BOOTSTRAP TESTING OF HYPOTHESES ON CO-INTEGRATION RELATIONS IN VAR MODELS

By Giuseppe Cavaliere, Heino Bohn Nielsen, and Anders Rahbek*

Abstract

It is well known that the finite-sample properties of tests of hypotheses on the co-integrating vectors in vector autoregressive models can be quite poor, and that current solutions based on Bartlett-type corrections or bootstrap based on unrestricted parameter estimators are unsatisfactory, in particular in those cases where also asymptotic χ^2 tests fail most severely. In this paper we solve this inference problem by showing the novel result that a bootstrap test where the null hypothesis is imposed on the bootstrap sample is asymptotically valid. That is, not only does it have asymptotically correct size but, in contrast to what is claimed in existing literature, it is consistent under the alternative. Compared to the theory for bootstrap tests on the co-integration rank (Cavaliere, Rahbek and Taylor, 2012, *Econometrica,* 80, 1721-1740), establishing the validity of the bootstrap in the framework of hypotheses on the co-integrating vectors requires new theoretical developments, including the introduction of multivariate Ornstein-Uhlenbeck processes with random (reduced rank) drift parameters. Finally, as documented by Monte Carlo simulations, the bootstrap test outperforms existing methods.

Keywords: Co-integration; Hypotheses on co-integrating vectors; Bootstrap.

1 Introduction

A key part of co-integration analysis in vector autoregressive (VAR) models is, upon determination of the co-integration rank, to make inference on the co-integrating relations. As originally proposed in Johansen (1991) this can be done using likelihood ratio tests. However, as is well-known, the finite-sample properties of these tests can be quite poor, much like in the case of rank determination. Currently there are two dominating approaches to this in the literature: one is to apply Bartlett corrections as in Johansen (2000, 2002a) and the other to base inference on a bootstrap scheme where the VAR model is estimated unrestrictedly (i.e., without imposing the null hypothesis), as in Omtzigt and Fachin (2006). Unfortunately, both approaches tend to be unreliable in terms of size, in particular in those cases where also the asymptotic test fails most severely. We propose to solve this problem by showing the novel result that a bootstrap test where the null hypothesis is imposed on the bootstrap sample is asymptotically valid. That is, not only does it have asymptotically correct size in the parameter region

* We thank the co-editor and three anonymous referees for their very helpful and constructive comments on a previous version of this paper. This paper also benefited from comments received by conference and seminar participants at the Chinese University of Hong Kong, the Hong Kong University of Science and Technology, Monash University, Shanghai Jao Tong University, Singapore Management University, the University of Oxford and the University of Zaragoza. Funding from the Danish National Research Foundation is gratefully acknowledged (Sapere Aude, Advanced Grant no. 12-124980). Rahbek is also affiliated with CREATES, funded by the Danish National Research Foundation.

under the null but, in contrast to what is claimed in the literature, it is also consistent under the alternative. The theoretical arguments used to establish this includes the introduction of Ornstein-Uhlenbeck (OU) processes with random drift parameters as originally discussed by Basawa *et al.* (1991), where univariate versions of such processes were used to prove that the standard bootstrap fails in the unit root case.

To initially fix ideas, assume that the data generating process (DGP) of the p-dimensional observations $\{X_t\}_{t=1,2,..,T}$ satisfies the kth order VAR model with r co-integrating relations as given by

$$\Delta X_t = \alpha\beta' X_{t-1} + \sum\nolimits_{i=1}^{k-1} \Gamma_i \Delta X_{t-i} + \varepsilon_t, \quad t = 1, ..., T, \tag{1}$$

with $\{\varepsilon_t\}$ independent and identically distributed (i.i.d.) with mean zero and full-rank variance matrix Ω, and where the initial values $X_{1-k}, ..., X_0$ are fixed in the statistical analysis. We assume the classic co-integration conditions hold, which we denote by the '$I(1, r)$ conditions'; that is, (a) the characteristic polynomial associated with (1) has $p - r$ roots equal to 1 and all other roots outside the unit circle, and (b) α and β have full column rank r. Under these conditions X_t is $I(1)$ with co-integration rank r, such that the co-integrating relations $\beta' X_t$ are stationary.

We are interested in hypothesis testing on the co-integrating vectors β, where we focus on testing 'known co-integrating relations', which corresponds to testing the null hypothesis $\mathsf{H}_0 : \beta = \tau$. That is, with τ a known $p \times r$ matrix of full column rank r, the subspaces spanned by β and τ, respectively, are identical. For the co-integrated VAR model, the most used test is the well-known likelihood ratio (LR) test of Johansen (1996), which rejects for large values of the LR statistic, say $Q_T(\tau)$, formally defined in Section 2 below. Johansen (1996) established that the $Q_T(\tau)$ is asymptotically χ^2 when H_0 is true. Moreover, when H_0 does not hold, as shown here $Q_T(\tau)$ diverges, see Remark 3.4 below. Hence, the LR test is consistent.

Bootstrap schemes based on restricted parameter estimates were originally proposed without theory in Fachin (2000) and Gredenhoff and Jacobson (2001); see also Li and Maddala (1996, 1997) for the case of co-integrating regressions. Based on simulations, they suggest to estimate (1) under H_0 and then use the corresponding (restricted) estimates $\tilde{\alpha}, \tau, \tilde{\Gamma}_i$, $i = 1, ..., k-1$, in the bootstrap generating process (BGP). Then, the null distribution of $Q_T(\tau)$ is approximated by the (conditional) distribution of $Q_T^*(\tau)$, i.e. the LR statistic for testing $\beta = \tau$, computed on the bootstrap sample. However, Omtzigt and Fachin (2006) conjecture that this bootstrap based on restricted parameter estimates is inadmissible under the alternative hypothesis.

Our proof of the validity of the bootstrap initially requires investigating the asymptotic behavior of the restricted parameter estimates. Although these estimates are consistent only when H_0 is true, in Lemma 1 below we establish that, nonetheless, even when H_0 does not hold these estimates converge to pseudo-true values which satisfy, for some $r^* < r$, the $I(1, r^*)$ conditions. Importantly, we show that this lemma does not imply that the resulting bootstrap sample is an $I(1, r^*)$ process: instead, we establish in Proposition 1 that the BGP has $p - r$ unit roots, and $(r - r^*)$ additional stochastic roots local-to-unity; i.e., stochastic roots which asymptotically tend to one at the rate of T. In this sense, the BGP has co-integrating rank r^* (which can be lower than the true co-integrating rank r) and $p - r^* = (p - r) + (r - r^*)$ common stochastic trends, which are of order $T^{1/2}$. This novel representation of the bootstrap sample allows us to establish in Theorem 1 that the bootstrap LR statistic is bounded in probability even when the incorrect null hypothesis is imposed on the bootstrap data, while being χ^2 under the true null hypothesis. Hence, the bootstrap LR test is first-order valid and consistent in the usual sense.

