

Bootstrapping non-stationary stochastic volatility

H. Peter Boswijk^a, Giuseppe Cavaliere^{b,c,*}, Iliyan Georgiev^{b,d}, Anders Rahbek^e

^aAmsterdam School of Economics and Tinbergen Institute, University of Amsterdam,
1018 WB Amsterdam, The Netherlands

^bDepartment of Economics, University of Bologna, 40126 Bologna, Italy

^cDepartment of Economics, Exeter Business School, Exeter EX4 4PU, UK

^dInstitut d'Anàlisi Econòmica (CSIC), Barcelona, Spain

^eDepartment of Economics, University of Copenhagen, 1353 Copenhagen K, Denmark

Received 29 November 2019; Received in revised form 30 September 2020;

Accepted 8 January 2021

Abstract

In this paper we investigate to what extent the bootstrap can be applied to conditional mean models, such as regression or time series models, when the volatility of the innovations is random and possibly non-stationary. In fact, the volatility of many economic and financial time series displays persistent changes and possible non-stationarity. However, the theory of the bootstrap for such models has focused on deterministic changes of the unconditional variance and little is known about the performance and the validity of the bootstrap when the volatility is driven by a non-stationary stochastic process. This includes near-integrated exogenous volatility processes as well as near-integrated GARCH processes, where the conditional variance has a diffusion limit; a further important example is the case where volatility exhibits infrequent jumps. This paper fills this gap in the literature by developing conditions for bootstrap validity in time series and regression models with non-stationary, stochastic volatility. We show that in such cases the distribution of bootstrap statistics (conditional on the data) is random in the limit. Consequently, the conventional approaches to proofs of bootstrap consistency, based on the notion of weak convergence in probability of the bootstrap statistic, fail to deliver the required validity results. Instead, we use the concept of 'weak convergence in distribution' to develop and establish novel conditions for validity of the wild bootstrap, conditional on the volatility process. We apply our results to several

*Corresponding author.

E-mail addresses: h.p.boswijk@uva.nl (H. P. Boswijk), giuseppe.cavaliere@unibo.it (G. Cavaliere), i.georgiev@unibo.it (I. Georgiev), anders.rahbek@econ.ku.dk (A. Rahbek).

testing problems in the presence of non-stationary stochastic volatility, including testing in a location model, testing for structural change using CUSUM-type functionals, and testing for a unit root in autoregressive models. Importantly, we work under sufficient conditions for bootstrap validity that include the absence of statistical leverage effects, i.e., correlation between the error process and its future conditional variance. The results of the paper are illustrated using Monte Carlo simulations, which indicate that a wild bootstrap approach leads to size control even in small samples.

Keywords: Bootstrap; Non-stationary stochastic volatility; Random limit measures; Weak convergence in distribution.

1 Introduction

In this paper we consider bootstrap and asymptotic inference on the conditional mean in econometric time series models when the (conditional) volatility is allowed to show a large degree of persistence due to possible permanent and stochastic changes, reflecting the well established fact that volatility in many economic and financial time series displays high persistence, and covariance non-stationarity.

Earlier references in macroeconomics include Kim and Nelson (1999) and McConnell and Perez-Quiroz (2000), who find evidence of an (unanticipated) structural change in the volatility of US GDP growth rates. Evidence of changes in the unconditional volatility appear in many key time series, such as aggregate consumption and income, in interest rate data and in nominal and real price variables; see Sensier and van Dijk (2004). Evidence on changes in the long-run component of volatility in stock and currency markets are initially reported in Loretan and Phillips (1994) and Hansen (1995), who show that when stochastic volatility [SV] models are taken to the data, the largest autoregressive root in the SV process is so close to one that the assumption of stationary volatility seems to be at odds with the data. Similarly, it is a well-known stylized fact that GARCH models fit to stock market returns display parameter estimates which reflect high persistence as they (nearly) violate covariance stationarity conditions (often referred to as “near-integrated GARCH”), and that such parameters are smaller when a slowly-varying long run component is accounted for in the model, see Engle and Rangel (2008). Harvey, Leybourne, Sollis and Taylor (2016) list a number of empirical studies that have found strong evidence of structural breaks in the unconditional variance of asset returns, with break dates driven by major financial and macroeconomic crises. Such (possibly random) volatility shifts are known to affect the asymptotic properties of estimators of the parameters of models for the conditional mean; see Cavaliere and Taylor (2007), Xu and Phillips (2008) and, for multivariate

models, Cavaliere, Rahbek and Taylor (2010a,b) and Boswijk, Cavaliere, Rahbek and Taylor (2016).

In the framework of a conditional mean, or a general (stationary, or non-stationary) regression type model, the *wild bootstrap* is an important tool to deliver consistent estimation of the asymptotic distributions of test statistics or parameter estimators. The wild bootstrap allows in particular to track changes in the quadratic variation of an econometric model by simply mimicking the (unknown) volatility dynamics through the squared model residuals, see Gonçalves and Kilian (2004, 2007) for applications to stationary time series models and Cavaliere *et al.* (2010a,b) for non-stationary multivariate models.

Consider the simple case where the volatility, say σ_t , can be approximated by a *non-stochastic* element of the space $\mathcal{D}[0, 1]$ of càdlàg functions on $[0, 1]$, such that $\sigma_t = \sigma(t/n)$ ($t = 1, \dots, n$, n denoting the sample size) with $\sigma \in \mathcal{D}[0, 1]$. Simple special cases are a single volatility break at time $\lfloor n\tau \rfloor$ (with $\lfloor \cdot \rfloor$ denoting the (floor) integer value), for some $\tau \in (0, 1)$, as given by (with $\mathbb{I}_A(\cdot)$ denoting the indicator function of the set A)

$$\sigma(u) := \sigma_A + (\sigma_B - \sigma_A)\mathbb{I}_{[\tau, 1]}(u), \text{ for } u \in [0, 1],$$

where $\sigma_A > 0$ and $\sigma_B > 0$; or the case of trending volatility,

$$\sigma(u) := \sigma_A + (\sigma_B - \sigma_A)u^\delta,$$

where $\sigma_A > 0$, $\sigma_B > 0$ and $\delta > 0$. A classic wild bootstrap, based on the resampling scheme $\varepsilon_t^* = \hat{\varepsilon}_t w_t^*$, where the $\hat{\varepsilon}_t$'s are the estimated residuals from the regression model and the w_t^* 's are i.i.d. $(0, 1)$ bootstrap shocks, independent of the original sample, is in general able to track the volatility path (in terms of quadratic variation) of the original data, without any assumption on the initial values of the volatility process as, loosely speaking, under standard assumptions

$$n^{-1} \sum_{t=1}^{\lfloor nu \rfloor} (\varepsilon_t^*)^2 = n^{-1} \sum_{t=1}^{\lfloor nu \rfloor} \hat{\varepsilon}_t^2 (w_t^*)^2 \approx n^{-1} \sum_{t=1}^{\lfloor nu \rfloor} \varepsilon_t^2 + o_p(1) = \int_0^u \sigma(s)^2 ds + o_p(1).$$

Existing theory of the bootstrap mainly focuses on such deterministic changes of the unconditional variance and little is known about the performance and the validity¹ of the bootstrap when the volatility is driven by a high-persistent, or (second order) non-stationary stochastic process. This includes leading key cases such as near-integrated exogenous volatility processes (as analyzed by Hansen, 1995), as well as near-integrated GARCH processes, where the conditional variance has a diffusion limit (Nelson, 1990).

¹Throughout the paper, with ‘validity’ of the bootstrap we mean that the associated bootstrap tests control size asymptotically. With ‘consistency of the bootstrap test’ we mean that the (bootstrap) test rejects with probability tending to one under the alternative.

This paper fills this gap in the literature by developing conditions for bootstrap validity and consistency of the associated bootstrap tests in regression and time series models with persistent stochastic volatility. That is, we replace the deterministic volatility assumption by allowing that volatility is the realization of a (non-stationary) stochastic process σ_t ; specifically, we derive results under the general assumption that, for the càdlàg version of the volatility, it holds that

$$\sigma_{\lfloor un+1 \rfloor} \xrightarrow{w} \sigma(u) \text{ for } u \in [0, 1], \quad (1)$$

where σ is some random element in $\mathcal{D}[0, 1]$.

The analysis of the bootstrap under a weak convergence assumption like (1) is not straightforward. As we show in the paper, a key fact under non-stationary stochastic volatility is that the distribution of bootstrap statistics (conditional on the data), rather than converging to the unconditional distribution of the statistic of interest, converges weakly to a *random* limit. By this we mean that the distribution function of the bootstrap statistic (conditional on the data) is stochastic not only for finite sample sizes n , but also in the limit as $n \rightarrow \infty$. Consequently, the conventional approach, based on the notion of weak convergence in probability of the bootstrap statistic to the limiting distribution of the original statistic (which is obviously non-stochastic), fails to deliver the required result of validity of the bootstrap. This problem is not new in the bootstrap literature, as it appears in various areas of application of the bootstrap; for example, in models with infinite variance innovations (Knight, 1989) and in autoregressive models with unit roots (Basawa, Mallik, McCormick, Reeves and Taylor, 1991; Cavaliere, Nielsen and Rahbek, 2015).

Specifically, in this paper we analyze the wild bootstrap under (non-stationary) stochastic volatility by adopting a new approach to assess bootstrap validity under random limit bootstrap measures. Thus, rather than focusing on the usual weak convergence in probability of the bootstrap conditional distribution, we make use of the concept of weak convergence in distribution (see Cavaliere and Georgiev, 2020, for a general introduction) to develop novel conditions for validity of the wild bootstrap, conditional on the volatility process. This allows us to establish that, although the presence of a *random* limiting distribution for the bootstrap statistic makes the bootstrap unable to estimate the unconditional distribution of the statistic of interest, the bootstrap can still deliver hypothesis tests with the desired size. In particular, we do this by establishing that the high-level conditions for bootstrap validity in Cavaliere and Georgiev (2020) can be shown to hold for a large class of models with stochastic volatility, including the aforementioned near-integrated GARCH model and the non-stationary stochastic volatility model. We do so by showing new weak convergence results conditional on volatility paths.

To illustrate our new approach and its applicability, we apply our results to three leading testing problems in the presence of non-stationary stochastic volatility, including testing a hypothesis on the location of a time series, testing for a unit root and testing for stability of the conditional mean using CUSUM-type statistics. These illustrative examples can easily be extended to cover more general cases, such as cointegration (as in Cavaliere *et al.*, 2010a, Cavaliere *et al.*, 2015 and Boswijk *et al.*, 2016) with multivariate stochastic volatility, or multivariate stability tests (see Perron, 2006 and Casini and Perron, 2019). Importantly, for all examples we show that conditions for conditional wild bootstrap validity include the absence of statistical leverage effects, i.e. correlation between the error process and its future conditional variance. The results of the paper are illustrated using Monte Carlo simulations, which indicate that under the conditions developed in our paper, the wild bootstrap leads to excellent size control even in small samples.

1.1 Structure of the paper

The structure of the paper is the following. In Section 2 we introduce the reference data generating process and our main assumptions, in particular on the volatility. Here we also introduce three examples which are used throughout the paper to illustrate the main results. We also derive the reference limit distribution for non-bootstrap statistics under non-stationary volatility. In Section 3 we introduce the main (wild) bootstrap algorithm. We show that when volatility is non-stationary, the bootstrap fails to mimic the asymptotic distribution of the corresponding statistics, and hence it is not valid in the usual sense. In Section 4 we discuss the wild bootstrap and prove, under proper assumptions, validity conditionally on the volatility path, as well as consistency of the bootstrap tests under the alternative. We first introduce in Section 4.1 the concept of weak convergence in distribution and discuss how to prove validity of the bootstrap in the presence of random limit bootstrap distributions, as it happens here under non-stationary volatility. Then in Section 4.2 we provide our main results under the required additional conditions on the original data. In Section 4.3 we apply our results about validity of the bootstrap in our three applications. Finally, in Section 4.4 we discuss consistency of the bootstrap tests under the alternative hypothesis. Results from a Monte Carlo study on the finite sample behavior of the bootstrap tests are reported in Section 5. Section 6 concludes. All proofs are reported in the online appendix.

1.2 Notation

The following (standard) notation is used throughout. With $x := y$ ($y =: x$) we mean that x is defined by y (y defined by x). For any $q \in \mathbb{R}$ (\mathbb{R} denoting the set of real numbers), $\lfloor q \rfloor$ denotes the integer part of q . For random elements (X_n, X) of a metric space, weak convergence of X_n to X is denoted by $X_n \xrightarrow{w} X$. Also, $\stackrel{d}{=}$ denotes equality in distribution. We use P^* , E^* and V^* respectively to denote probability, expectation and variance, conditional on the original sample. With $\xrightarrow{w^*}_p$ we denote weak convergence in probability; in particular, for random variables (X_n^*, X) , the notation $X_n^* \xrightarrow{w^*}_p X$ means that, as the sample size n diverges, the cumulative distribution function [cdf] G_n^* of X_n^* , conditional on the original data, converges in probability to the cdf G of X , at all continuity points of G . For a given sequence X_n^* of random elements, computed from the bootstrap data, $X_n^* - X = o_p^*(1)$, in probability, or $X_n^* \xrightarrow{p^*}_p X$, means that for any $\epsilon > 0$, $P^*(\|X_n^* - X\| > \epsilon) \xrightarrow{p} 0$, as $n \rightarrow \infty$. Similarly, $X_n^* = O_p^*(1)$, in probability, means that, for every $\epsilon > 0$, there exists a constant $M > 0$ such that, for all large n , $P^*(\|X_n^*\| > M) < \epsilon$ is arbitrarily close to one. Weak convergence in distribution and the related notation are introduced in Section 4.1. The Skorokhod spaces of càdlàg functions $[0, 1] \rightarrow \mathbb{R}^{m \times n}$ and $[0, 1] \rightarrow \mathbb{R}^n$ are denoted by $\mathcal{D}_{m \times n}[0, 1]$ and $\mathcal{D}_n[0, 1]$, respectively; for the latter, when $n = 1$ the subscript is suppressed. The Skorokhod space of càdlàg functions $\mathbb{R} \rightarrow \mathbb{R}$ is denoted by $\mathcal{D}(\mathbb{R})$.

