Gaussian vectors such that

$$\left|\sup_{x\in\mathcal{X}}\big|\alpha_n(x)^\top\widehat{\omega}_n^*\big| - \sup_{x\in\mathcal{X}}\big|\alpha_n(x)^\top\big(\widehat{H}_n^*\big)^{1/2}\xi_n^*\big|\right|$$
$$= O_p\big(\zeta_{0,n}m_n^{1/2}(\log n)^{3/2}n^{-1/2} + \zeta_{0,n}^{1/2}m_n^{1/4}(\log n)^{5/4}n^{-1/4} + \zeta_{0,n}^{1/3}m_n^{1/6}\log(n)n^{-1/6}\big)$$
$$= o_p\big((\log n)^{-1/2}\big), \tag{A.26}$$

where the $o_p((\log n)^{-1/2})$ statement follows from Assumption 2(ii).

Under the null hypothesis, Lemma S7 and Lemma S12 imply that

$$\big\|\widehat{H}_n^* - A_n\big\|_S = O_p\big(\delta_{A,n} + \zeta_{1,n}m_n^{1/2}n^{-1/2} + \zeta_{0,n}m_n^{1/2}n^{-1/4}\big) + O_p\big(\zeta_{0,n}m_n n^{-1/2}\big)$$
$$= O_p\big(\delta_{A,n} + \zeta_{1,n}m_n^{1/2}n^{-1/2} + \zeta_{0,n}m_n^{1/2}n^{-1/4}\big). \tag{A.27}$$

By (A.12) and (A.27),

$$\sup_{x\in\mathcal{X}}\big|\alpha_n(x)^\top\big(\big(\widehat{H}_n^*\big)^{1/2} - A_n^{1/2}\big)\xi_n^*\big|$$
$$= O_p\big(m_n^{1/2}\big(\delta_{A,n} + \zeta_{1,n}m_n^{1/2}n^{-1/2} + \zeta_{0,n}m_n^{1/2}n^{-1/4}\big)\big)$$
$$= o_p\big((\log n)^{-1/2}\big). \tag{A.28}$$

Together with (A.26), this estimate further implies

$$\left|\sup_{x\in\mathcal{X}}\big|\alpha_n(x)^\top\widehat{\omega}_n^*\big| - \sup_{x\in\mathcal{X}}\big|\alpha_n(x)^\top A_n^{1/2}\xi_n^*\big|\right| = o_p\big((\log n)^{-1/2}\big), \tag{A.29}$$

as asserted in (A.23).

Step 4. By (A.15) and (A.23), we complete the proof of part (a) as follows:

$$\left| \sup_{x \in \mathcal{X}} \frac{|n^{1/2} P(x)^\top (\widehat{b}_n^* - \widehat{b}_n)|}{\widehat{\sigma}_n^*(x)} - \sup_{x \in \mathcal{X}} \frac{|P(x)^\top Q_n^{-1} A_n^{1/2} \xi_n^*|}{\sigma_n(x)} \right|$$

$$\leq \left| \sup_{x \in \mathcal{X}} \frac{|n^{1/2} P(x)^\top (\widehat{b}_n^* - \widehat{b}_n)|}{\widehat{\sigma}_n^*(x)} - \sup_{x \in \mathcal{X}} \frac{|P(x)^\top Q_n^{-1} \widehat{\omega}_n^*|}{\sigma_n(x)} \right|$$

$$+ \left| \sup_{x \in \mathcal{X}} \frac{|P(x)^\top Q_n^{-1} \widehat{\omega}_n^*|}{\sigma_n(x)} - \sup_{x \in \mathcal{X}} \frac{|P(x)^\top Q_n^{-1} A_n^{1/2} \xi_n^*|}{\sigma_n(x)} \right|$$

$$= o_p\big((\log n)^{-1/2}\big).$$

(b) The size property of the test follows from part (a) of Theorem 2. It remains to show the claimed power property. Observe

$$\left| \sup_{x \in \mathcal{X}} \frac{|n^{1/2} \widehat{h}_n(x)|}{\widehat{\sigma}_n(x)} - \sup_{x \in \mathcal{X}} \frac{|n^{1/2} h(x)|}{\widehat{\sigma}_n(x)} \right| \leq \sup_{x \in \mathcal{X}} \frac{n^{1/2} |\widehat{h}_n(x) - h(x)|}{\widehat{\sigma}_n(x)}. \tag{A.30}$$