The paper is organized as follows. In Section 2 we outline the reference LR test as well as the main bootstrap algorithm(s). The large sample properties of the bootstrap are established in Section 3. In Section 4, we report some finite-sample comparisons from a Monte Carlo study which suggests that the proposed bootstrap test allows for substantial improvements not only over the finite sample properties of the asymptotic LR test, but also relative to the corresponding bootstrap procedure of Omtzigt and Fachin (2006) and the Bartlett-corrected LR test of Johansen (2000, 2002*a*). In Section 5 we discuss inclusion of deterministic components, with focus on the model with an intercept. Section 6 concludes. Mathematical proofs, additional Monte Carlo results and a full treatment of the model with an intercept are collected in the accompanying Supplement (Cavaliere, Nielsen and Rahbek, 2014, CNR hereafter).

Notation. We use P^*, E^* and Var^* respectively to denote probability, expectation and variance, conditional on the original sample. With $\xrightarrow{w}$, $\xrightarrow{p}$ and $\xrightarrow{w^*}_p$ we denote weak convergence, convergence in probability and weak convergence in probability, respectively. The Euclidean norm of the vector x is $\|x\| := (x'x)^{1/2}$, where $x := y$ indicates that x is defined by y. For a given sequence X_T^* computed on the bootstrap data, $X_T^* = o_p^*(1)$ and $X_T^* \xrightarrow{p^*}_p X$ mean that for any $\epsilon > 0$, $P^*(||X_T^*|| > \epsilon) \xrightarrow{p} 0$ and $P^*(||X_T^* - X|| > \epsilon) \xrightarrow{p} 0$, respectively, as $T \to \infty$. Similarly, $X_T^* = O_p^*(1)$ means that, for every $\epsilon > 0$, there exists a constant $M > 0$ such that, for all large T, $P(P^*(||X_T^*|| > M) < \epsilon)$ is arbitrarily close to one. Also, $\mathbb{I}(\cdot)$ denotes the indicator function; $\lfloor \cdot \rfloor$ denotes the integer part of its argument; I_k denotes the $k \times k$ identity matrix and $0_{j \times k}$ the $j \times k$ matrix of zeroes. The space spanned by the columns of any $m \times n$ matrix a is denoted as $\text{span}(a)$; if a is of full column rank $n < m$, then $\bar{a} := a\,(a'a)^{-1}$ and $a_\perp$ is an $m \times (m - n)$ full column rank matrix satisfying $a_\perp' a = 0$; for any square matrix a, $\det(a)$ denotes its determinant. Finally, with $q \in \mathbb{R}$, $(q)^+ := \max\{0, q\}$. The space of $m \times 1$ vectors of càdlàg functions on the unit interval $[0, 1]$ is denoted by $\mathscr{D}^m$.

2 The LR Test and Bootstrap Algorithms

The (quasi) LR test statistic for the null hypothesis $\mathsf{H}_0 : \beta = \tau$ is based on maximization of the Gaussian likelihood of (1) with and without the restriction imposed. We use H_1 to refer to the unrestricted model where β is a freely varying $(p \times r)$ matrix. For estimation under H_1, let $S_{ij} := T^{-1} \sum_{t=1}^T R_{it} R_{jt}'$, $i, j = 0, 1$, with R_{0t} and R_{1t} respectively denoting ΔX_t and X_{t-1}, corrected (by OLS) for the lagged differences, $\Delta \mathbb{X}_{2t} := \left(\Delta X_{t-1}', ..., \Delta X_{t-k+1}'\right)'$. As in Johansen (1996), with $\hat{\lambda}_1 > \ldots > \hat{\lambda}_p$ the ordered solutions of the eigenvalue problem

$$\det(\lambda S_{11} - S_{10} S_{00}^{-1} S_{01}) = 0, \tag{2}$$

$\hat{\beta}$ is found as the eigenvectors corresponding to the r largest eigenvalues. The (unrestricted) maximized likelihood function is found, up to a constant, from $L_{\max}^{-2/T}(\hat{\beta}) = \det(S_{00})\prod_{i=1}^{r}(1-\hat{\lambda}_i) = \det(\hat{\Omega})$, where $\hat{\Omega} = T^{-1}\sum_{t=1}^{T}\hat{\varepsilon}_t\hat{\varepsilon}_t' = S_{00} - S_{01}\hat{\beta}(\hat{\beta}'S_{11}\hat{\beta})^{-1}\hat{\beta}'S_{10}$ is the sample covariance matrix of the residuals. Under the restriction H_0: $\beta = \tau$, likelihood estimation in (1) reduces to the regression of ΔX_t on $(X_{t-1}'\tau,\ \Delta\mathbb{X}_{2t}')'$, and the corresponding maximized likelihood function is $L_{\max}^{-2/T}(\tau) = \det(\tilde{\Omega})$, with $\tilde{\Omega} = T^{-1}\sum_{t=1}^{T}\tilde{\varepsilon}_t\tilde{\varepsilon}_t' = S_{00} - S_{01}\tau(\tau'S_{11}\tau)^{-1}\tau'S_{10}$. The LR statistic is then given by

$$Q_T(\tau) = -2\log Q(\mathsf{H}_0|\mathsf{H}_1) = -T\log\det(\tilde{\Omega}^{-1}\hat{\Omega}) \tag{3}$$

and, under H_0, $Q_T(\tau)$ converges to a χ^2 distribution with $(p-r)\,r$ degrees of freedom.

The bootstrap implementation of the LR test for $\mathsf{H}_0 : \beta = \tau$ which we advocate is based around a bootstrap recursion which mimics the DGP under the null hypothesis. To that end, let $\Psi := (\Gamma_1, ..., \Gamma_{k-1})$ collect the short run parameters and recall that when $\beta = \tau$, the ML estimators of α and Ψ are given by $\tilde{\alpha} = S_{01}\tau(\tau'S_{11}\tau)^{-1}$ and $\tilde{\Psi} := (\tilde{\Gamma}_1, .., \tilde{\Gamma}_{k-1}) = M_{02}M_{22}^{-1} - \tilde{\alpha}\tau'M_{12}M_{22}^{-1}$ respectively, with $M_{ij} := T^{-1}\sum_{t=1}^{T} Z_{it}Z_{jt}'$, $i, j = 0, 1, 2$, and Z_{0t}, Z_{1t} and Z_{2t} denoting ΔX_t, X_{t-1} and $\Delta\mathbb{X}_{2t}$, respectively.