2 Set-up and preliminaries

In this section we introduce our reference class of models for the conditional mean under stochastic volatility as well as the (test) statistics of interest. In Section 2.1 we focus on statistics which can be expressed (at least when the associated null hypothesis holds true) as functionals of the partial sum of the innovations and of the partial sum of the squared innovations. To illustrate ideas, we consider three simple univariate cases (which can easily be extended to multivariate cases) throughout: (i) testing a hypothesis on the mean in a simple location model; (ii) CUSUM testing for parameter constancy in a location model; (iii) testing for an autoregressive unit root in an AR(1) model. The main assumption on the volatility — which, *inter alia*, allows for non-stationary stochastic volatility or near-integrated GARCH dynamics — is discussed next in Section 2.2. Under the assumptions in Sections 2.1 and 2.2, the asymptotic (null) distributions can be derived. We do this in Section 2.3, where we show that the limiting distribution can be expressed in terms of a continuous martingale and its quadratic variation process. The implications of these results on bootstrap inference and hypothesis testing are the focus of the main Sections 3 and 4.

2.1 Model and hypotheses of interest

We are concerned with inference and hypothesis testing on the regression parameters of a heteroskedastic time series regression model in a triangular array form:

$$y_{n,t} = \beta' x_{n,t} + \varepsilon_{n,t}, \quad t = 1, \dots, n; \quad n = 1, 2, \dots \quad (2)$$

where $\varepsilon_{n,t}$ is a martingale difference sequence (mds) relative to a suitable filtration $\mathcal{F}_{n,t}$, with conditional variance $\sigma_{n,t}^2 = E(\varepsilon_{n,t}^2 | \mathcal{F}_{n,t-1})$. To simplify notation, unless strictly required we simply write (2) as $y_t = \beta' x_t + \varepsilon_t$, with $\sigma_t^2 := E(\varepsilon_t^2 | \mathcal{F}_{t-1})$.

Inference focuses on test statistics, which we assume can be expressed (at least under the null hypothesis) as functionals of partial sum processes in terms of the innovations and squared innovations

$$(M_n(u), U_n(u)) := \left(n^{-1/2} \sum_{t=1}^{\lfloor nu \rfloor} \varepsilon_t, n^{-1} \sum_{t=1}^{\lfloor nu \rfloor} \varepsilon_t^2 \right), \quad u \in [0, 1]. \quad (3)$$

as is the case for many testing problems, see also the discussion and examples below.

Defining $\mathcal{F}_n(u) := \mathcal{F}_{\lfloor nu \rfloor}$, the mds assumption implies that $\{M_n(u), \mathcal{F}_n(u)\}_{u \in [0,1]}$ is a martingale for all n , and $U_n(u)$ is its quadratic variation process, i.e.,

$$U_n(u) = [M_n](u) = \sum_{t=1}^{\lfloor nu \rfloor} \left(M_n \left(\frac{t}{n} \right) - M_n \left(\frac{t-1}{n} \right) \right)^2, \quad u \in [0, 1]. \quad (4)$$

Throughout it will also be useful to define the predictable quadratic variation or angle bracket process (see Jacod and Shiryaev, 2003):

$$V_n(u) := \langle M_n \rangle(u) = n^{-1} \sum_{t=1}^{\lfloor nu \rfloor} \sigma_t^2, \quad u \in [0, 1], \quad (5)$$

with the defining property that $\{M_n^2(u) - \langle M_n \rangle(u), \mathcal{F}_n(u)\}_{u \in [0,1]}$ is a martingale.

The following three testing problems are discussed in the paper. These are all special cases of (2) where the statistic of interest is indeed a functional of (M_n, U_n) .

Example 1 (Testing in a location model). Consider the location model $y_t = \theta + \varepsilon_t$, which is trivially obtained from (2) by setting $\beta = \theta$ and $x_t = 1$. The true location parameter is denoted by θ_0 . Suppose that interest is in testing the simple null hypothesis $\theta = \bar{\theta}$. Then, one can consider the test statistic $S_n := \sqrt{n}(\bar{y}_n - \bar{\theta})$, where $\bar{y}_n = n^{-1} \sum_{t=1}^n y_t$, or, alternatively, its studentized version $T_n := \sqrt{n}(\bar{y}_n - \bar{\theta})/s_n$, with $s_n^2 = n^{-1} \sum_{t=1}^n (y_t - \bar{y}_n)^2$. It is not difficult to see that, under the null hypothesis, it holds that S_n and T_n can be expressed in terms of M_n and U_n defined in (3) as

$$S_n = \sqrt{n}(\bar{y}_n - \bar{\theta}) = M_n(1), \quad T_n = \sqrt{n} \frac{(\bar{y}_n - \bar{\theta})}{s_n} = \frac{M_n(1)}{\sqrt{U_n(1) - n^{-1} M_n(1)^2}}.$$

If s_n is constructed with the null imposed, i.e. $s_n^2 = n^{-1} \sum_{t=1}^n (y_t - \bar{\theta})^2$, then under the null hypothesis T_n simplifies to

$$T_n = \frac{M_n(1)}{\sqrt{U_n(1)}}.$$

Example 2 (CUSUM test in a location model). Consider the time-varying location model $y_t = \theta_t + \varepsilon_t$, and suppose that interest is in testing the null hypothesis of a constant location parameter, i.e. $H_0 : \theta_t = \theta_1, t = 2, \dots, n$. A standard CUSUM test can be constructed by considering the statistic (see e.g. Deng and Perron, 2008, and the references therein)

$$CS_n := \frac{1}{n^{1/2}} \max_{t=1, \dots, n} \left| \sum_{i=1}^t (y_i - \bar{y}_n) \right|,$$

or its studentized version,

$$CT_n := \frac{1}{s_n n^{1/2}} \max_{t=1, \dots, n} \left| \sum_{i=1}^t (y_i - \bar{y}_n) \right|,$$

which, as in Example 1, reduce to

$$CS_n = \sup_{u \in [0,1]} |M_n(u) - uM_n(1)|, \quad CT_n = \frac{\sup_{u \in [0,1]} |M_n(u) - uM_n(1)|}{\sqrt{U_n(1) - n^{-1}M_n^2(1)}}$$

under H_0 .

Example 3 (Testing for a unit root). Consider the first-order autoregression $y_t = (1 + \theta)y_{t-1} + \varepsilon_t$, with $y_0 = 0$ (which again follows from (2) by setting $\beta = 1 + \theta$ and $x_t = y_{t-1}$). A test of the unit root hypothesis $\theta = 0$ can be based on the Dickey-Fuller ‘coefficient’ statistic $R_n := n\hat{\theta}_n$, where $\hat{\theta}_n = \sum_{t=1}^n y_{t-1} \Delta y_t / \sum_{t=1}^n y_{t-1}^2$ is the least-squares estimator from the regression of Δy_t on y_{t-1} . Under the null hypothesis, $\hat{\theta}_n = \sum_{t=1}^n \varepsilon_t (\sum_{i=1}^{t-1} \varepsilon_i) / \sum_{t=1}^n (\sum_{i=1}^{t-1} \varepsilon_i)^2$ and the test statistic may be expressed as

$$R_n = \frac{\int_0^1 M_n(u) dM_n(u)}{\int_0^1 M_n^2(u) du} = \frac{\frac{1}{2} (M_n^2(1) - U_n(1))}{\int_0^1 M_n^2(u) du}.$$

If the test is based on the Dickey-Fuller ‘ratio’ statistic $W_n := \hat{\theta}_n (s_n / (\sum_{t=1}^n y_{t-1}^2))^{-1/2}$, where $s_n^2 := n^{-1} \sum_{t=1}^n (\Delta y_t - \hat{\theta} y_{t-1})^2$, then

$$W_n = \frac{\int_0^1 M_n(u) dM_n(u)}{\sqrt{\int_0^1 M_n^2(u) du}} \frac{1}{\sqrt{U_n(1) - n^{-1} (\int_0^1 M_n(u) dM_n(u))^2 / \int_0^1 M_n^2(u) du}}$$

under the null hypothesis.

Some remarks are in order.

Remark 2.1. Although for fixed n , U_n can be determined from M_n as seen in (4), it does not define a continuous function $h : \mathcal{D}[0,1] \rightarrow \mathcal{D}[0,1]$. Therefore, limit results for U_n cannot be obtained from weak convergence of M_n together with the continuous mapping theorem [CMT].

Joint weak convergence of (M_n, U_n) is required to obtain the asymptotic null distribution of the statistics T_n , CT_n , R_n and W_n . This is related to the well-known fact that weak convergence of $\int_0^1 M_n dM_n$ to the stochastic integral $\int_0^1 M dM$ does not follow from $M_n \xrightarrow{w} M$ and the CMT; see e.g. Chan and Wei (1988).

Remark 2.2. The above examples involve single-parameter models. In more general testing situations, such as testing for a unit root in higher-order autoregressive models, the statistic of interest may be written as a functional of (M_n, U_n) plus an asymptotically negligible term. The theory developed in this paper can be extended to cover such cases.

2.2 Non-stationary stochastic volatility

We now introduce our basic hypotheses on the dynamic behavior of the conditional volatility σ_t^2 of the shocks ε_t . More specifically, we will allow volatility to be a persistent stochastic process, with a stochastic volatility weak limit, as formulated in the next two assumptions. These are in the spirit of the seminal paper by Hansen (1995), who considers conditional variances driven by nearly-integrated autoregressive shocks, although we do not constrain the behavior of the conditional variance to be of the autoregressive type.

Assumption 1. *In (2), we have $\varepsilon_t = \sigma_t z_t$, where z_t is a martingale difference sequence relative to $\mathcal{F}_t = \sigma(\{z_i\}_{i=1}^t, \{\sigma_i\}_{i=1}^{t+1})$, satisfying $E(z_t^2 | \mathcal{F}_{t-1}) = 1$.*

Define now the $\mathcal{D}[0, 1]$ version of the partial sum of the z_t 's as $B_{z,n}(u) := n^{-1/2} \sum_{t=1}^{\lfloor nu \rfloor} z_t$, $u \in [0, 1]$, and $\sigma_n(u)$, $u \in [0, 1]$, the $\mathcal{D}[0, 1]$ version of σ_t , as:

$$\sigma_n(u) := \sigma_{\lfloor nu+1 \rfloor}, \text{ for } u \in [0, 1], \quad (6)$$

with $\sigma_n(1) := \sigma_n$. In the following, the process $\sigma_n(u)$, $u \in [0, 1]$, will be referred to as σ_n ; this should not be confused with σ_t for $t = n$, as also will be clear from the context where the notation is used.

Assumption 2. *As $n \rightarrow \infty$, $(\sigma_n(u), B_{z,n}(u)) \xrightarrow{w} (\sigma(u), B_z(u))$ in $\mathcal{D}_2[0, 1]$, where $\sigma \in \mathcal{D}[0, 1]$ is a stochastic process satisfying $\inf_{u \in [0, 1]} \sigma(u) > 0$ a.s., and B_z is a standard Brownian motion on $[0, 1]$.*

While the convergence of the partial sum $B_{z,n} \in \mathcal{D}[0, 1]$ is standard, the requirement on $\sigma_n \in \mathcal{D}[0, 1]$ is not. More specifically, this assumption requires the conditional variance process $\{\sigma_t^2\}_{t=1}^n$ to be persistent enough such that its behavior can be approximated by an element of the space of càdlàg functions $\mathcal{D}[0, 1]$. No higher moment conditions are required on $\{\sigma_t^2\}_{t=1}^n$ or $\{\varepsilon_t\}_{t=1}^n$. Some examples of processes satisfying Assumptions 1 and 2 are presented next. These

will be analyzed in detail throughout the paper; for additional cases and discussions see e.g. Cavaliere and Taylor (2009).

Example V1 (Stochastic volatility). Let σ_t^2 be generated by

$$\log \sigma_t^2 = \phi_n \log \sigma_{t-1}^2 + (1 - \phi_n) \log \bar{\sigma}^2 + n^{-1/2} \eta_{t-1}, \quad t = 1, 2, \dots$$

where $\sigma_0^2 = \bar{\sigma}^2$ for some $\bar{\sigma} > 0$, $\phi_n = e^{-\kappa/n}$ for some $\kappa \geq 0$, and $\eta_t \sim \text{i.i.d. } N(0, \sigma_\eta^2)$, independent of $z_t \sim \text{i.i.d. } N(0, 1)$. Then $(\sigma_n, B_{z,n}) \xrightarrow{w} (\sigma, B_z)$ in $\mathcal{D}_2[0, 1]$, where:

$$d \log \sigma^2(u) = \kappa(\log \sigma^2(u) - \log \bar{\sigma}^2) du + \sigma_\eta dB_\eta(s), \quad u \in [0, 1],$$

with $\sigma^2(0) = \bar{\sigma}^2$, and where (B_η, B_z) is a bivariate standard Brownian motion.