Note that $\sup_{x \in \mathcal{X}} \widehat{\sigma}_n(x) \leq \sup_{x \in \mathcal{X}} \|P(x)\| \lambda_{\max}^{1/2}(\widehat{\Sigma}_n) = O_p(\zeta_{0,n} m_n^{1/2})$. Under the alternative hypothesis, we have $\sup_{x \in \mathcal{X}} |h(x)| > 0$. Therefore, $\sup_{x \in \mathcal{X}} |n^{1/2} h(x)|/\widehat{\sigma}_n(x)$ diverges to infinity in probability at a rate that is at least $n^{1/2}/(\zeta_{0,n} m_n^{1/2})$. In addition, we observe

$$\sup_{x \in \mathcal{X}} \frac{n^{1/2} |\widehat{h}_n(x) - h(x)|}{\widehat{\sigma}_n(x)}$$

$$\leq O_p(n^{1/2}) \left(\sup_{x \in \mathcal{X}} \frac{\|P(x)\| \|\widehat{b}_n - b_n^\star\|}{\|P(x)\|} + \sup_{x \in \mathcal{X}} \frac{|P(x)^\top b_n^\star - h(x)|}{\|P(x)\|} \right)$$

$$= O_p(n^{1/2} \delta_{b,n}), \tag{A.31}$$

where the inequality follows from the triangle inequality and the Cauchy–Schwarz inequality, and the latter rate statement follows from Assumption 2 and Lemma S6 (with $\delta_{b,n} \equiv \zeta_{1,n} m_n^{1/2} n^{-1/2} + \zeta_{0,n} m_n^{1/2} n^{-3/4}$). Under Assumption 2, $n^{1/2} \delta_{b,n} = o(n^{1/2}/(\zeta_{0,n} m_n^{1/2}))$. Therefore, the sup-t statistic $\sup_{x \in \mathcal{X}} |n^{1/2} \widehat{h}_n(x)|/\widehat{\sigma}_n(x)$ also diverges to infinity in probability at a rate that is at least $n^{1/2}/(\zeta_{0,n} m_n^{1/2})$.

From (A.1), Lemma S9, and Lemma S11, it is easy to see that $\lambda_{\min}^{-1}(\widehat{\Sigma}_n^*) = O_p(1)$. Therefore,

$$\sup_{x \in \mathcal{X}} \frac{|n^{1/2} P(x)^\top (\widehat{b}_n^* - \widehat{b}_n)|}{\widehat{\sigma}_n^*(x)} \leq \lambda_{\min}^{-1/2}(\widehat{\Sigma}_n^*) \|n^{1/2}(\widehat{b}_n^* - \widehat{b}_n)\|$$

$$= O_p\big(\zeta_{0,n} \log(m_n)^{1/2} m_n^{1/2}\big)$$

$$= o_p\big(n^{1/2}/(\zeta_{0,n} m_n^{1/2})\big)$$

where the first equality is by Lemma S10, and the second equality holds under the maintained rate requirement in Assumption 2. In view of the fact that the sup-t statistic diverges to infinity at rate that is at least $n^{1/2}/(\zeta_{0,n} m_n^{1/2})$, we can conclude that the test rejects the alternative hypothesis with probability approaching 1. $\qquad \square$

References

Adrian, T. and M. K. Brunnermeier (2016), "CoVaR." *American Economic Review*, 106 (7), 1705. [125, 127, 129, 139, 140, 141, 142]

Andrews, D. W. K. (1994), "Empirical process methods in econometrics." In *Handbook of Econometrics*, Vol. 4, 2247–2294, Elsevier. [132]

Baker, S. R., N. Bloom, and S. J. Davis (2016), "Measuring economic policy uncertainty." *Quarterly Journal of Economics*, 131 (4), 1593–1636. [140]

Basel Committee (2010), "Basel III: A global regulatory framework for more resiliant banks and banking systems, bank for international settlements." http://www.bis.org/publ/bcbs189.pdf. [126]