Using these estimates, the bootstrap algorithm we consider is based on the recursion

$$\Delta X_t^* = \tilde{\alpha}\tau'X_{t-1}^* + \sum\nolimits_{i=1}^{k-1}\tilde{\Gamma}_i\Delta X_{t-i}^* + \varepsilon_t^*, \tag{4}$$

where the bootstrap shocks ε_t^* are obtained by re-sampling (after re-centering) from the corresponding restricted residuals, say $\tilde{\varepsilon}_t := \Delta X_t - \tilde{\alpha}\tau'X_{t-1} - \tilde{\Psi}\Delta\mathbb{X}_{2t}$, obtained from estimating (1) under the null hypothesis H_0. The bootstrap LR statistic for $\mathsf{H}_0 : \beta = \tau$ is simply the LR statistic computed on the bootstrap sample, $Q_T^*(\tau) := -T\log\det(\tilde{\Omega}^{*-1}\hat{\Omega}^*)$, where $\hat{\Omega}^* = T^{-1}\sum_{t=1}^{T}\hat{\varepsilon}_t^*\hat{\varepsilon}_t^{*\prime} = S_{00}^* - S_{01}^*\hat{\beta}^*(\hat{\beta}^{*\prime}S_{11}^*\hat{\beta}^*)^{-1}\hat{\beta}^{*\prime}S_{10}^*$ and $\tilde{\Omega}^* = T^{-1}\sum_{t=1}^{T}\tilde{\varepsilon}_t^*\tilde{\varepsilon}_t^{*\prime} = S_{00}^* - S_{01}^*\tau(\tau'S_{11}^*\tau)^{-1}\tau'S_{10}^*$; here the '*' superscript on product moment matrices and parameter estimates refer to the bootstrap sample.

We now detail in Algorithm 1 our bootstrap LR test for $\mathsf{H}_0 : \beta = \tau$.

Algorithm 1:

(i) Estimate model (1) under H_0, producing estimates $\{\tilde{\alpha}, \tilde{\Psi} = (\tilde{\Gamma}_1, ..., \tilde{\Gamma}_{k-1})\}$ and corresponding residuals, $\tilde{\varepsilon}_t$.

(ii) Construct the bootstrap sample recursively from (4) initialized at $X_j^* = X_j$, $j = -k+1, ..., 0$, and with the T bootstrap errors ε_t^* obtained by i.i.d. re-sampling of the re-centered residuals, $\tilde{\varepsilon}_t^c := \tilde{\varepsilon}_t - T^{-1}\sum_{i=1}^{T}\tilde{\varepsilon}_i$;

(iii) Using the bootstrap sample, $\{X_t^*\}$, compute the bootstrap analogue of the LR statistic, $Q_T^*(\tau)$, and define the corresponding p-value as $p_T^* := 1 - G_T^*(Q_T(\tau))$, with $G_T^*(\cdot)$ denoting the conditional (on the original data) cdf of $Q_T^*(\tau)$.

(iv) The bootstrap test of H_0 at level η rejects H_0 if $p_T^* \leq \eta$.

Remark 2.1 It may happen that in finite samples the estimates obtained in step (i) do not satisfy the $I(1, r)$ conditions; if this is the case, the bootstrap sample may be explosive. Hence, after step (i) one may check that the roots of the equation $\det(\tilde{A}(z)) = 0$, where $\tilde{A}(z) := (1-z)I_p - \tilde{\alpha}\tau'z - \sum_{i=1}^{k-1}\tilde{\Gamma}_i z^i(1-z)$, are either equal to 1 or outside the unit circle. This step has been advocated for bootstrap co-integration rank testing by Swensen (2006, 2009) and Cavaliere, Rahbek and Taylor (2012). By Lemma 1 below it holds that $\tilde{A}(z)$ always satisfies this condition as T diverges even when H_0 is false, so that checking this condition becomes irrelevant in large samples. Monte Carlo simulations show that it is also irrelevant in samples of moderate size.

Remark 2.2 The cdf $G_T^*(\cdot)$ required in Step (iv) can be approximated numerically by generating B (conditionally) independent bootstrap statistics, $Q_{T:b}^*$, $b = 1, ..., B$, computed as in Algorithm 1. The bootstrap p-value is then approximated by $\tilde{p}_{T,B}^* := B^{-1}\sum_{b=1}^{B}\mathbb{I}(Q_{T:b}^*(\cdot) > Q_T(\cdot))$, and is such that $\tilde{p}_{T,B}^* \overset{a.s.}{\to} p_T^*$ as $B \to \infty$; cf. Hansen (1996), Andrews and Buchinsky (2000) and Davidson and MacKinnon (2000).

Remark 2.3 The key feature of this bootstrap scheme is that the estimated parameters used in constructing the bootstrap sample, are obtained under the restriction of the null hypothesis. As advocated by Omtzigt and Fachin (2006), it is also possible to base the bootstrap on unrestricted parameter estimates, in line with much statistical literature on bootstrap hypothesis testing. That is, with $\hat{\alpha}, \hat{\beta}, \hat{\Gamma}_i$, $i = 1, ..., k-1$ denoting the unrestricted estimates, the bootstrap sample is generated according to the recursion

$$\Delta X_t^* = \hat{\alpha}\hat{\beta}' X_{t-1}^* + \sum\nolimits_{i=1}^{k-1}\hat{\Gamma}_i \Delta X_{t-i}^* + \varepsilon_t^*. \tag{5}$$

In this case, the bootstrap test statistic is the previous LR statistic but with $\tau := \hat{\beta}$; that is, the LR statistic for the null hypothesis that β equals the unrestricted estimate $\hat{\beta}$. We label this approach 'unrestricted bootstrap' in the following. When H_0 holds, the restricted estimates will be more efficient than the unrestricted estimates, and one would expect the bootstrap test from Algorithm 1 to display superior finite sample size properties than the unrestricted bootstrap test based on (5). This prediction on the size of the bootstrap tests is supported by the Monte Carlo comparisons in Section 4.

3 Asymptotic Analysis

The focus in this section is to establish that the bootstrap LR test from Algorithm 1 is asymptotically valid, i.e. that it is asymptotically correctly sized under the null and consistent under the alternative. The proof of this requires three distinct steps.