Example V2 (Near-integrated GARCH). Consider the case where σ_t^2 is generated by the standard GARCH recursion

$$\sigma_t^2 = \omega_n + \alpha_n \varepsilon_{t-1}^2 + \beta_n \sigma_{t-1}^2 = \omega_n + \alpha_n \sigma_{t-1}^2 z_{t-1}^2 + \beta_n \sigma_{t-1}^2, \quad t = 1, 2, \dots$$

where $\sigma_0^2 = \bar{\sigma}^2$ for some $\bar{\sigma} > 0$, $\alpha_n + \beta_n = 1 - n^{-1}\kappa$ for some $\kappa \geq 0$, $\alpha_n = (2n)^{-1/2}\sigma_\eta$ for some $\sigma_\eta > 0$, $\omega_n = n^{-1}\bar{\sigma}^2\kappa$ and $z_t \sim \text{i.i.d. } N(0, 1)$. Then it follows from Nelson (1990) that $(\sigma_n, B_{z,n}) \xrightarrow{w} (\sigma, B_z)$ in $\mathcal{D}_2[0, 1]$, where

$$d\sigma^2(u) = \kappa(\sigma^2(u) - \bar{\sigma}^2) du + \sigma_\eta \sigma^2(u) dB_\eta(u), \quad u \in [0, 1],$$

with $\sigma^2(0) = \bar{\sigma}^2$, and where (B_η, B_z) is a bivariate standard Brownian motion. \square

Remark 2.3. In both examples, the process generating σ_t^2 depends on the sample size n , so that $\{\sigma_{nt}^2\}_{1 \leq t \leq n; n \geq 1}$ is actually a triangular array. We will not make this explicit in the notation in this section.

Remark 2.4. The initial condition σ_0^2 equals $\bar{\sigma}^2$ in both examples; this value represents the unconditional variance in the GARCH model, and $\log \bar{\sigma}^2$ is the (unconditionally) expected log-variance in the stochastic volatility model, in both cases assuming $\kappa > 0$. However, this has only been assumed for simplicity: the examples can be generalized to allow for an arbitrary initial condition, fixed or random, as long as it satisfies $\sigma_0^2 \xrightarrow{w} \sigma^2(0)$ for some strictly positive random variable $\sigma^2(0)$, independent of B_η .

Remark 2.5. The main difference between the examples is that in Example V1, the volatility shocks $\{\eta_t\}_{t \geq 1}$ are independent of $\{z_t\}_{t \geq 1}$, whereas in Example V2, the variance is driven by $\eta_t = (z_t^2 - 1)/\sqrt{2}$, which is fully determined by (although uncorrelated with) z_t . In both examples, however, the limit volatility process $\sigma(\cdot)$ is independent of the Brownian motion B_z generated by $\{z_t\}$. If we think of ε_t as the deviation of a financial return from its conditional

expectation, and $\{\sigma_t\}$ as its conditional volatility, then this rules out so-called leverage effects, i.e., asymmetric effects of positive and negative return shocks ε_t on future volatility σ_{t+h} , $h > 0$. Although the results given in this section also apply to processes with leverage, we will assume (asymptotic) independence in order to establish bootstrap validity.

Remark 2.6. In Example V1, the log-volatility follows a near-integrated first-order autoregression, converging weakly to an Ornstein-Uhlenbeck [OU] process. Similarly, in Hansen (1995) $\{\sigma_t\}$ satisfies Assumption 2 with $\sigma(\cdot)$ a (possibly nonlinear) transformation of an OU process (or Brownian motion). Cases where the volatility is allowed to jump at a countable number of times (while being constant between these jump times) are also allowed by our assumption. For instance, let $\sigma_t = \exp(\omega_0 + \omega_1 J_t)$, $J_t := \sum_{i=1}^t \delta_i \eta_i$, $J_0 = 0$ a.s., where for all t , δ_t is a Bernoulli random variable which equals one if and only if a volatility jump occurs at time t . If the η_t 's (which denote the random jump sizes) are i.i.d., independent of δ_t 's, and if $P(\delta_t = 1) = \lambda n^{-1}$, then in $\mathcal{D}[0, 1]$ (see Georgiev, 2008) $J_n(u) := J_{\lfloor nu \rfloor}$ converges weakly to the compound Poisson process $C_\lambda(u) := \sum_{i=1}^{N(u)} \eta_i$, where N is a Poisson process in $\mathcal{D}[0, 1]$ with intensity parameter λ . As expected, the limiting volatility process is $\sigma(u) = \exp(\omega_0 + \omega_1 C_\lambda(u))$, a piecewise constant process with number of discontinuities given by $N(1)$.

2.3 Standard asymptotics under non-stationary stochastic volatility

Assumptions 1 and 2 allow to analyze the asymptotic behavior of the functional (M_n, U_n) , as is done in the following Lemma.

Lemma 1. *Under Assumptions 1 and 2, we have in $\mathcal{D}_2[0, 1]$ as $n \rightarrow \infty$,*

$$(M_n(u), U_n(u)) \xrightarrow{w} (M(u), V(u)) := \left(\int_0^u \sigma(s) dB_z(s), \int_0^u \sigma^2(s) ds \right), \quad u \in [0, 1], \quad (7)$$

with $V(u) = \langle M \rangle(u)$. Furthermore,

$$\sup_{u \in [0, 1]} |U_n(u) - V(u)| \xrightarrow{p} 0. \quad (8)$$

The implications for the testing problems in Section 2.1 are given next.

Example 1 (cont'd). Consider the location model of Example 1. A straightforward application of Lemma 1 along with the CMT yields that, under H_0 , $S_n = M_n(1) \xrightarrow{w} M(1)$, which corresponds to the mixed normal distribution $N(0, \int_0^1 \sigma^2(u) du)$ and hence is non-pivotal. For T_n it holds that

$$T_n = \frac{M_n(1)}{\sqrt{U_n(1)}} + o_p(1) \xrightarrow{w} \frac{M(1)}{\sqrt{V(1)}} = \frac{\int_0^1 \sigma(u) dB_z(u)}{\sqrt{\int_0^1 \sigma^2(u) du}}. \quad (9)$$

Notice that the limit distribution in (9) is non-pivotal in cases where σ and B_z are not stochastically independent. In contrast, should independence hold, then (9) corresponds to a standard Gaussian distribution.

Example 2 (cont'd). For the CUSUM test statistics of Example 2 it holds that, again by Lemma 1 and the CMT, that under H_0

$$CS_n = \sup_{u \in [0,1]} |M_n(u) - uM_n(1)| \xrightarrow{w} \sup_{u \in [0,1]} |M(u) - uM(1)|,$$

$$CT_n = \frac{\sup_{u \in [0,1]} |M_n(u) - uM_n(1)|}{\sqrt{U_n(1) - n^{-1}M_n^2(1)}} \xrightarrow{w} \frac{\sup_{u \in [0,1]} |M(u) - uM(1)|}{\sqrt{V(1)}}.$$

Both statistics have a non-pivotal asymptotic null distribution, even in cases where the limit stochastic volatility process and the limit Brownian motions are stochastically independent.

Example 3 (cont'd). Finally, as shown in Cavaliere and Taylor (2009), for the unit root testing problem the presence of non-stationary volatility renders the null distribution of the Dickey-Fuller coefficient and t -statistics non-pivotal. More specifically, under the unit root null hypothesis it holds that

$$R_n = \frac{\frac{1}{2}(M_n^2(1) - U_n(1))}{\int_0^1 M_n^2(u) du} \xrightarrow{w} \frac{\frac{1}{2}(M(1)^2 - V(1))}{\int_0^1 M^2(u) du} = \frac{\int_0^1 M(u) dM(u)}{\int_0^1 M^2(u) du},$$

and

$$W_n = \frac{\int_0^1 M_n(u) dM_n(u)}{\sqrt{\int_0^1 M_n^2(u) du}} \frac{1}{\sqrt{U_n(1) + o_p(1)}} \xrightarrow{w} \frac{\int_0^1 M(u) dM(u)}{\sqrt{V(1) \int_0^1 M^2(u) du}}$$

as the sample size diverges.

3 Bootstrap under non-stationary stochastic volatility

Consider the standardized sample mean statistic S_n for the location model, see Example 1. Lemma 1 implies that under the null hypothesis $S_n \xrightarrow{w} M(1) = \int_0^1 \sigma(u) dB_z(u)$, see (7). The distribution of $M(1)$ depends on the limit volatility process σ , implying that critical values cannot be tabulated without providing a complete specification of this process. When the Brownian motion B_z and the limit volatility process σ are independent, then in the location model of Example 1 this problem can be avoided by considering the studentized test statistic T_n , which has a standard normal limit distribution under the null. However, in other testing problems (such as those in Examples 2 and 3) it is generally not possible to find such asymptotically pivotal statistics.

This motivates the development of bootstrap tests. Following much of the literature (e.g. Cavaliere and Taylor, 2008, 2009), we consider the wild bootstrap, which replicates the volatility

patterns in the original data. Let w_t^* be an i.i.d. sequence with mean zero and variance one², independent of $\{\sigma_t, z_t\}_{t \geq 1}$, and define the bootstrap shocks as

$$\varepsilon_t^* = \varepsilon_t w_t^*, \quad t = 1, 2, \dots, n.$$

Accordingly, we can define the bootstrap partial sum and the bootstrap partial sum of squares,

$$(M_n^*(u), U_n^*(u)) = \left(n^{-1/2} \sum_{t=1}^{\lfloor nu \rfloor} \varepsilon_t^*, n^{-1} \sum_{t=1}^{\lfloor nu \rfloor} (\varepsilon_t^*)^2 \right), \quad u \in [0, 1].$$

These processes are the bootstrap analogs of the processes M_n and U_n of Section 2. Notice that this implementation of the bootstrap assumes that ε_t is observed under the null hypothesis, which is the case in Examples 1 and 3 (location and unit root test). In more general testing problems, including Example 2 (CUSUM test), ε_t will be replaced by some residuals $\hat{\varepsilon}_t$ (either restricted by the null hypothesis or unrestricted).

3.1 Failure of classic bootstrap validity

Classic validity of the bootstrap (usually denoted as ‘bootstrap consistency’) is usually understood as the convergence in probability (or almost surely) of the conditional (on the original data) cdf of the bootstrap statistic to the limit cdf of the original statistic. We show here that, in the presence of stochastic volatility as in the previous section, in general classic validity of the bootstrap fails. This is essentially because the conditional cdf of the bootstrap statistic remains random in the limit. In this section we discuss this fact and its implications on bootstrap inference using, as the reference bootstrap algorithm, a wild bootstrap scheme as is typically applied when the data are heteroskedastic.

Focusing again on the location statistic S_n , its bootstrap counterpart is $S_n^* = n^{-1/2} \sum_{t=1}^n \varepsilon_t^* = M_n^*(1)$ where M_n^* is as previously defined. Define P^* as the bootstrap measure conditional on the original data, D_n . The bootstrap (conditional) cdf is

$$F_n^*(x) := P^*(S_n^* \leq x) = P(S_n^* \leq x | D_n).$$

The classical condition for bootstrap validity is that, as $n \rightarrow \infty$, $S_n^* \xrightarrow{w^*} P S := M(1)$. If the limit cdf $F(x) := P(S \leq x)$ is continuous, then this weak convergence in probability corresponds to the property

$$\sup_{x \in \mathbb{R}} |F_n^*(x) - F(x)| \xrightarrow{P} 0, \tag{10}$$

where F denotes the cdf of the asymptotic distribution of S_n :

$$F(x) := P(N(0, V(1)) \leq x) = \int \Phi(V(1)^{-1/2}x) dP(V(1)),$$

²Further conditions on the moments of w_t^* may be required in some specific applications.

where $\Phi(\cdot)$ is the standard normal cdf. Because $F(x)$ is the marginal cdf of $S = M(1)$, under independence of the processes B_z and σ it corresponds to the cdf of the mixed normal random variable $S = V(1)^{1/2}Z$, where $Z \sim N(0, 1)$, independent of $V(1)$.

However, under Assumption 2, condition (10) fails to hold, which is seen as follows. Choosing $w_t^* \sim$ i.i.d. $N(0, 1)$ for convenience, it is seen that

$$S_n^*|D_n \sim N(0, U_n(1))|U_n(1). \quad (11)$$

This follows because, conditional on the original data,

$$S_n^* = n^{-1/2} \sum_{t=1}^n \varepsilon_t w_t^* \sim N\left(0, n^{-1} \sum_{t=1}^n \varepsilon_t^2\right) \sim N(0, U_n(1)).$$

In terms of the conditional distribution F_n^* of S_n^* given the data D_n , (11) corresponds to

$$\begin{aligned} F_n^*(x) := P^*(S_n^* \leq x) &= P(N(0, U_n(1)) \leq x | U_n(1)) \\ &= P(N(0, 1) \leq U_n(1)^{-1/2} x | U_n(1)) = \Phi(U_n(1)^{-1/2} x). \end{aligned}$$

Letting $n \rightarrow \infty$, the limit distribution of the bootstrap statistic given the data follows from Lemma 1 and the CMT. Specifically, we have that

$$F_n^*(x) \xrightarrow{w} \Phi(V(1)^{-1/2} x) \quad (12)$$

in $\mathcal{D}(\mathbb{R})$; eq. (12) implies that the limit distribution of the bootstrap statistic is in fact *random*. That is, the bootstrap cdf F_n^* does not converge in probability but weakly, and the limiting cdf is random, as it depends on the random variable $V(1)$. Therefore, there is no reason to expect that the difference between the random function F_n^* and the non-random function F converges in probability to 0, as required for standard bootstrap validity to apply.

The fact that the conditional cdf F_n^* converges weakly (in $\mathcal{D}(\mathbb{R})$), rather than in probability, to a random cdf, will be referred to as ‘weak convergence in distribution’, and denoted as ‘ \xrightarrow{w} ’. More specifically, for sequences of random variables (Z_n, Y_n) and (Z, Y) (possibly defined on different probability spaces), the notation $Z_n|Y_n \xrightarrow{w} Z|Y$, when the conditional distribution of $Z|Y$ is diffuse (non-atomic), means that

$$F_n(\cdot|Y_n) := P(Z_n \leq \cdot | Y_n) \xrightarrow{w} P(Z \leq \cdot | Y) =: F(\cdot|Y), \quad \text{in } \mathcal{D}(\mathbb{R}). \quad (13)$$

A more general definition of $Z_n|Y_n \xrightarrow{w} Z|Y$, which includes non-diffuse conditional distributions and that is applicable to the case where (Z_n, Y_n) and (Z, Y) are random elements of the metric spaces $\mathcal{S}_Z \times \mathcal{S}_{Y_n}$ and $\mathcal{S}_Z \times \mathcal{S}_Y$, respectively (and hence to stochastic processes), is that $E(g(Z_n)|Y_n) \xrightarrow{w} E(g(Z)|Y)$ for all bounded continuous functions $g : \mathcal{S}_Z \rightarrow \mathbb{R}$, see Kallenberg

(1997), Cavaliere and Georgiev (2020) and the references therein. We generalize this definition to multivariate (joint) convergence in Section 4.1 below and employ it to prove bootstrap validity conditionally on the volatility path.