Belloni, A., V. Chernozhukov, D. Chetverikov, and K. Kato (2015), "Some new asymptotic theory for least squares series: Pointwise and uniform results." *Journal of Econometrics*, 186 (2), 345–366. [126, 130]

Bierens, H. J. (1982), "Consistent model specification tests." *Journal of Econometrics*, 20, 105–134. [128]

Bierens, H. J. (1990), "A consistent conditional moment test of functional form." *Econometrica*, 58, 1443–1458. [128]

Bierens, H. J. and D. K. Ginther (2001), "Integrated conditional moment testing of quantile regression models." *Empirical Economics*, 26, 307–324. [128]

Bierens, H. J. and W. Ploberger (1997), "Asymptotic theory of integrated conditional moment tests." *Econometrica*, 65 (5), 1129–1151. [128]

Bontemps, C. (2019), "Moment-based tests under parameter uncertainty." *Review of Economics and Statistics*, 101 (1), 146–159. [137]

Chen, X. (2007), "Large sample sieve estimation of semi-nonparametric models." In *Handbook of Econometrics* (J. Heckman and E. Leamer, eds.). vol. 6B, Chapter 76, 1 edn. Elsevier. [131, 134]

Chen, X. and T. M. Christensen (2015), "Optimal uniform convergence rates and asymptotic normality for series estimators under weak dependence and weak conditions." *Journal of Econometrics*, 188 (2), 447–465. [134]

Chernozhukov, V., D. Chetverikov, and K. Kato (2013), "Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors." *The Annals of Statistics*, 41 (6), 2786–2819. [147, 148]

Chernozhukov, V., S. Lee, and A. M. Rosen (2013), "Intersection bounds: Estimation and inference." *Econometrica*, 81 (2), 667–737. [126, 128, 130]

Christoffersen, P. F. (1998), "Evaluating interval forecasts." *International Economic Review*, 39 (4), 841–862. [127]

Dedecker, J., E. Rio, and F. Merlevède (2014), "Strong approximation of the empirical distribution function for absolutely regular sequences in $\mathbb{R}^d$." *Electronic Journal of Probability*, 19, 1–56. [138]

Dehling, H., O. Durieu, and D. Volny (2009), "New techniques for empirical processes of dependent data." *Stochastic Processes and their Applications*, 119 (10), 3699–3718. [138]

Engle, R. F. and S. Manganelli (2004), "CAViaR: Conditional autoregressive value at risk by regression quantiles." *Journal of Business and Economic Statistics*, 22, 367–381. [127]

Escanciano, J. C. and C. Velasco (2010), "Specification tests of parametric dynamic conditional quantiles." *Journal of Econometrics*, 159, 209–221. [128]

Giacomini, R. and H. White (2006), "Tests of conditional predictive ability." *Econometrica*, 74 (6), 1545–1578. [131]

Horvath, P., J. Li, Z. Liao, and A. J. Patton (2022), "Supplement to 'A consistent specification test for dynamic quantile models'." *Quantitative Economics Supplemental Material*, 13, https://doi.org/10.3982/QE1727. [128]

Jurado, K., S. C. Ludvigson, and S. Ng (2015), "Measuring uncertainty." *American Economic Review*, 105 (3), 1177–1216. [140]

Koenker, R. (2005), *Quantile Regression*. Cambridge University Press. [125]

Koenker, R. and G. Basset (1978), "Quantile regression." *Econometrica*, 46 (1), 33–50. [125]

Komunjer, I. (2005), "Quasi-maximum likelihood estimation for conditional quantiles." *Journal of Econometrics*, 128 (1), 137–164. [129]

Komunjer, I. (2013), *Quantile Prediction*. Handbook of Economic Forecasting, Vol. 2. Elsevier. [125, 127]

Li, J. and Z. Liao (2020), "Uniform nonparametric inference for time series." *Journal of Econometrics*, 219 (1), 38–51. [126, 127, 130, 132, 135, 136, 141]

McNeil, A. J., R. Frey, and P. Embrechts (2015), *Quantitative Risk Management*, Revised edition. Princeton University Press. [126]

Rio, E. (2017), *Asymptotic Theory of Weakly Dependent Random Processes*. Springer-Verlag, Berlin. [138]

van der Vaart, A. W. (1998), *Asymptotic Statistics*. Cambridge University Press. [132]