First, we derive in Lemma 1 the limiting behavior of the restricted Gaussian quasi maximum likelihood estimator (QMLE) under the alternative hypothesis. Specifically, we show that, as the sample size increases, the QMLE converges to pseudo-true parameters which satisfy the $I(1, r^*)$ conditions, where $r^* < r$ is formally defined below. That is, as $p - r^* > p - r$, we find that under the alternative, the pseudo-true parameters allow for more common stochastic trends than in the DGP. While this implies that the $Q_T(\tau)$ test is consistent (as we formally prove in Remark 3.4 below), somewhat surprisingly it does not imply that the bootstrap statistic $Q_T^*(\tau)$ diverges. To see this, as a second step we show in Proposition 1 that the bootstrap sample indeed has $p - r^*$ stochastic common trends which are of order $T^{1/2}$ (as I(1) common trends), and r^* components which are of the order of classic stationary relations. Thus, in this sense the bootstrap sample mimics the behavior of the system generated by the pseudo-true parameters derived in

Lemma 1. By applying this result in a detailed likelihood-based asymptotic analysis, in the third step, we show that both $\tilde{\Omega}^*$ and $\hat{\Omega}^*$ converge (in probability) to the same pseudo true value Ω_0^*, at a rate such that, crucially, $Q_T^*(\tau) = -T \log\det(\tilde{\Omega}^{*-1}\hat{\Omega}^*)$ is asymptotically *bounded*, irrespectively of the null hypothesis to be true or not. Moreover, since under the null hypothesis, $Q_T^*(\tau)$ is asymptotically χ^2 (in probability), these two results guarantee that the bootstrap test is asymptotically correctly sized and consistent under the alternative.

The results established in this section hold for any $I(1,r)$ DGP satisfying the assumptions stated e.g. in Johansen (1996), coupled with a finite fourth order moment condition as required for the bootstrap co-integration tests of Swensen (2006) and Cavaliere *et al.* (2012). Henceforth, with the subscript '0' denoting the true parameter values, we make the following assumptions.

ASSUMPTION 1 *The true parameters $\{\alpha_0, \beta_0, \Psi_0\}$ in (1) satisfy the I(1,r_0) conditions.*
ASSUMPTION 2 *The innovations $\{\varepsilon_t\}$ in (1) form an i.i.d. sequence with $E(\varepsilon_t) = 0$, $E(\varepsilon_t\varepsilon_t') = \Omega_0$ with Ω_0 positive definite, and $E\|\varepsilon_t\|^4 \leq K < \infty$.*

3.1 Behavior of Restricted Parameter Estimators

When H_0 holds, $\text{span}(\tau) = \text{span}(\beta_0)$ and all r_0 linear combinations $\tau' X_t$ are stationary, as is well-known from classic co-integration analysis. Conversely, when H_0 does not hold, $\text{span}(\tau) \neq \text{span}(\beta_0)$. While the two subspaces are different, they may share (co-integrating) vectors. Thus, with r^* denoting the dimension (i.e., the 'pseudo-rank') of $\text{span}(\beta_0) \cap \text{span}(\tau)$, there exist ϕ and ξ, both of dimension $r_0 \times r^*$ such that $\beta_0\phi = \tau\xi$. Furthermore, we can set $\beta_0^* := \beta_0\phi \in \text{span}(\tau)$, such that the r^* linear combinations $\beta_0^{*\prime} X_t$ are stationary. In other words, if $r^* = 0$, all r_0 linear combinations $\tau' X_t$ will be integrated of order one, i.e. I(1), while if $r^* > 0$, r^* linear combinations of $\tau' X_t$ are stationary, while $r_0 - r^*$ are not. Note that $r^* = r_0$ only when the null hypothesis is true. In addition, $r^* = 0$ requires that $r_0 \leq p - r_0$, or $2r_0 - p \geq 0$. In terms of possible values, r^* may thus take values in the closed interval defined as $\{(2r_0 - p)^+, ..., r_0\}$.

The pseudo-rank r^* is a fundamental quantity to determine the (large sample) behavior of estimates used to generate the bootstrap sample. In the next lemma, we characterize the limiting system generated by the pseudo-true parameters (i.e., the limit of the QML estimators under the alternative) as an $I(1, r^*)$ system.

LEMMA 1 *With $\tilde{\Pi} = \tilde{\alpha}\tau'$, $\tilde{\Psi}$, and $\tilde{\Omega}$ the restricted QML estimators it follows that under Assumptions 1 and 2, as $T \to \infty$,*

$$\tilde{\Pi} \xrightarrow{p} \Pi_0^* = \alpha_0^*\beta_0^{*\prime}, \quad \tilde{\Psi} \xrightarrow{p} \Psi_0^* = \left(\Gamma_{0,1}^*, ..., \Gamma_{0,k-1}^*\right) \quad \text{and} \quad \tilde{\Omega} \xrightarrow{p} \Omega_0^* > \Omega_0,$$

where the pseudo-true parameters α_0^*, β_0^* and Ψ_0^* and Ω_0^*, defined in Section C.1 of CNR, satisfy the $I(1, r^*)$ conditions.

REMARK 3.1 By Lemma 1, estimation under the wrong restrictions on β implies that the estimated Π matrix has, in the limit, rank *lower* than the true co-integrating rank r_0. In terms of the characteristic polynomial associated to the pseudo-true parameters,

$A^*(z) := (1-z)I_p - \alpha_0^*\beta_0^{*\prime}z - \sum_{i=1}^{k-1}\Gamma_{0,i}^* z^i(1-z)$, we have that $A^*(z)$ has $p-r^*$ unit roots and all other roots outside the unit circle. As $p-r^* \geq p-r_0$, we have $p-r_0$ unit roots corresponding to those characterizing the true DGP, and r_0-r^* additional unit roots generated by the fact that, when imposing the incorrect null in the estimation, there can at most be $r^* < r_0$ co-integrating relations. Stated differently, imposing incorrect co-integrating vectors, (asymptotically) decreases the co-integration rank.

Remark 3.2 It is a simple consequence of Lemma 1 that if, and only if, H_0 is true, we have that $r^* = r_0$, $\alpha_0^*\beta_0^{*\prime} = \alpha_0\beta_0'$ and the restricted estimators $\tilde{\alpha}$, $\tilde{\Psi}$, $\tilde{\Omega}$ are consistent.

Remark 3.3 As noted, Lemma 1 implies that, in the limit, the estimates $\{\tilde{\alpha}, \tau, \tilde{\Psi}\}$ satisfy the $I(1, r^*)$ conditions, even if r^* is lower than the true rank r_0, and consequently in the limit, the root check as in Remark 2.1 becomes redundant.

Remark 3.4 An important corollary of Lemma 1, which will be a key ingredient for the proof of the consistency of the bootstrap LR test is that, when H_0 does not hold, $Q_T(\tau)$ diverges. To see this, it suffices to notice that $Q_T(\tau) = T\log\det(\hat{\Omega}^{-1}\tilde{\Omega})$, with $\hat{\Omega}$ the covariance matrix of the estimated residuals from the unrestricted model. As is well known (Johansen, 1996), $\hat{\Omega} \xrightarrow{p} \Omega_0$ while from Lemma 1, $\tilde{\Omega} \xrightarrow{p} \Omega_0^* > \Omega_0$.