When Z_n represents a bootstrap statistic and the conditioning set Y_n is the original data D_n , we use the notation $\overset{w^*}{\rightarrow}_w$. Hence, eq. (12) corresponds to the weak convergence in distribution

$$S_n^* \overset{w^*}{\rightarrow}_w N(0, V(1))|V(1).$$

Unless $V(1)$ is non-random (which is not the case under stochastic volatility as considered here), this convergence shows that the limit bootstrap measure is indeed a random measure. Hence, the bootstrap cannot be valid in the usual sense of weak convergence in probability of F_n^* to F .

3.2 Examples (continued)

The result in Section 3.1 applies to the other examples considered, except for the asymptotically pivotal statistic T_n .

Example 1 (cont'd). Consider the location model example and assume that the bootstrap data are generated as $y_t^* = \varepsilon_t^*$, with ε_t^* as defined above. The bootstrap test statistics are $S_n^* := \sqrt{n}\bar{\varepsilon}_n^*$ and $T_n^* := \sqrt{n}\bar{\varepsilon}_n^*/s_n^*$, $s_n^* = (n^{-1}\sum_{t=1}^n(\varepsilon_t^* - \bar{\varepsilon}_n^*)^2)^{1/2}$. Using the argument discussed above, we have that $S_n^* \overset{w^*}{\rightarrow}_w M^*(1)|V(1)$, where $M^*(u) := \int_0^u \sigma(s)dB_z^*(s)$ with B_z^* a standard Brownian motion, stochastically independent of σ , and $V(1) := \int_0^1 \sigma(u)^2 du$. The conditional distribution of $M^*(1)|V(1)$ is equal (a.s.) to the conditional normal distribution $N(0, V(1))|V(1)$. In contrast, for T_n^* it holds that

$$T_n^* \overset{w^*}{\rightarrow}_w \frac{M^*(1)}{\sqrt{V(1)}} \Big| V(1),$$

where the limit corresponds to a conditional $N(0, 1)$ distribution (independent of $V(1)$), following from the conditional $N(0, V(1))$ distribution of $M^*(1)|V(1)$. Since weak convergence of a conditional distribution to a non-random cdf corresponds to weak convergence in probability, in this special case $T_n^* \overset{w^*}{\rightarrow}_p N(0, 1)$.

Two facts are worth stressing. First, in the above representations of the limit conditional distribution of the bootstrap statistic, σ and B_z^* are independent, even if the original processes σ and B_z are not. This result stems from the assumption that the wild bootstrap shocks w_t^* are independent of the original data. Second, if σ and B_z are stochastically independent, then the conditional distributions of $M^*(1)|V(1)$ and $M(1)|V(1)$ are equal (a.s.). This distributional equality is crucial to determine validity of the bootstrap when the convergence of S_n^* and S_n can be shown to be joint (see Section 4.1).

Example 2 (cont'd). For the bootstrap CUSUM statistics, suppose that the bootstrap data are generated as $\varepsilon_t^* = \hat{\varepsilon}_t w_t^*$ where $\hat{\varepsilon}_t := y_t - \bar{y}_n$, such that when the null hypothesis is true $\hat{\varepsilon}_t = \varepsilon_t - \bar{\varepsilon}_n$. The bootstrap statistics are defined as

$$CS_n^* = \sup_{u \in [0,1]} |M_n^*(u) - uM_n^*(1)|, \quad CT_n^* = \frac{\sup_{u \in [0,1]} |M_n^*(u) - uM_n^*(1)|}{\sqrt{U_n^*(1) - n^{-1}M_n^*(1)^2}},$$

with $M_n^*(u) := n^{-1/2} \sum_{t=1}^{\lfloor nu \rfloor} \varepsilon_t^*$ as above and $U_n^*(1) := n^{-1} \sum_{t=1}^n (\varepsilon_t^*)^2$. Under the null hypothesis,

$$n^{-1/2} \sum_{t=1}^{\lfloor nu \rfloor} \varepsilon_t^* = n^{-1/2} \sum_{t=1}^{\lfloor nu \rfloor} \hat{\varepsilon}_t w_t^* = n^{-1/2} \sum_{t=1}^{\lfloor nu \rfloor} \varepsilon_t w_t^* + O_p^*(n^{-1/2}),$$

uniformly in $u \in [0, 1]$, and it holds that

$$CS_n^* \xrightarrow{w^*} \sup_{u \in [0,1]} |M^*(u) - uM^*(1)| \Big| \sigma,$$

$$CT_n^* \xrightarrow{w^*} \frac{\sup_{u \in [0,1]} |M^*(u) - uM^*(1)|}{\sqrt{V(1)}} \Big| \sigma,$$

where again $M^*(u) := \int_0^u \sigma(s) dB_z^*(s)$ with B_z^* a standard Brownian motion, stochastically independent of σ , and $V(1) := \int_0^1 \sigma(u)^2 du$. Both bootstrap statistics have a *random* non-pivotal asymptotic null distribution.

Example 3 (cont'd). Finally, consider the unit root example. To avoid the problems described in Basawa et al. (1991), the bootstrap data are generated with the unit root imposed, i.e. $y_t^* = y_{t-1}^* + \varepsilon_t^*$, with $y_0^* = 0$ and $\varepsilon_t^* := (\Delta y_t) w_t^*$; see e.g. Cavaliere and Taylor (2008). Under the null, clearly $\varepsilon_t^* := \varepsilon_t w_t^*$. As discussed earlier for the non-bootstrap case, we have that

$$R_n^* = \frac{\frac{1}{2} (M_n^{*2}(1) - U_n^*(1))}{\int_0^1 M_n^{*2}(u) du}, \quad (14)$$

and, up to a negligible term,

$$W_n^* = \frac{\frac{1}{2} (M_n^{*2}(1) - U_n^*(1))}{\sqrt{U_n^*(1) \int_0^1 M_n^{*2}(u) du}}.$$

In this case it holds that

$$R_n^* \xrightarrow{w^*} \frac{\int_0^1 M^*(u) dM^*(u)}{\int_0^1 M^{*2}(u) du} \Big| \sigma, \quad W_n^* \xrightarrow{w^*} \frac{\int_0^1 M^*(u) dM^*(u)}{\sqrt{V(1) \int_0^1 M^{*2}(u) du}} \Big| \sigma, \quad (15)$$

where $M^*(u) := \int_0^u \sigma(s) dB_z^*(s)$ with B_z^* a standard Brownian motion, stochastically independent of σ , and $V(u) := \int_0^u \sigma(s)^2 ds$. The asymptotic distributions in (15) are random, except in the special case where σ is non-stochastic.

Cavaliere and Georgiev (2020) provide a number of other examples where the bootstrap validity condition (10) fails for any non-random cdf F , and develop an alternative criterion for *conditional* bootstrap validity. We will apply this concept to the present situation and extend it to the analysis of consistency of the bootstrap tests in the next section.

4 Validity of the bootstrap

Despite the fact that under non-stationary stochastic volatility the bootstrap is unable to consistently estimate the limiting distribution of the original statistic, it can still be valid, in the sense that it delivers control over type one error probabilities as n diverges. This can be seen by focusing on the bootstrap p -value. Taking the statistic S_n and associated bootstrap analog S_n^* to illustrate, the bootstrap p -value is defined as

$$p_n^* := P(S_n^* \leq S_n | D_n) = F_n^*(x)|_{x=S_n},$$

where $F_n^*(\cdot)$ is the cdf of S_n^* , conditional on the data (this definition of the p -value assumes a left-tailed test, which will be assumed below unless indicated otherwise). As in Cavaliere and Georgiev (2020), we say that the bootstrap based on S_n, S_n^* is valid conditionally on the volatility process $\{\sigma_t\}_{t=1}^n$ if p_n^* is asymptotically $U(0, 1)$ distributed conditionally on $\{\sigma_t\}_{t=1}^n$, i.e.

$$P(p_n^* \leq q | \{\sigma_t\}_{t=1}^n) \xrightarrow{P} q, \quad q \in (0, 1). \quad (16)$$

If this is the case, even if (as shown in the previous section) the limiting conditional distribution of S_n^* is random, the bootstrap test is still correctly sized in large samples. In fact, proofs of validity in the form of (16) also imply³ that, unconditionally,

$$P(p_n^* \leq q) \rightarrow q, \quad q \in (0, 1).$$

In the next subsections we discuss a set of sufficient conditions for (16) to hold. These are new in the literature on bootstrapping conditional mean models when the volatility can be stochastic. First, in Section 4.1 we provide our strategy to assess bootstrap validity. Our main results are given in Section 4.2. Application to our examples are provided in Section 4.3. Finally, the behavior under the alternative is analyzed and applied to our examples in Section 4.4.

4.1 Weak convergence in distribution and bootstrap validity

In this section, we summarize the approach developed by Cavaliere and Georgiev (2020), applied here to establish conditional bootstrap validity in the presence of non-stationary stochastic volatility. Before turning to bootstrap validity, we introduce a multivariate version of the concept of weak convergence in distribution that we anticipated in Section 3.1.

³For fixed $q \in (0, 1)$, $Q_n := P(p_n^* \leq q | \{\sigma_t\}_{t=1}^n)$, is a random sequence in $[0, 1]$, and hence uniformly integrable. Result (16), i.e. $Q_n \xrightarrow{P} q$, together with uniform integrability of Q_n , implies L^1 convergence and hence $P(p_n^* \leq q) = E(Q_n) \rightarrow q$ (Kallenberg, 1997, Lemma 3.11).

Precisely, let $Z_n := (Z_n^{(1)}, Z_n^{(2)})$ and $Z := (Z^{(1)}, Z^{(2)})$ be random elements of the (complete and separable) metric space $\mathcal{S}_{Z^{(1)}} \times \mathcal{S}_{Z^{(2)}}$, and $Y_n := (Y_n^{(1)}, Y_n^{(2)})$ and $Y := (Y^{(1)}, Y^{(2)})$ be random elements of the (complete and separable) metric spaces $\mathcal{S}'_{Y^{(1)}} \times \mathcal{S}'_{Y^{(2)}}$ and $\mathcal{S}_{Y^{(1)}} \times \mathcal{S}_{Y^{(2)}}$, respectively. We say that

$$(Z_n^{(1)}|Y_n^{(1)}, Z_n^{(2)}|Y_n^{(2)}) \xrightarrow{w} (Z^{(1)}|Y^{(1)}, Z^{(2)}|Y^{(2)}) \quad (17)$$

jointly if, for all bounded continuous $g : \mathcal{S}_{Z^{(1)}} \rightarrow \mathbb{R}$ and $h : \mathcal{S}_{Z^{(2)}} \rightarrow \mathbb{R}$,

$$\left(E(g(Z_n^{(1)})|Y_n^{(1)}), E(h(Z_n^{(2)})|Y_n^{(2)}) \right) \xrightarrow{w} \left(E(g(Z^{(1)}|Y^{(1)}), E(h(Z^{(2)}|Y^{(2)})) \right);$$

see Cavaliere and Georgiev (2020). If $Z_n^{(i)}, Z^{(i)}, i = 1, 2$ are random variables and the conditional distributions $Z^{(1)}|Y^{(1)}, Z^{(2)}|Y^{(2)}$ are diffuse, then the above convergence is equivalent to

$$\begin{aligned} & \left(P(Z_n^{(1)} \leq \cdot | Y_n^{(1)}), P(Z_n^{(2)} \leq \cdot | Y_n^{(2)}) \right) \\ & \xrightarrow{w} \left(P(Z^{(1)} \leq \cdot | Y^{(1)}), P(Z^{(2)} \leq \cdot | Y^{(2)}) \right) \text{ in } \mathcal{D}(\mathbb{R}) \times \mathcal{D}(\mathbb{R}). \end{aligned} \quad (18)$$

Finally, if $Z^{(1)}$ and $Z^{(2)}$ are the same random element (and similiary for $Y^{(1)}$ and $Y^{(2)}$), then we use the following alternative notation for (17):

$$(Z_n^{(1)}|Y_n^{(1)}, Z_n^{(2)}|Y_n^{(2)}) \xrightarrow{w} (Z|Y, Z|Y) =: (1, 1)Z|Y,$$

and for $Z_n = (Z_n^{(1)}, Z_n^{(2)})$ a bivariate random variable, (18) becomes

$$\left(P(Z_n^{(1)} \leq \cdot | Y_n^{(1)}), P(Z_n^{(2)} \leq \cdot | Y_n^{(2)}) \right) \xrightarrow{w} ((P(Z \leq \cdot | Y), (P(Z \leq \cdot | Y)) \text{ in } \mathcal{D}(\mathbb{R}) \times \mathcal{D}(\mathbb{R}).$$

Given this definition, we turn the attention to a statistic τ_n which is a function of the data D_n , which in our general set-up may be represented by (M_n, U_n) ; that is, $\tau_n = \tau(M_n, U_n)$. Its bootstrap equivalent is $\tau_n^* = \tau(M_n^*, U_n^*)$, and we let $\tau = \tau(M, V)$, with (M, V) denoting the weak limit of (M_n, U_n) , see Section 2.3.