3.2 Behavior of the Bootstrap Sample

A direct implication of Lemma 1 for the bootstrap recursion in (4) would seem to be that the bootstrap sample generated by (4) is $I(1)$ with co-integration rank r^* in large samples. That is, in large samples X_t^* will behave like $X_t^\dagger$ defined by the recursion

$$\Delta X_t^\dagger = \alpha_0^*\beta_0^{*\prime}X_{t-1}^\dagger + \sum\nolimits_{i=1}^{k-1}\Gamma_{0,i}^*\Delta X_{t-i}^\dagger + \varepsilon_t^* \tag{6}$$

which generates $p-r^*$ unit roots, all other roots being outside the unit circle. However, as we will show next, this is only true for the stationary transformations $\beta_0^{*\prime}X_t^*, \Delta X_t^*$. Conversely, in the common trend directions $\alpha_{0\perp}^{*\prime}\Gamma_0^*X_t^*$ (with $\Gamma_0^* := I_p - \sum_{i=1}^{k-1}\Gamma_{0,i}^*$) this is not true, as in the limit the bootstrap data X_t^* have r_0-r^* additional roots *local-to-unity* rather than additional r_0-r^* unit roots.

The next proposition states this formally and is a fundamental part of our asymptotic analysis as it provides the theory needed for controlling the asymptotic behavior of the bootstrap sample and therefore of the bootstrap $Q_T^*(\tau)$ statistic.

Proposition 1 *Under Assumptions 1 and 2, with $\{X_t^*\}$ generated as in Algorithm 1:*
(i) *The bootstrap process X_t^* has the common trend representation*

$$X_t^* = C_z^*Z_t^* + S_t^*, \tag{7}$$

with $C_z^* = \beta_{0\perp}^*(\alpha_{0\perp}^{*\prime}\Gamma_0^*\beta_{0\perp}^*)^{-1}$ and $\max_{t=1,..,T}||S_t^*|| = o_p^*(T^{1/2})$. The process Z_t^* satisfies

$$T^{-1/2}Z_{\lfloor T\cdot\rfloor}^* \xrightarrow{w^*}_p Z(\cdot) \tag{8}$$

on $\mathscr{D}^{p-r^*}$. Here Z is a $(p-r^*)$-dimensional Ornstein-Uhlenbeck process satisfying

$$dZ(u) = \pi^*Z(u)\,du + \alpha_{0\perp}^{*\prime}dW^*(u), \tag{9}$$

with π^* a random drift parameter of dimension $(p-r^*)\times(p-r^*)$ of reduced rank $r_0 - r^*$, a.s., and W^* the p-dimensional Brownian motion with covariance matrix Ω_0^*, defined by the functional limit of the bootstrap errors, $T^{-1/2}\sum_{t=1}^{\lfloor T\cdot\rfloor}\varepsilon_t^* \xrightarrow{w^*}_p W^*(\cdot)$ on $\mathscr{D}^p$.
(ii) *With $X_t^\dagger$ generated recursively as in (6) with initial values $X_t^\dagger = X_t$, $t = -k+1, ..., 0$,*

$$\max_{t=1,...,T} ||\operatorname{Var}^*(R_t^* - R_t^\dagger)|| = O_p(T^{-1}) \tag{10}$$

where $R_t^* := (X_t^{*\prime}\beta_0^*, \Delta X_t^{*\prime})'$ and $R_t^\dagger := (X_t^{\dagger\prime}\beta_0^*, \Delta X_t^{\dagger\prime})'$.

Remark 3.5 To illustrate the results in Proposition 1, consider a VAR(1) with $r_0 = 1$. Estimation under $\mathsf{H}_0 : \beta = \tau$ requires fitting the regression $\Delta X_t = \alpha\tau' X_{t-1} + \varepsilon_t$, yielding $\tilde{\alpha} = M_{01}\tau(\tau' M_{11}\tau)^{-1}$ and $\tilde{\Pi} = \tilde{\alpha}\tau'$. If H_0 does not hold, the regressor $\tau' X_t$ is $I(1)$ and therefore, even though $\alpha_0 \neq 0$, we have that $\tilde{\alpha} \xrightarrow{p} 0$ and that $T\tilde{\alpha}$ has a unit-root type distribution. Consequently, the restricted estimator of Π satisfies (as in Lemma 1) $\tilde{\Pi} := \tilde{\alpha}\tau' \xrightarrow{p} 0_{p\times p}$ (of rank $r^* = 0$) and (as in Proposition 1) $T\tilde{\Pi} \xrightarrow{p} \pi^*$, where π^* is a random $p\times p$ matrix of rank $r - r^* = 1$. As a result, the bootstrap recursion $\Delta X_t^* = \tilde{\alpha}\tau' X_{t-1}^* + \varepsilon_t^*$ generates a co-integrated VAR(1) with $r = 1$ for any finite T, but since $\tilde{\alpha}\tau' = O_p(T^{-1})$, it is also *local* to a VAR(1) with rank $r^* = 0$. As a consequence, $T^{-1/2}X_{\lfloor T\cdot\rfloor}^* \xrightarrow{w^*}_p Z(\cdot)$, where $Z(u)$ solves $dZ(u) = \pi^* Z(u)\,du + dW^*(u)$.

Remark 3.6 The behavior of X_t^* in the $(p-r^*)$ non-stationary directions given by $\beta_{0\perp}^{*\prime}X_t^*$, is determined by Z_t^* which, normalized by $T^{1/2}$, converges weakly (in probability) to the multivariate OU process with random drift Z. In terms of Z, observe that π^* has rank $r_0 - r^*$, such that the original $(p - r_0)$ unit roots are preserved and, in addition, $r_0 - r^*$ extra (stochastic) roots local-to-unity are entering, see (C.20) in CNR for the definition of the random drift parameter π^*. That is, the convergence rate of the asymptotic additional $(r_0 - r^*)$ local to unit roots is such that they enter the distribution of the limit OU process Z. On the other hand, when $r^* = r_0$ such that the null hypothesis applies, $\beta_{0\perp}^* = \beta_{0\perp}$, $\alpha_{0\perp}^* = \alpha_{0\perp}$ and $\pi^* = 0$, which implies $Z = \alpha_{0\perp}' W^*$.

Remark 3.7 Multivariate OU processes with reduced rank drift as Z in (9) appear, in the special case of non-stochastic drift, in Kessler and Rahbek (2001, 2004) for the analysis of co-integrated diffusion processes and in Johansen (1996) for the derivation of the local power function of the trace test for co-integration. However, in our case the drift parameter π^* is stochastic and, in this sense, Z has a *random* distribution. Interestingly, a special case of Z with $p = 1$ and $r^* = 0$ has been found by Basawa *et al.* (1991), who show that the usual bootstrap estimator of the autoregressive coefficient in AR(1) processes with a unit root is not Dickey-Fuller distributed but instead has a random distribution, described in terms of a univariate OU process with random drift.