Recall that, with $\sigma_n \in \mathcal{D}[0, 1]$ the volatility process defined in (6), conditioning on σ_n is equivalent to conditioning on $\{\sigma_t\}_{t=1}^n$, and moreover that $\sigma \in \mathcal{D}[0, 1]$ is the weak limit of σ_n by Assumption 2. It then follows by Cavaliere and Georgiev (2020, Corollary 3.2) that if the condition

$$(\tau_n|\{\sigma_t\}_{t=1}^n, \tau_n^*|D_n) = (\tau_n|\sigma_n, \tau_n^*|D_n) \xrightarrow{w} (\tau|\sigma, \tau|\sigma) = (1, 1)\tau|\sigma \quad (19)$$

is satisfied, with the random cdf of $\tau|\sigma$ being sample-path continuous, then

$$\sup_{x \in \mathbb{R}} |P(\tau_n \leq x | \{\sigma_t\}_{t=1}^n) - P(\tau_n^* \leq x | D_n)| \xrightarrow{P} 0.$$

This means that the bootstrap consistently estimates the distribution of the original statistic conditional on the volatility process, which in turn implies that the bootstrap is conditionally valid, i.e.,

$$P(p_n^* \leq q | \{\sigma_t\}_{t=1}^n) \xrightarrow{p} q$$

for all $q \in (0, 1)$, where $p_n^* = P(\tau_n^* \leq \tau_n | D_n)$ is the bootstrap p -value. The key condition to verify is therefore the one given in (19), along with continuity of the limiting (random) cdf.

Because $\tau_n = \tau(M_n, U_n)$ with $(M_n, U_n) \in \mathcal{D}_2[0, 1]$, proving (19) involves proving conditional functional limit theorems. It is known, see Goggin (1994) and Crimaldi and Pratelli (2005), that joint weak convergence of e.g. $((M_n, U_n), \sigma_n)$ is not sufficient for conditional weak convergence. For example, Goggin (1994) shows that a sufficient condition is that (M_n, U_n) is independent of σ_n , or that a change of measure can be found (with weakly convergent Radon-Nikodym derivative) under which this independence holds. These conditions do not seem to be directly applicable to the present case.

Remark 4.1. An alternative approach to proving conditional limit theorems would be to strengthen the joint weak convergence of $(\sigma_n, B_{z,n})$ in Assumption 2 to \mathcal{G} -stable convergence, see Häusler and Luschgy (2015), with \mathcal{G} the σ -algebra generated by the limit volatility process $\sigma(\cdot)$. This approach would require σ_n and σ to be defined on a common probability space, which is a natural assumption in the analysis of bootstrapping realized volatility, see e.g. Dovonon, Gonçalves and Meddahi (2013); here σ_n is a discretization of σ , converging in probability to σ by in-fill asymptotics. However, it does not cover Example V2, where only weak convergence of σ_n to σ may be established, and the two processes are defined on different probability spaces.

In contrast, our approach to proving (19) involves Skorokhod's representation theorem, and in particular the version of Kallenberg (1997), see Corollary A.1 in the online appendix. This allows us to obtain limit results "as if" the conditioning element σ_n converges almost surely to σ . By restricting the dependence between (M_n, U_n) and σ_n , we may then fix a realization of the volatility process and prove an unconditional functional limit theorem for each convergent sequence of realizations (except on a set with measure zero).

4.2 Main results

Recall that the main assumption used to derive the limiting distribution of the original statistic and the limiting (conditional) distribution of the bootstrap statistic is that the errors form a mds with respect to the past information set. This condition, however, is not sufficient for conditional bootstrap validity, unless the volatility is deterministic or stationary. In the presence of non-stationary stochastic volatility, further conditions are required. A sufficient set

of conditions is provided in the following assumption.

Assumption 3. Define $\mathcal{G}_{nt} := \sigma(\{z_i\}_{i=1}^t, \{\sigma_i\}_{i=1}^n)$ and hence $\mathcal{G}_{n0} := \sigma(\{\sigma_i\}_{i=1}^n)$, and define $\psi_{nt}^2 := E(z_t^2 | \mathcal{G}_{n,t-1})$ and $v_{nt} := z_t / \psi_{nt}$. Then:

(a) for all n , $\{v_{nt}\}_{t=1}^n$ is independent of \mathcal{G}_{n0} , and $\{\psi_{nt}\}_{t=1}^n$ is \mathcal{G}_{n0} -measurable.

(b) $\{z_t, \mathcal{G}_{nt}\}_{1 \leq t \leq n, n \geq 1}$ is a martingale difference array (mda), satisfying for all $\epsilon > 0$:

$$n^{-1} \sum_{t=1}^n E\left(z_t^2 \mathbb{I}_{\{|z_t| > \sqrt{n}\epsilon}\} | \mathcal{G}_{n0}\right) \xrightarrow{p} 0. \quad (20)$$

A few remarks are in order.

Remark 4.2. If $\{z_t\}_{t \geq 1}$ is independent of $\{\sigma_t\}_{t \geq 1}$, as in Example V1, then Assumption 3 is trivially satisfied with $\psi_{nt} = 1$ and $v_{nt} = z_t$ (the Lindeberg condition (20) is implicitly assumed in Assumption 2, to guarantee $B_{z,n} \xrightarrow{w} B_z$). The dependence allowed by the assumption is needed to cover situations such as the GARCH process in Example V2, where for all $t < n$,

$$z_t^2 = \frac{\sigma_{t+1}^2 - \omega_n - \beta_n \sigma_t^2}{\alpha_n \sigma_t^2},$$

which is known given $\{\sigma_t\}_{t=1}^n$, such that $\psi_{nt} = |z_t|$ and hence $v_{nt} = z_t / |z_t| = \text{sgn}(z_t)$, which because of symmetry of the Gaussian distribution will indeed be independent of $|z_t|$ and hence $\{\sigma_t\}_{t=1}^n$.

Remark 4.3. The mda assumption $E(z_t | \mathcal{G}_{n,t-1}) = 0$ rules out leverage effects, such as implied by non-zero correlation between z_t and the volatility shocks η_t in the stochastic volatility model of Example V1. It may be possible to weaken this assumption for the results to follow, and allow for dependencies for finite n , as long as they vanish asymptotically (such that σ and B_z are independent). In the latter case it would be guaranteed that at least the unconditional validity property $P(p_n^* \leq q) \rightarrow q$ for $q \in (0, 1)$ holds for the bootstrap, by Theorem 3.1 of Cavaliere and Georgiev (2020).

Remark 4.4. Part (a) of Assumption 3 implies that we may recover

$$M_n(\cdot) = n^{-1/2} \sum_{t=1}^{\lfloor n \cdot \rfloor} \sigma_t \psi_{nt} v_{nt}, \quad U_n(\cdot) = n^{-1} \sum_{t=1}^{\lfloor n \cdot \rfloor} \sigma_t^2 \psi_{nt}^2 v_{nt}^2,$$

and $V_n(\cdot)$ from the two independent sequences $\{\sigma_t\}_{t=1}^n$ and $\{v_{nt}\}_{t=1}^n$. This independence facilitates the analysis of conditional distributions, as will be evident from the proof of Theorem 1.

The main result in this section is given in Theorem 1 and Corollary 1 below.

Theorem 1. *Under Assumptions 1–3, we have as $n \rightarrow \infty$,*

$$\left(\begin{array}{c} (M_n, U_n) | \sigma_n \\ (M_n^*, U_n^*) | D_n \end{array} \right) \xrightarrow{w} \binom{1}{1} (M, V) | \sigma.$$

The key result of Theorem 1 is that the bootstrap processes M_n^* and U_n^* , conditionally on the data, replicate in the limit the distribution of the original processes M_n and U_n , conditionally on the volatility process σ_n (or, equivalently, on $\{\sigma_t\}_{t=1}^n$). The implication of Theorem 1 on the behaviour of the bootstrap p -values is provided in the following corollary, which applies to a statistic $\tau_n = \tau(M_n, U_n)$ (which, under the null, converges weakly to $\tau = \tau(M, V)$) and its bootstrap equivalent $\tau_n^* = \tau(M_n^*, U_n^*)$.

Corollary 1. *Under the conditions of Theorem 1, the bootstrap is valid conditionally on $\{\sigma_t\}_{t=1}^n$, i.e. with $p_n^* := P^*(\tau_n^* \leq \tau_n)$,*

$$p_n^* | \{\sigma_t\}_{t=1}^n \xrightarrow{w} U(0, 1),$$

provided that the conditional distribution of $\tau = \tau(M, V)$ given σ is sample-path continuous and the function τ is itself continuous.

4.3 Examples revisited

In this section we check whether the conditions for bootstrap validity hold for the examples. We assume throughout that Assumptions 1 and 2, strengthened by 3, hold.

Example 1 (cont'd). As earlier, the bootstrap statistics are given by $S_n^* := \sqrt{n}\bar{\varepsilon}_n^*$ and $T_n^* := \sqrt{n}\bar{\varepsilon}_n^*/s_n^*$. Under the null hypothesis, the original statistics are given by $S_n^* = \sqrt{n}\bar{\varepsilon}_n$ and $T_n^* = \sqrt{n}\bar{\varepsilon}_n/\hat{s}_n$. The original (bootstrap) statistics obtain as a continuous transformation of (M_n, U_n) (of (M_n^*, U_n^*)). Hence by Theorem 1 and a version of the CMT (see Cavaliere and Georgiev, 2020, Theorem A.1) it holds that

$$(S_n | \{\sigma_t\}_{t=1}^n, S_n^* | D_n) = (S_n | \sigma_n, S_n^* | D_n) \xrightarrow{w} (1, 1)M(1) | \sigma,$$

with $M(1) = \int_0^1 \sigma(u)dB_z(u)$. With $V(1) = \int_0^1 \sigma(u)^2 du$, the cdf of $M(1) | \sigma$ is given by $\Phi(uV(1)^{-1/2})$, which is sample-path continuous with probability 1. Hence, by Corollary 1, the bootstrap is valid conditionally on the volatility path σ_n . For the studentized statistic it holds that

$$(T_n | \sigma_n, T_n^* | D_n) \xrightarrow{w} (1, 1)Z | \sigma \stackrel{a.s.}{=} (1, 1)Z, \quad (21)$$

where $Z \sim N(0, 1)$, independent of σ ; this implies that the bootstrap is conditionally valid. Notice also that, see the discussion in Section 4.1, (21) and continuity of the limiting Gaussian cdf Φ imply that $\sup_{x \in \mathbb{R}} |P^*(T_n^* \leq x) - \Phi(x)| = o_p(1)$. As $T_n \xrightarrow{w} Z$, the bootstrap test based

on the studentized statistics T_n, T_n^* is also valid in the classic sense, i.e. the distribution of the bootstrap statistic T_n^* conditional on the data consistently estimates the (unconditional) distribution of the original statistic T_n . The same type of result does not hold for the bootstrap based on S_n, S_n^* ; however, bootstrap conditional validity is guaranteed by Corollary 1.

Example 2 (cont'd). As for the previous example, since the CUSUM (bootstrap) statistics are continuous transformations of (M_n, U_n) (of (M_n^*, U_n^*)), from Theorem 1 and the CMT in Cavaliere and Georgiev (2020) we have that, for $\tau_S := \sup_{u \in [0,1]} |M(u) - uM(1)|$

$$(CS_n | \sigma_n, CS_n^* | D_n) \xrightarrow{w} (1, 1) \tau_S | \sigma;$$

similarly, for $\tau_T := V(1)^{-1/2} \sup_{u \in [0,1]} |M(u) - uM(1)|$,

$$(CT_n | \sigma_n, CT_n^* | D_n) \xrightarrow{w} (1, 1) \tau_T | \sigma.$$

Both conditional asymptotic distributions are continuous with probability 1. As discussed in Andrews (1997), this holds using the results in Lifshits (1982) because the limiting random distribution corresponds (up to an almost surely strictly positive term) to the supremum of a conditionally Gaussian process with conditional covariance function which is nonsingular almost surely. Hence, by Corollary 1 the bootstrap is valid conditionally on the volatility process σ_n .

Example 3 (cont'd). Finally, in the unit root example we have that, under the stated assumption and if the null hypothesis holds, with $\tau_R := (\int_0^1 M^2(u) du)^{-1} \int_0^1 M(u) dM(u)$,

$$(R_n | \sigma_n, R_n^* | D_n) \xrightarrow{w} (1, 1) \tau_R | \sigma.$$

Similarly, for the t ratio test, with $\tau_W := (\int_0^1 M^2(u) du)^{-1/2} \int_0^1 M(u) dM(u)$,

$$(W_n | \sigma_n, W_n^* | D_n) \xrightarrow{w} (1, 1) \tau_W | \sigma.$$

As proved in Lemma A.1 in the online appendix, the limiting conditional cdfs have almost surely continuous sample paths, and hence by Corollary 1, the bootstrap is valid conditionally on σ_n .

4.4 Power considerations

We now briefly discuss the behavior of the bootstrap tests under the alternative hypothesis when the stochastic volatility process induces randomness of the limiting distribution of the bootstrap statistic. As before, consider a left-sided test based on the statistic $\tau_n = \tau(M_n, U_n)$ and its bootstrap equivalent $\tau_n^* = \tau(M_n^*, U_n^*)$. Suppose that under the alternative the original statistic diverges, say to $-\infty$, while the bootstrap statistic satisfies

$$\tau_n^* \xrightarrow{w^*} \tau^* | \sigma \tag{22}$$

for some random variable τ^* . Then, the following lemma holds.

Lemma 2. *Suppose that (22) holds and that $\tau_n \xrightarrow{p} -\infty$ as $n \rightarrow \infty$. Then, with $p_n^* := P^*(\tau_n^* \leq \tau_n)$, it holds that $p_n^* \xrightarrow{p} 0$.*

Lemma 2 shows that the fact that the limit distribution of the bootstrap statistic is random and depends on the volatility path does not affect the consistency of the bootstrap test. Essentially, weak convergence in distribution of τ_n^* given the data implies that the bootstrap statistic is $O_p^*(1)$, in probability. If the original statistic diverges to $-\infty$, it then holds that the bootstrap test rejects with probability converging to 1.⁴ We now apply this result to the three leading examples.

Example 1 (cont'd). Consider the location model example, where the econometrician is interested in testing the simple null hypothesis $\theta = \bar{\theta}$ when $\bar{\theta} > \theta_0$, θ_0 being the true parameter value. A wild bootstrap with the null imposed generates bootstrap data as $\varepsilon_t^* := \hat{\varepsilon}_t w_t^*$ with w_t^* i.i.d. $N(0, 1)$ and $\hat{\varepsilon}_t := y_t - \bar{\theta} = \varepsilon_t + \delta$, $\delta := \theta_0 - \bar{\theta} < 0$. It follows that, conditionally on the data, $M_n^*(\cdot) := n^{-1/2} \sum_{t=1}^{\lfloor n \cdot \rfloor} \varepsilon_t^*$ is a zero-mean Gaussian process with independent increments and with conditional variance function $\hat{U}_n(\cdot)$, where

$$\hat{U}_n(u) := n^{-1} \sum_{t=1}^{\lfloor nu \rfloor} \hat{\varepsilon}_t^2 = U_n(u) + u\delta^2 + O_p(n^{-1/2}) \xrightarrow{w} V(u) + u\delta^2, \quad \text{in } \mathcal{D}[0, 1].$$

Hence, under Assumptions 1 and 2, we have, as $n \rightarrow \infty$, that⁵ $M_n^* \xrightarrow{w^*} \tilde{M}^* | \sigma$, where $\tilde{M}^*(u) := \int_0^u \tilde{\sigma}(s) dB_z^*(s)$, with $\tilde{\sigma}(s) := (\sigma(s)^2 + \delta^2)^{1/2}$. This implies that $S_n^* \xrightarrow{w^*} \tilde{M}^*(1) | \sigma$. As $S_n \rightarrow -\infty$ as $n \rightarrow \infty$, the conditions of Lemma 2 are satisfied and $p_n^* \xrightarrow{p} 0$. Consistency of the test based on T_n follows by standard arguments as $T_n^* \xrightarrow{w^*} N(0, 1)$.

Example 2 (cont'd). For the bootstrap CUSUM statistics, consider the alternative $\theta_t = \theta_1 + g(t/n)$, where $g : [0, 1] \rightarrow \mathbb{R}$ is an arbitrary function satisfying $0 < \int_0^1 g^2(u) du < \infty$, see Ploberger and Krämer (1992). The wild bootstrap partial sum process is $M_n^*(\cdot) := n^{-1/2} \sum_{t=1}^{\lfloor n \cdot \rfloor} \varepsilon_t^*$, a zero-mean Gaussian process with independent increments and with conditional variance function $\hat{U}_n(\cdot) := n^{-1} \sum_{t=1}^{\lfloor n \cdot \rfloor} \hat{\varepsilon}_t^2 = U_n(\cdot) + G_n(\cdot) + O_p(n^{-1/2})$, with

$$G_n(u) := n^{-1} \sum_{t=1}^{\lfloor nu \rfloor} \left(g(t/n) - n^{-1} \sum_{t=1}^n g(t/n) \right)^2 \rightarrow \int_0^u \left(g(s) - \int_0^1 g(r) dr \right)^2 ds =: G(u),$$

in $\mathcal{D}[0, 1]$. This implies that

$$M_n^*(u) \xrightarrow{w^*} \tilde{M}^*(u) := \int_0^u \tilde{\sigma}(s) dB_z^*(s) \Big| \sigma, \quad \text{in } \mathcal{D}[0, 1],$$

⁴Notice that a consistent right-sided test can be obtained by focusing on the bootstrap p -value $\tilde{p}_n^* := 1 - p_n^*$.

⁵The proof reduces to the standard characterization of weak convergence to a Gaussian process by means of a Skorokhod representation.

with $\tilde{\sigma}(s) := (\sigma(s)^2 + (g(s) - \int_0^1 g(r) dr)^2)^{1/2}$ and B_z^* a standard Brownian motion, stochastically independent of σ . The limiting distribution of the bootstrap statistic CS_n^* is then given by

$$CS_n^* \xrightarrow{w^*} \sup_{u \in [0,1]} |\tilde{M}^*(u) - u\tilde{M}^*(1)| \sigma.$$

In order to analyze the CT_n^* statistic, notice that its denominator satisfies, in probability,

$$U_n^*(1) - n^{-1}M_n^*(1)^2 = U_n^*(1) + O_p^*(n^{-1}) = \hat{U}_n(1) + O_p^*(n^{-1/2}) \xrightarrow{w} V(1) + G(1).$$

By Lemma A.2 in the online appendix, the weak convergence in distribution of the numerator CS_n^* , together with the weak convergence of the denominator $\sqrt{U_n^*(1) - n^{-1}M_n^*(1)^2}$, implies the pairwise weak convergence in distribution of the numerator and denominator of CT_n^* , which in turn implies

$$CT_n^* \xrightarrow{w^*} \frac{\sup_{u \in [0,1]} |\tilde{M}^*(u) - u\tilde{M}^*(1)|}{\sqrt{V(1) + G(1)}} \sigma.$$

As both S_n and T_n diverge under the alternative considered, Lemma 2 applies and for both tests $p_n^* \xrightarrow{p} 0$.

Example 3 (cont'd). Consider the unit root example with wild bootstrap shocks generated with the null hypothesis imposed, i.e. $\varepsilon_t^* := (\Delta y_t)w_t^*$. The bootstrap R_n^* statistic is given as in (14) with $M_n^*(u) := n^{-1/2} \sum_{t=1}^{\lfloor nu \rfloor} (\Delta y_t)w_t^*$ and $U_n^*(u) := n^{-1} \sum_{t=1}^{\lfloor nu \rfloor} (\Delta y_t)^2 w_t^{*2}$. Conditionally on the data, M_n^* is a zero-mean Gaussian process with independent increments and conditional variance function $\hat{U}_n(\cdot) := n^{-1} \sum_{t=1}^{\lfloor n\cdot \rfloor} (\Delta y_t)^2$. Under the alternative that $y_t = (1+\theta)y_{t-1} + \varepsilon_t$ with $\theta \in (-2, 0)$, Δy_t can be written as the linear process with exponentially decaying coefficients $\Delta y_t = \sum_{i=0}^{t-1} \psi_i \varepsilon_{t-i}$ with $\psi_0 = 1$ and $\psi_i = \theta(1+\theta)^{i-1}$, $i = 1, 2, \dots$. Hence, by standard decompositions for squared stationary autoregressions it holds that (the proof is reported in the online appendix)

$$\hat{U}_n(u) = \bar{\psi}^2 U_n(u) + o_p(1), \quad \bar{\psi}^2 := \sum_{i=0}^{\infty} \psi_i^2, \quad (23)$$

where the $o_p(1)$ term is uniform in $u \in [0, 1]$, which implies that $\hat{U}_n \xrightarrow{w} \bar{\psi}^2 V$. Hence, $M_n^* \xrightarrow{w^*} \bar{\psi} M^* | \sigma$. Finally, using the fact that $U_n^* = \hat{U}_n + o_p^*(1)$, in probability, and invoking again Lemma A.2 from the online appendix to imply pairwise weak convergence in distribution of the numerator and denominator, it holds that

$$R_n^* \xrightarrow{w^*} \frac{\frac{1}{2}(\bar{\psi}^2 M^{*2}(1) - \bar{\psi}^2 V(1))}{\bar{\psi}^2 \int_0^1 M^{*2}(u) du} \sigma = \frac{\frac{1}{2}(M^{*2}(1) - V(1))}{\int_0^1 M^{*2}(u) du} \sigma.$$

Hence, under the alternative the bootstrap replicates the null distribution of the original statistic conditional on the volatility process and consistency of the bootstrap test follows from Lemma 2. An identical result holds for the t -ratio test based on W_n .

Remark 4.5. While in this section we focused on asymptotic power against fixed alternatives, it is possible to extend our analysis to cover power against local alternatives. To illustrate, consider the test based on S_n for the hypothesis $H_0 : \theta = \bar{\theta}$ in the location model $y_t = \theta + \varepsilon_t$ (Example 1). Under a sequence of local alternatives of the form $H_n : \theta_n = \bar{\theta} + \delta_n$ with $\delta_n = n^{-1/2}c$, it is straightforward to show that $S_n = c + M_n(1)$, which converges weakly to $c + M(1)$ under Assumptions 1–2. For the bootstrap statistic, the results obtained above with $\delta_n \rightarrow 0$ imply $S_n^* \xrightarrow{w^*} M(1)^* | V(1) = V(1)^{1/2} Z^* | V(1)$, with Z^* being $N(0, 1)$ (independent of $V(1)$); hence, the bootstrap distribution under the local alternative is the same as under the null. Suppose now that Assumption 3 holds; then, by Theorem 1,

$$(S_n | \sigma_n, S_n^* | D_n) \xrightarrow{w} (c + M(1) | \sigma, M(1) | \sigma).$$

Hence the limiting cdf of $S_n | \{\sigma_t\}_{t=1}^n$ is given by $F_c(x) = \Phi((x - c)V(1)^{-1/2})$, which is continuous with probability 1, while the bootstrap cdf $F_n^*(x)$ converges weakly to $F^*(x) = \Phi(xV(1)^{-1/2})$. Then, by application of Theorem 3.3 in Cavaliere and Georgiev (2020) it holds that the power of the bootstrap test at the $100\alpha\%$ nominal level, conditionally on the volatility process, is given by

$$\begin{aligned} P(p_n^* \leq \alpha | \sigma_n) &= P(F_n^*(S_n) \leq \alpha | \sigma_n) \xrightarrow{w} F_c(F^{*-1}(\alpha)) \\ &= \Phi((V(1)^{1/2}\Phi^{-1}(\alpha) - c)V(1)^{-1/2}) \\ &= \Phi(\Phi^{-1}(\alpha) - cV(1)^{-1/2}). \end{aligned}$$

By construction, the local power function depends on c ; we observe that it also depends on $V(1)$, and hence is random in the limit. In more general testing problems, the conditional local power function will depend on the entire volatility process. In the next section, we illustrate, by Monte Carlo simulations, the dependence on c as well as on (the limit of) σ_n .

5 Numerical results

In this section we analyze finite sample size and power properties of bootstrap tests under non-stationary stochastic volatility using Monte Carlo simulations. To study the behavior of the tests from Examples 1–3 under the null hypothesis, we report the Monte Carlo (empirical) cdfs of bootstrap p -values, both unconditionally over all Monte Carlo replications and conditionally on specific simulated volatility paths. Following Cavaliere and Georgiev (2020), we report the results in the form of fan charts of the conditional cdfs, displayed together with the unconditional cdf and the theoretical cdf of the $U(0, 1)$ distribution. Similarly, we display conditional power curves of the tests under local alternatives in fan charts.

In all experiments, we draw observations $\{\varepsilon_t\}_{t=1}^n$ from the GARCH(1,1) process from Example V2, with

$$\omega_n = 1 - \alpha_n - \beta_n = n^{-1}\kappa, \quad \alpha_n = (2n)^{-1/2}\sigma_\eta,$$

corresponding to a limit process $\sigma^2(u)$ with unit unconditional variance $\bar{\sigma}^2 = 1$, mean-reversion parameter κ , and volatility-of-volatility parameter σ_η . The conditional variance sequence is initialized at the unconditional variance, i.e., $\sigma_1^2 = 1$. We set $\kappa = 5$ and $\sigma_\eta = \sqrt{10}$, corresponding to a rather persistent volatility process with a reasonable amount of short-run variability of the volatility, which we know from earlier simulation studies to lead to substantial size distortions in tests based on standard (constant-volatility) asymptotic critical values. We expect similar results from stochastic volatility processes (Example V1) with the same type of persistence and volatility-of-volatility properties. We report results for two sample sizes, $n \in \{100, 500\}$. The standardized errors z_t are drawn from three different distributions, discussed below.

For each distribution and sample size, we first simulate 100 different realizations of the volatility path $\{\sigma_t\}_{t=1}^n$. For each of these paths, we draw 50,000 replications from the conditional distribution of $\{\varepsilon_t = \sigma_t z_t\}_{t=1}^n$ given $\{\sigma_t\}_{t=1}^n$. As discussed in Remark 4.2, this is equivalent to drawing $v_{nt} = \text{sgn}(z_t)$ conditional on $\psi_{nt} = |z_t|$ for $t = 1, \dots, n-1$, and drawing z_n from its unconditional distribution (independent of $\{\sigma_t\}_{t=1}^n$). For each choice of the distribution of $\{z_t\}_{t=1}^n$, we can check the conditions of Assumption 3 for conditional validity of the bootstrap.

The first data-generating process, labelled DGP 1, is defined by $z_t \sim N(0, 1)$. In that case the conditional distribution of the signs v_{nt} is discrete uniform over $\{-1, 1\}$, independent of $|z_t|$. This in turn implies that the mda condition of Assumption 3 is satisfied (as well as the independence, measurability and Lindeberg conditions), such that the bootstrap is conditionally valid.