Remark 3.8 In the proof of Proposition 1 we show that the bootstrap process can be transformed into a recursion with unit roots and additional (stochastic) near unit roots. More precisely, let Y_t^* denote an appropriate rotation of the bootstrap process X_t^* (see Section C.2 of CNR for details), with ε_{yt}^* the correspondingly rotated bootstrap errors. Then Y_t^* can be given by the recursion

$$\Delta Y_t^* = (\tilde{a}_0^*\tilde{b}_0^{*\prime} + T^{-1}\tilde{a}_1^*\tilde{b}_1^{*\prime})Y_{t-1}^* + \sum_{i=1}^{k-1}\tilde{g}_i^*\Delta Y_{t-i}^* + \varepsilon_{yt}^*,$$

where $\tilde{a}_0^*\tilde{b}_0^{*\prime}$ and $\tilde{a}_1^*\tilde{b}_1^{*\prime}$ are stochastic and of reduced rank r^* and (r_0-r^*) a.s., respectively. With $\tilde{a}_1^*$, $\tilde{b}_1^{*\prime}$ and $\tilde{g}_i^*$ fixed (that is, conditional on the original sample), such recursion appears in the study of local alternatives for the co-integration rank statistic of rank r^* with $r_0 - r^*$ additional co-integrating vectors, as studied in Johansen (1996). Also, it is known from the literature on Wald tests on the co-integration parameters in the presence of near integrated regressors, see Elliott (1998), where the rank r^* here corresponds to the dimension of the co-integrating parameter, whereas the additional $r_0 - r^*$ near unit roots are due to the fact that some of the variables in the system are assumed near integrated. A crucial difference is that, in our bootstrap setting, the local alternative parameters are random, even in the limit, and therefore the limiting processes are stated in terms of a OU process with random drift.

Remark 3.9 It follows from Proposition 1 that the behavior of X_t^* in the r^* stationary directions given by $\beta_0^{*\prime}X_t^*$ as well as of the differenced process ΔX_t^* can be approximated by $\beta_0^{*\prime}X_t^\dagger$ and $\Delta X_t^\dagger$, see (10). Thus, indeed as claimed before Proposition 1, the limiting $I(1,r^*)$ process $X_t^\dagger$ (approximately) describes the stationary transformations of the bootstrap process X_t^*. This will be used repeatedly when deriving the limiting properties of the bootstrap statistic $Q_T^*(\tau)$ in the next section.

3.3 Behavior of the Bootstrap LR Statistic

We can now establish the asymptotic behavior of the bootstrap statistic $Q_T^*(\tau)$ of Algorithm 1, see Theorem 1 below. The proof of Theorem 1 is based on the asymptotic properties of the restricted ($\tilde{\alpha}^*\tau'$ and $\tilde{\Gamma}_i^*$, $i = 1, ..., k-1$) and unrestricted ($\hat{\alpha}^*\hat{\beta}^{*\prime}$ and $\hat{\Gamma}_i^*$, $i = 1, ..., k-1$) parameter estimators computed on the bootstrap sample, as implied by Lemma 1 and Proposition 1. In particular, these asymptotic properties imply that the restricted and unrestricted bootstrap residual covariance matrices $\tilde{\Omega}^*$ and $\hat{\Omega}^*$ are not only consistent for the pseudo-true covariance Ω_0^* but also, importantly, that $\tilde{\Omega}^* - \hat{\Omega}^*$ is of order $O_p^*\left(T^{-1}\right)$. This is used in a Taylor expansion to show that $Q_T^*(\tau)$ is bounded in probability under H_1, while when H_0 holds, we strengthen this result by showing that the limiting distribution is χ^2, as required.

Theorem 1 *Let the bootstrap statistic $Q_T^*(\tau)$ be generated as in Algorithm 1. Then, under the conditions of Proposition 1, $Q_T^*(\tau) = O_p^*(1)$. In addition, if H_0 holds, $Q_T^*(\tau) \xrightarrow{w}_p \chi^2(r(p-r))$.*

An immediate corollary of Theorem 1 is that the bootstrap test based on $Q_T^*(\tau)$ will be asymptotically correctly sized under the null hypothesis and consistent under H_1. These two results follow since, under Assumptions 1 and 2, $Q_T(\tau) \xrightarrow{w} \chi^2(r(p-r))$ when H_0 holds coupled with the fact that $Q_T(\tau)$ diverges under H_1, see Remark 3.4.

Corollary 1 *Under the conditions of Theorem 1, if H_0 holds the bootstrap p-value satisfies $p_T^* \xrightarrow{w} U[0,1]$. If H_1 holds, $p_T^* \xrightarrow{p} 0$.*

Remark 3.10 A crucial part of the proof of Theorem 1 is to establish the behavior of bootstrap unrestricted estimators $\hat{\alpha}^*$ and $\hat{\beta}^*$, obtained by reduced rank regression under H_1. We show that only r^* of the corresponding limiting eigenvalues are non-zero, and hence only r^* eigenvectors in $\hat{\beta}^*$ are identified in the limit. In fact, the normalized r^* eigenvectors, see (C.39) in CNR, have a weak limit (in probability) essentially of the form

$$\alpha_0^{*\prime}\Omega_0^{*-1}\int_0^1 dW^*(u)\,Z(u)'\left(\int_0^1 Z(u)\,Z(u)'\,du\right)^{-1},$$

which is mixed Gaussian as $\alpha_0^{*\prime}\Omega_0^{*-1}W^*$ and Z are independent and W^* Gaussian. Hence, inference based on these eigenvectors would be expected to be χ^2, in line with results in Elliott (1998), where an identified triangular system is considered under local to unity. However, under H_1, it holds that $r > r^*$ and therefore $\hat{\beta}^*$ is not mixed Gaussian and $Q_T^*(\tau)$ not χ^2-distributed under H_1. For this reason, in Theorem 1 we state the main result under H_1 in terms of the order of $Q_T^*(\tau)$.

Remark 3.11 The key fact that $Q_T^*(\tau) = O_p^*(1)$ under the alternative, while $Q_T(\tau)$ diverges is an implication of the fact that by Proposition 1 the bootstrap sample has $p - r^*$ linear combinations which are of order $T^{1/2}$, and r^* which are of order one (stationary linear combinations). This is unlike the original sample which, as the $I(1, r_0)$ conditions hold, has $p - r_0 \leq p - r^*$ stochastic trends which are of order $T^{1/2}$. Thus, the bootstrap sample mimics the limiting $I(1, r^*)$ system in Lemma 1 in the sense that it shares the number of $T^{1/2}$ stochastic trends and hence, as shown in the proof of Theorem 1, $Q_T^*(\tau)$ is bounded in probability. In contrast, the original data, under the $I(1, r_0)$ conditions, $r_0 > r^*$, has too few $T^{1/2}$ trends and the limiting LR statistic $Q_T(\tau)$ therefore diverges, as expected.