In DGP 2, z_t is drawn from the following mixed normal density

$$f(z) = \frac{1}{3}\phi(z; \mu_1, \sigma_1) + \frac{2}{3}\phi(z; \mu_2, \sigma_2), \quad (24)$$

where $\phi(z; \mu, \sigma)$ is the pdf of the $N(\mu, \sigma^2)$ distribution, and where $\mu_1 = -2a, \sigma_1 = a, \mu_2 = a, \sigma_2 = a\sqrt{2}$, with $a = \sqrt{3/11}$. This distribution was constructed by Meijer (2000) to be asymmetric but with skewness 0 (and with mean zero and unit variance). Because $v_{nt} = \text{sgn}(z_t)$ in this case has a conditional distribution depending on $\psi_{nt} = |z_t|$, with $P(v_{nt} = 1 | \psi_{nt}) = f(\psi_{nt}) / (f(\psi_{nt}) + f(-\psi_{nt})) \neq \frac{1}{2}$, it follows that Assumption 3 is violated, and conditional validity of the bootstrap is not guaranteed. On the other hand, the zero skewness implies that the limit result of Example V2 still applies, with B_z independent of B_η and hence σ . As conjectured in Remark 4.3, we may expect unconditional bootstrap validity in this case.

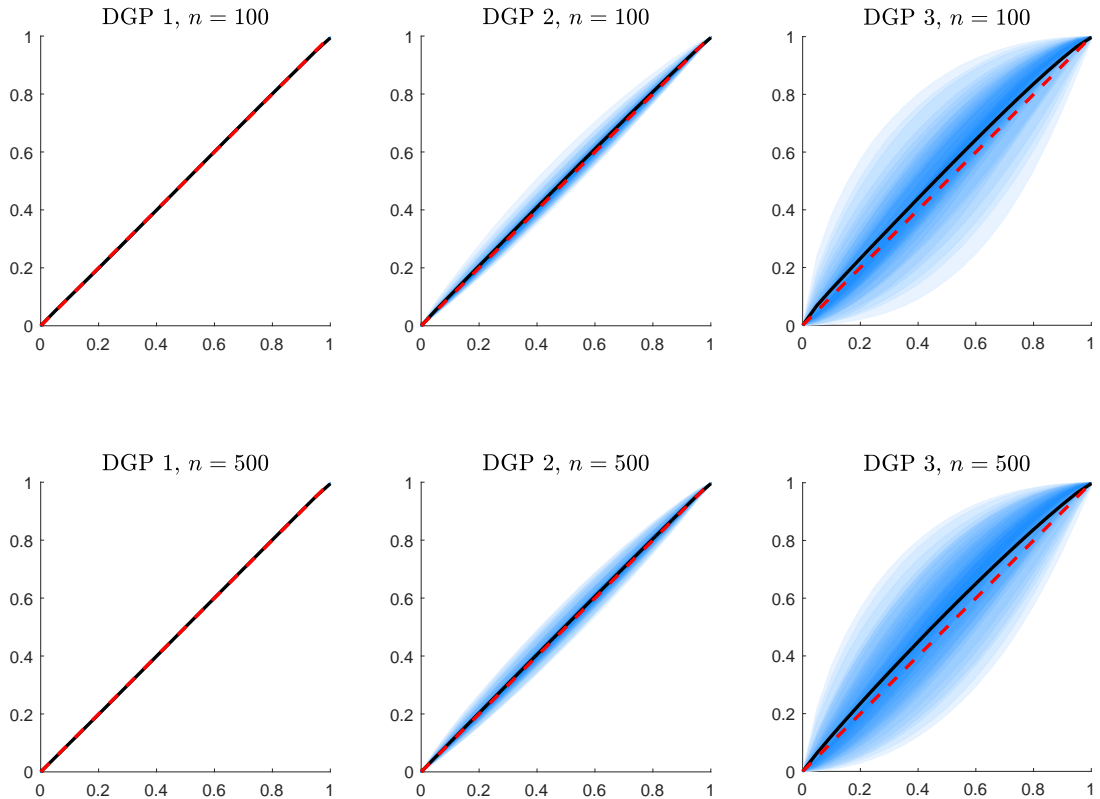


Figure 1: Monte Carlo conditional cdfs of bootstrap p -values of T_n

In DGP 3, z_t is drawn from another version of (24), but now with $\mu_1 = -2b, \sigma_1 = b\sqrt{2}, \mu_2 = b, \sigma_2 = b$, with $b = \sqrt{3/10}$. This is a distribution with mean zero, unit variance and negative skewness, so that B_z and B_η in Example V2 have a negative correlation, corresponding to long-run leverage effects. This implies that the wild bootstrap is invalid in this case, both conditionally and unconditionally.

Figures 1–3 display the results for the behavior of bootstrap p -values (based on 199 bootstrap replications) under the null hypothesis, for the studentized tests based on T_n , CT_n and W_n , respectively. Unreported results for the other three test statistics S_n , CS_n and R_n are very similar to the results for the corresponding studentized tests.

From Figure 1, we observe that when the standardized errors z_t are standard normal (DGP 1, left panels), then the conditional distribution of the bootstrap p -values is very close to uniform, and appears to be independent of σ for both sample sizes considered. Thus the theoretical conditional validity of the bootstrap in this case is clearly reflected in finite-sample behavior. When the distribution of z_t is asymmetric with zero skewness (DGP 2, centre panels), then the bootstrap appears to be valid on average (indicated by the solid line almost coinciding with the $U(0,1)$ cdf, especially for $n = 500$), but the conditional cdfs of bootstrap p -values do depend

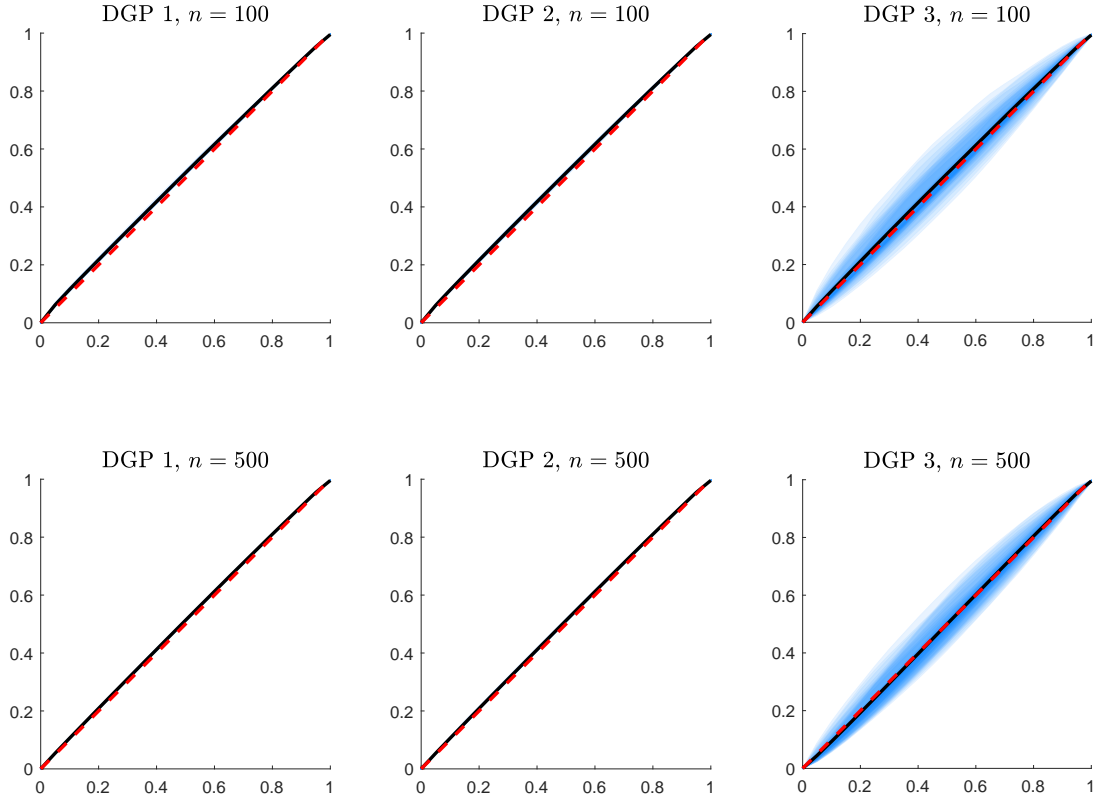


Figure 2: Monte Carlo conditional cdfs of bootstrap p -values of CT_n

on the volatility path and deviate from the uniform cdf, illustrating the conjectured violation of conditional bootstrap validity. Finally, for DGP 3 (right panels, skewed z_t), we observe more extreme dependence of bootstrap p -values on the volatility path. In this case the bootstrap does not appear to be valid on average either, as predicted by the dependence between B_z and σ implied by this DGP, which is not replicated by the wild bootstrap.

Figure 2 displays the results for the studentized CUSUM test based on CT_n . For this test, the finite-sample size distortion (indicated by the difference between the solid and dashed line) is more pronounced than for the location test, in particular for the smaller sample size ($n = 100$). Unreported additional simulations show that these size distortions are even stronger for the test based on CS_n . The results improve when the sample size increases, and it should be noted that the rejection frequencies at the 5% significance level are still fairly close to 0.05; the deviations are larger at the centre of the distribution. For this test, the dependence of the conditional cdf of p -values on σ is much weaker than for the location test. For DGP 2, we do not observe any deviation of conditional cdfs from their average; in case of DGP 3, there is a clear violation of conditional bootstrap validity, but the deviations are less pronounced than for T_n .

The results for the unit root test based on W_n are given in Figure 3. For DGP 1 and 2, the

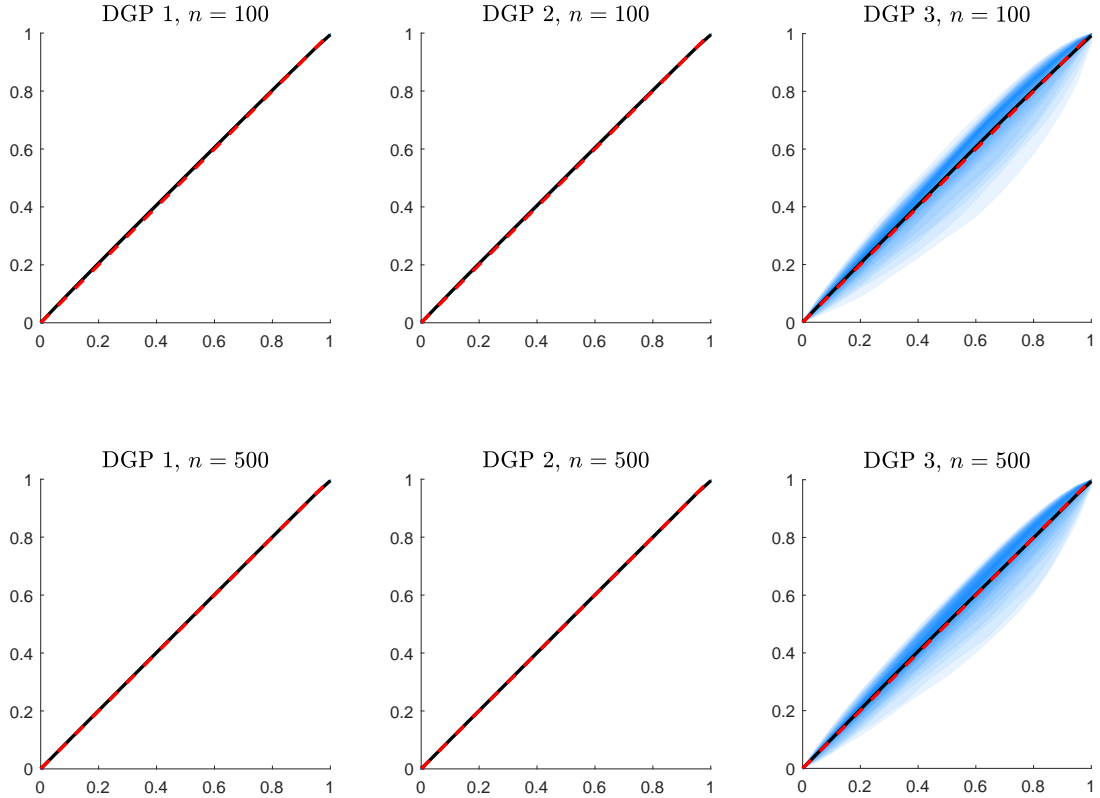


Figure 3: Monte Carlo conditional cdfs of bootstrap p -values of W_n

size distortions appear to be negligible for both sample sizes. Similarly to the CUSUM test, the dependence of bootstrap p -values on the volatility path for DGP 1 and 2 appears to be very weak. On the other hand, for DGP 3 we find this dependence to be clearly present, illustrating again a violation of conditional bootstrap validity. On average, the bootstrap appears to be valid even for DGP 3, although theoretically we would not expect this to be the case because of the dependence between B_z and σ . Unreported simulations have shown that in case of stronger leverage effects (i.e., a stronger correlation between B_z and B_η), the unconditional cdf of bootstrap p -values does differ from the $U(0, 1)$ cdf, as predicted by the theoretical results.

Next, we investigate the local power of the bootstrap tests, again conditional on the same realizations of the volatility path as considered for the size of the tests. As in Remark 4.5, for the location tests we evaluate the rejection frequency of the test for $H_0 : \theta = 0$ against local alternatives $\theta_n = -n^{-1/2}c$, with $c \in [0, 8]$. For the CUSUM tests, the local alternative is a break in the mean of the series, at $t = n/2$, from $\theta_{n,t} = 0$ to $\theta_{n,t} = n^{-1/2}c$, with $c \in [0, 15]$. For the unit root tests, we consider local alternatives $\theta_n = -n^{-1}c$, with $c \in [0, 20]$. We provide results for the tests based on the studentized statistics T_n , CT_n and W_n , and for the sample size $n = 100$.