Remark 3.12 As demonstrated, $Q_T^*(\tau)$ is (asymptotically) χ^2 under H_0 and of order $O_p^*(1)$ – but not χ^2 – under H_1. Therefore, it may be conjectured that the bootstrap $Q_T^*(\tau)$ test has less power than the unrestricted bootstrap test of Remark 2.3, which is (asymptotically) χ^2 both under H_0 and under H_1. Indeed, as illustrated by Monte Carlo simulations in Section 4, the distribution of $Q_T^*(\tau)$ is shifted to the right with respect to the χ^2 distribution, but the corresponding power loss is difficult to deem relevant, given the massive size improvement of the bootstrap test based on $Q_T^*(\tau)$ with respect to the other methods. That is, even when the unrestricted bootstrap is size corrected – which is unfeasible in practice – the power loss is negligible.

Remark 3.13 Following the previous remark, in order to design a bootstrap test with χ^2 asymptotic distribution both under H_0 and H_1, one may think of a 'hybrid' bootstrap based on $Q_T^*(\tau)$ computed on a bootstrap sample satisfying

$$\Delta X_t^* = \hat{\alpha}\tau' X_{t-1}^* + \sum\nolimits_{i=1}^{k-1}\hat{\Gamma}_i \Delta X_{t-i}^* + \varepsilon_t^*,$$

where $\mathsf{H}_0 : \beta = \tau$ is imposed on the BGP but unrestricted estimates of the remaining parameters are employed. Since the unrestricted estimators are consistent irrespectively of H_0 being true or not, $Q_T^*(\tau)$ will indeed be asymptotically χ^2, when the limiting parameters $\{\alpha, \tau, \Gamma_1, ..., \Gamma_{k-1}\}$ satisfy the $I(1, r)$ conditions. While under H_0 this is obviously implied by Assumptions 1 and 2, under H_1 this is however not true in general. Consequently, the BGP may not be I(1), hence invalidating the bootstrap.

4 Numerical Results

Using Monte Carlo simulations we compare the finite sample properties of the proposed bootstrap LR test (denoted '*Bootstrap*' in the following) with: (i) the asymptotic LR test ('*Asymptotic*'), where the obtained LR statistic is compared with the χ^2 limiting null distribution; (ii) the Bartlett corrected LR test ('*Bartlett*') as proposed in Johansen (2000, 2002*a*), where the Bartlett correction is implemented using the unrestricted parameter estimates, as in Omtzigt and Fachin (2006); (iii) the bootstrap algorithm based on unrestricted parameter estimates ('*Unrestricted bootstrap*'), see Remark 2.3. Unlike existing simulations, the reported Monte Carlo results allow a detailed comparison between all four approaches. The simulations cover co-integrated systems of different dimensions ($p = 2, 4$), ranks ($r = 1, 2$) and lags ($k = 1, 2$) as well as the inclusion of an intercept (see Section 5). This section presents findings for a specific class of co-integrated processes (that is, $p = 4$, $r = 1$, $k = 1$; see below), while general and exhaustive simulations are reported in Section B of CNR. Importantly, the specific case presented here is representative for the conclusions based on all simulations.

The DGPs are as in (1) with Gaussian i.i.d. innovations, dimension $p = 4$, number of lags $k = 1$, co-integration rank $r = 1$, initialized at $X_0 = 0$. Sample sizes $T \in [40, 1000]$ are considered and different parameter configurations for α, β and Ω are chosen by setting $\alpha = (a_1, a_2, 0, 0)'$, $\beta = (1, b_1, 0, 0)'$ and $\Omega = I_4$, with a_1, a_2 and b_1 varying such that the process is $I(1)$ and co-integrated. Tests of the hypothesis $\mathsf{H}_0 : \beta = \tau$, are considered with $\tau = (1, 0, 0, 0)'$. To study empirical size, $b_1 = 0$ in the DGPs such that by varying (a_1, a_2) we fully characterize the relevant part of the parameter space under the null, see Johansen (2000, 2002*b*). The parameter grid applied is $(-a_1, a_2) \in A \times (A \cup \{0\})$, with $A := \{0.05, 0.10, ..., 0.9\}$. Behavior under the alternative is examined by letting $b_1 > 0$ vary and considering different combinations of (a_1, a_2) such that the $I(1, r)$ conditions hold with $r = 1$. We report here results for a 10% nominal significance level; qualitatively similar results are obtained for the 5% level. All experiments are run over 10,000 Monte Carlo replications. For the bootstrap tests, the bootstrap distribution of $Q_T^*(\tau)$ is simulated as in Remark 2.2 with $B = 999$ bootstrap replications. Ox code for the bootstrap test is available from the authors upon request.

Properties Under the Null. Taking the case $T = 100$ to illustrate, Panel I of Figure 1 reports empirical rejection frequencies (ERFs) for the four tests considered. The simulations show excellent size control of our proposed bootstrap, with ERFs close to the nominal 10% over the entire parameter space. The asymptotic test, on the other hand, displays severe size distortions, with ERFs above 50% for parameter values with slow error-correction, i.e. (a_1, a_2) close to the origin. Notice that the ERFs for the flat areas of the graph are above 13%. The Bartlett corrected test is an improvement over the asymptotic test, but the ERFs are again severely inflated for cases with slow error correction. The unrestricted bootstrap is quite similar to the Bartlett correction, and delivers inferior size control, as compared with the proposed bootstrap. Overall, in terms of ERFs under the null, the proposed bootstrap appears to be the only valid approach, as it allows for a proper size control over the entire parameter space, with the other approaches showing severe size distortions. The same result carries through when different sample sizes are considered, see Section B of CNR.

[Figure 1 around here]

Properties Under the Alternative. The large differences in empirical size make it difficult to meaningfully compare the tests under parameter configurations not satisfying the null (see e.g. Davidson and MacKinnon, 2006). In order to shed some light on the power properties of the considered tests we present ERFs obtained after size-adjusting the tests pointwise. To do so, for each given point in the parameter space (a_1, a_2, b_1, T), we first perform the simulation under the null (i.e., for $(a_1, a_2, b_1, T) = (a_1, a_2, 0, T)$) and record the nominal level that would have given an ERF equal to the desired 10%. Next, we use this adjusted nominal level in the simulations under the alternative hypothesis, where $b_1 \neq 0$. Panel II of Figure 1 shows the rejection frequencies for tests of H_0 against a sequence of DGPs with $b_1 > 0$, for $(a_1, a_2) = (-0.4, 0)$.