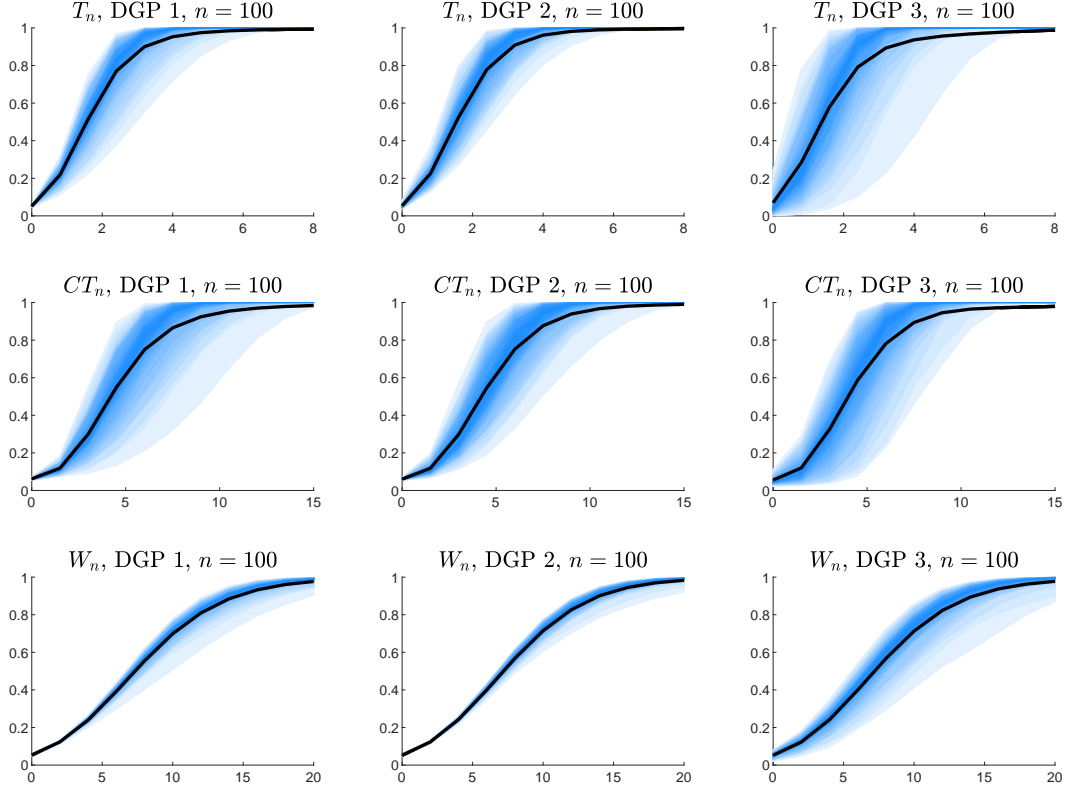


Figure 4: Monte Carlo conditional rejection frequencies of bootstrap tests

Figure 4 displays the rejection frequencies of the bootstrap tests, based on 10,000 replications of the test for each volatility path, plotted against c . We observe that the conditional rejection probabilities under the alternative hypothesis depend on both the non-centrality parameter c and the volatility process σ , for each test and DGP 1–3. While the dependence on the volatility is as expected for DGP 2 and 3, we note that for DGP 1, where the rejection probabilities under the null hypothesis are conditionally independent of the volatility process, the power of the tests clearly depends on the volatility (as discussed in Remark 4.5).

To gain some insight in the sensitivity of the simulation results to the chosen parameter values, we have repeated the Monte Carlo experiments with $\kappa = 0$, implying an integrated GARCH process ($\alpha_n + \beta_n = 1$) with infinite unconditional variance. As in the original simulations, the initial condition is $\sigma_1^2 = 1$, the volatility of volatility is $\sigma_\eta = \sqrt{10}$, and the same three DGPs for z_t have been used. The corresponding figures (available from the authors upon request) show qualitatively very similar results as those displayed in Figures 1–4. The increased persistence and variation in the realizations of $\{\sigma_t^2\}_{t=1}^n$ leads to a bigger bias in DGP 3 for the W_n test statistic, and a bit more variation in the conditional distributions for DGP 3 for the W_n and CT_n test statistics. Furthermore, we find more variation in conditional power functions and a

slightly higher average (unconditional) power.

6 Conclusions

In this paper we have analyzed the properties of wild bootstrap inference in time series models for the conditional mean under non-stationary stochastic volatility. In our setting, we do not make any specific assumption on the volatility process, other than assuming that it admits a weak limit in $\mathcal{D}[0, 1]$. The additional advantage of this semi-parametric nature of our analysis is that we do not need any assumptions on the existence of higher moments, or on the initial condition of the conditional variance process. On the other hand, if the parametric form of the volatility process was known to the econometrician, then other model-based estimators and tests (such as Seo (1999), for the unit root testing problem under GARCH(1,1) errors) and related bootstrap methods could be implemented. These approaches, although particularly interesting and likely to deliver power improvements over our wild bootstrap methods, are beyond the scope of this paper.

A central ingredient in our analysis is that the bootstrap distribution (the distribution of a bootstrap statistic conditional on the data) has a random limit. In our case, this random limit can be characterized as the conditional distribution of the limiting statistic conditional on a continuous-time volatility process (which is itself the weak limit of the finite-sample volatility process). The random nature of the limit of bootstrap distributions is not new: a well-known example in the unit-root literature is Basawa et al. (1991), and Sen, Banerjee and Woodroffe (2010) provide another example from non-parametric statistics involving cube-root asymptotics. Cavaliere and Georgiev (2020) analyze a number of other examples of this phenomenon. In the examples encountered so far, this random limit distribution can always be expressed as a conditional distribution, but its specific form is determined by a combination of model assumptions, statistical methods and bootstrap implementation details. Some of these combinations lead to bootstrap invalidity, which we may try to resolve by another bootstrap implementation. For example, the randomness of the bootstrap limit distribution in the unit-root testing problem, and the associated invalidity of the standard recursive-design bootstrap, is resolved by the use of restricted residuals (with the unit root imposed) in the bootstrap algorithm. The challenging part of our work is to show that the bootstrap distribution matches a particular conditional distribution of the original statistic, in our case conditional on the volatility process $\{\sigma_t\}_{t=1}^n$. If we were to change the bootstrap scheme (e.g., by using an i.i.d. or ‘ m out of n ’ bootstrap), then such a match would not occur, and the bootstrap would be invalid.

Our results can be generalized in several directions. First, our applications deal with uni-

variate time series models and it is naturally of interest to apply our results to multivariate (time series) models, where volatilities and correlations are time-varying, stochastic and non-stationary. In particular, in Boswijk *et al.* (2016) the bootstrap was considered for multivariate cointegrated vector autoregressions in the presence of stationary volatility, in combination with possible deterministic changes in the volatility; we conjecture that our results obtained here also apply to the case of non-stationary multivariate stochastic volatility. Second, it would be important to understand how to bootstrap conditional mean time series models in the presence of leverage. Although, as we have shown, the wild bootstrap is not valid in this context, our theory may be useful for assessing validity of other bootstrap methods when the volatility displays leverage effects.

Acknowledgements

We thank the guest Co-Editor and two anonymous referees for important suggestions on an earlier draft of the paper. We also thank participants of the conference in honour of Pierre Perron (Boston University, 14–15 March 2019), the 32nd European Meeting of Statisticians (Palermo, 22–26 July 2019), the 12th World Congress of the Econometric Society (Bocconi University, 17–21 August 2020), and of workshops and seminars in Amsterdam, Canterbury, Ílhavo, Nottingham, Oxford, Rotterdam and Vienna, for useful feedback on earlier versions of the paper. This research was supported by the Danish Council for Independent Research (DSF Grant 015-00028B), by the University of Bologna (ALMA IDEA 2017 Grant), and by the Italian Ministry of University and Research (PRIN 2017 Grant 2017TA7TYC).

References

- Andrews, D. W. K. (1997). A conditional Kolmogorov Test. *Econometrica* 65, 1097–1128.
- Andrews, D. W. K. (2000). Inconsistency of the bootstrap when a parameter is on the boundary of the parameter space. *Econometrica* 68, 399–405.
- Basawa, I., Mallik, A., McCormick, W., Reeves, J., Taylor, R. (1991). Bootstrapping unstable first-order autoregressive processes. *The Annals of Statistics* 19, 1098–1101.
- Boswijk, H. P., Cavaliere, G., Rahbek, A., Taylor, A. M. R. (2016). Inference on co-integration parameters in heteroskedastic vector autoregressions. *Journal of Econometrics* 192, 64–85.
- Casini, A., Perron, P. (2019). Structural breaks in time series. In: *Oxford Research Encyclopedia of Economics and Finance*. Oxford University Press, Oxford.

- Cavaliere, G., Georgiev, I. (2020). Inference under random limit bootstrap measures. *Econometrica* 88, 2547–2574.
- Cavaliere, G., Nielsen, H. B., Rahbek, A. (2015). Bootstrap testing of hypotheses on cointegration relations in VAR models. *Econometrica* 83, 813–831.
- Cavaliere, G., Nielsen, H. B., Rahbek, A. (2017). On the consistency of bootstrap testing for a parameter on the boundary of the parameter space. *Journal of Time Series Analysis* 38, 513–534.
- Cavaliere, G., Rahbek, A., Taylor, A. M. R. (2010a). Testing for co-integration in vector autoregressions with non-stationary volatility. *Journal of Econometrics* 158, 7–24.
- Cavaliere, G., Rahbek, A., Taylor, A. M. R. (2010b). Cointegration rank testing under conditional heteroskedasticity. *Econometric Theory* 26, 1719–1760.
- Cavaliere, G., Taylor, A. M. R. (2007). Testing for unit roots in time series models with non-stationary volatility. *Journal of Econometrics* 140, 919–947.
- Cavaliere, G., Taylor, A. M. R. (2008). Bootstrap unit root tests for time series with nonstationary volatility. *Econometric Theory* 24, 43–71.
- Cavaliere, G., Taylor, A. M. R. (2009). Heteroskedastic time series with a unit root. *Econometric Theory* 25, 1228–1276.
- Chan, N. H., Wei, C. Z. (1988). Limiting distributions of least squares estimates of unstable autoregressive processes. *Annals of Statistics* 16, 367–401.
- Crimaldi, I., Pratelli, L. (2005). Convergence results for conditional expectations. *Bernoulli* 11, 737–745.
- Deng, A., Perron, P. (2008). A non-local perspective on the power properties of the CUSUM and CUSUM of squares tests for structural change. *Journal of Econometrics* 142, 212–240.
- Dovonon, P., Gonçalves, S., Meddahi, N. (2013). Bootstrapping realized multivariate volatility measures. *Journal of Econometrics* 172, 49–65.
- Engle, R. F., Rangel, J. G. (2008). The spline-GARCH model for low-frequency volatility and its global macroeconomic causes. *The Review of Financial Studies* 21, 1187–1222.
- Georgiev, I. (2008) Asymptotics for cointegrated processes with infrequent stochastic level shifts and outliers. *Econometric Theory* 24, 587–615.
- Goggin, E. M. (1994). Convergence in distribution of conditional expectations. *Annals of Probability* 22, 1079–1114.
- Gonçalves, S., Kilian, L. (2004). Bootstrapping autoregressions with conditional heteroskedasticity of unknown form. *Journal of Econometrics* 123, 89–120.
- Gonçalves, S., Kilian, L. (2007). Asymptotic and bootstrap inference for $AR(\infty)$ processes with

- conditional heteroskedasticity. *Econometric Reviews* 26, 609–641.
- Hansen, B. E. (1995). Regression with nonstationary volatility. *Econometrica* 63, 1113–1132.
- Harvey, D., Leybourne, S. J., Sollis, R., Taylor, A. M. R. (2016). Tests for explosive financial bubbles in the presence of non-stationary volatility. *Journal of Empirical Finance* 38, 548–574.
- Häusler, E., Luschgy, H. (2015). *Stable Convergence and Stable Limit Theorems*. Springer, Cham.
- Jacod, J., Shiryaev, A. N. (2003). *Limit Theorems for Stochastic Processes* (2nd ed.). Springer, Berlin.
- Kallenberg, O. (1997). *Foundations of Modern Probability*. Springer, New York.
- Kim, C., Nelson, C. (1999). Has the U.S. economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle. *The Review of Economics and Statistics* 81, 608–616.
- Knight, K. (1989). On the bootstrap of the sample mean in the infinite variance case. *The Annals of Statistics* 17, 1168–1175.
- Lifshits, M. A. (1982). On the absolute continuity of distributions of functionals of random processes. *Theory of Probability and Its Applications* 27, 600–607.
- Loretan, M., Phillips, P. C. B. (1994). Testing the covariance stationarity of heavy-tailed time series: An overview of the theory with applications to several financial datasets. *Journal of Empirical Finance* 1, 211–248.
- McConnell, M. M., Perez-Quiros, G. (2000). Output fluctuations in the United States: What has changed since the early 1980's? *American Economic Review* 90, 1464–1476.
- Meijer, E. (2000). An asymmetric distribution with zero skewness. Working paper. <http://dx.doi.org/10.2139/ssrn.2531847>.
- Nelson, D. B. (1990). ARCH models as diffusion approximations. *Journal of Econometrics* 45, 7–38.
- Perron, P. (2006). Dealing with structural breaks. In: Patterson, K., Mills, T. C. (Eds.). *Palgrave Handbook of Econometrics, Vol. 1: Econometric Theory*. Palgrave Macmillan, pp. 278–352.
- Ploberger, W., Krämer, W. (1992). The CUSUM test with OLS residuals. *Econometrica* 60, 271–285.
- Sen, B., Banerjee, M., Woodroffe, M. (2010). Inconsistency of bootstrap: The Grenander estimator. *The Annals of Statistics* 38, 1953–1977.
- Sensier, M., van Dijk, D. (2004). Testing for volatility changes in U.S. macroeconomic time series. *The Review of Economics and Statistics* 86, 833–839.
- Seo, B. (1999). Distribution theory for unit root tests with conditional heteroskedasticity. *Jour-*

nal of Econometrics 91, 113–144.

Xu, K.-L., Phillips, P. C. B. (2008). Adaptive estimation of autoregressive models with time-varying variances. *Journal of Econometrics* 142, 265–280.