Results for $T = 60$ and 100 are summarized in Graphs (A)-(B), respectively. The results illustrate that the suggested bootstrap is very close, in terms of ERFs under the alternative, to the unfeasible, size-adjusted asymptotic test. In particular, the ERFs of the two tests differ only marginally, and become virtually identical as T increases. The reason seems to be that the distribution of $Q_T^*(\tau)$ under the alternative is shifted to the right with respect to the asymptotic (χ^2) null distribution, cf. Theorem 1 and Remark 3.12. Note, however, that the marginally higher rejection rates of the asymptotic tests are not attainable in practice—due to the severe size distortion under a true null.

To further illustrate the behavior of the tests as a function of the number of observations, in Graphs (C)-(D) we fix b_1 to 0.04 and 0.2, respectively, while letting T ranging from 40 to 1000. For small deviations from the null, the rejection frequencies of the proposed bootstrap test are indistinguishable from the asymptotic test, while they are only marginally lower for larger deviations from the null.

5 The Model with an Intercept

In this section we discuss how the results generalize to VAR models with deterministic trends. To illustrate, we focus on the case of an intercept; other cases of interest, such as linear trends, can be handled similarly. The model with an intercept is given by

$$\Delta X_t = \alpha\beta' X_{t-1} + \sum_{i=1}^{k-1} \Gamma_i \Delta X_{t-i} + \mu + \varepsilon_t, \quad t = 1, ..., T. \tag{11}$$

Under the $I(1, r)$ conditions, the co-integrating relations $\beta_0' X_t$ are stationary while X_t is $I(1)$, with a deterministic linear trend provided $C\mu_0 = C_g \alpha_{0\perp}' \mu_0 \neq 0$, with $C_g := \beta_{0\perp} (\alpha_{0\perp}' \Gamma_0 \beta_{0\perp})^{-1}$, $\Gamma_0 := I_p - \sum_{i=1}^{k-1} \Gamma_{0i}$. The LR test statistic for the null hypothesis $\mathsf{H}_0 : \beta = \tau$ is computed as in Section 2, but with X_t and ΔX_{t-i} $(i = 0, ..., k-1)$ corrected by OLS for a constant. Under H_0, the LR statistic $Q_T(\tau)$ converges to a χ^2 distribution with $(p - r) r$ degrees of freedom; this result is proved in Johansen (1996) for the case where $C\mu_0 \neq 0$, and extended in Lemma D.2 in CNR to the case $C\mu_0 = 0$.

Our suggested bootstrap LR test for $\mathsf{H}_0 : \beta = \tau$ is again implemented by estimating all parameters under H_0. In particular, $\tilde{\mu}$ is added to the right hand side of the bootstrap recursion in (4), and the bootstrap shocks ε_t^* are obtained by i.i.d. re-sampling of the restricted residuals, $\tilde{\varepsilon}_t := \Delta X_t - \tilde{\alpha}\tau' X_{t-1} - \tilde{\Psi}\Delta\mathbb{X}_{2t} - \tilde{\mu}$. The bootstrap LR statistic $Q_T^*(\tau)$ is the LR statistic for $\mathsf{H}_0 : \beta = \tau$, computed on the bootstrap sample.

In Section D of CNR the validity of the bootstrap is established by extending the results of Lemma 1 and Proposition 1 to the intercept case. Specifically, Lemma D.1 shows that $\tilde{\alpha}\tau', \tilde{\Psi}, \tilde{\Omega}$ and $\tilde{\mu}$ converge to pseudo-true values $\alpha_0^*\beta_0^{*\prime}, \Psi_0^*, \Omega_0^*$ and μ_0^* which correspond to an $I(1, r^*)$ system without, or with, a linear trend, depending on whether $C\mu_0 = 0$ or not. In addition, Proposition D.1 states that X_t^*, if $C\mu_0 \neq 0$, has a representation in terms of a linear trend and a $(p - r^* - 1)$-dimensional limiting OU process with random drift parameters, while if $C\mu_0 = 0$, the linear trend vanishes, and instead the limiting OU process is of dimension $(p - r^*)$. This result shows that, for any possible choice of τ, the bootstrap sample mimics the original sample in terms of deterministic terms. Lemma D.1 and Proposition D.1 allow – as done for the proof of Theorem 1 – to show that $Q_T^*(\tau) = O_p^*(1)$ in probability when H_0 does not hold, and χ^2 distributed otherwise. Monte Carlo simulations (see Section B in CNR) show that the finite sample results of Section 4 remain unchanged when an intercept is included.

6 Conclusions

We have discussed bootstrap implementations of the likelihood ratio test for hypotheses on the co-integrating vectors β based on restricted estimates of the underlying VAR model. Bootstrap testing of hypotheses on β turns out to be a highly non-standard inference problem, in particular in comparison to bootstrap tests on the co-integration rank r, see Cavaliere *et al.* (2012), Swensen (2006, 2009) and Trenkler (2009). We have argued that in the case of tests on β, the imposition of the incorrect null hypothesis affects the co-integration rank of the bootstrap DGP, which displays, in addition to the $(p - r)$ unit roots characterizing the true DGP, additional stochastic roots which are local-to-unity. As a consequence, in large samples the bootstrap DGP asymptotically behaves as an Ornstein-Uhlenbeck process with *random* drift matrix. This contrasts with the case of tests on the co-integration rank, where the multiplicity of unit roots in the bootstrap sample is always as implied by the null hypothesis being tested (i.e., $p - r$), regardless of the null hypothesis being true or not in the original sample.

Despite these findings, which may suggest the failure of the bootstrap in the framework of tests on β, we have shown that the bootstrap test statistic has the correct (χ^2) limiting distribution under the null, while remains bounded (in probability) when the null is not true. Hence, the bootstrap test based on restricted parameter estimates is asymptotically correctly sized and, strikingly, consistent under the alternative.

Monte Carlo comparisons suggest that this bootstrap procedure works extremely well in finite samples and outperforms the asymptotic test, procedures based on Bartlett adjustment, and the bootstrap test based on unrestricted estimates as proposed in Omtzigt and Fachin (2006).

Two obvious extensions of this work may be considered. First, although being based on the assumption of i.i.d. errors, our analysis can be generalized to the case of martingale difference sequences with possible conditional and unconditional heteroskedasticity (as done in Cavaliere, Rahbek and Taylor, 2010*a,b*, for tests on the co-integration rank). Second, since we have focused our analysis on the hypothesis of known co-integrating vectors, the case of more general hypotheses (see e.g. Johansen, 1996; Boswijk and Doornik, 2004) represents an obvious and important development of the results obtained in this paper. Both extensions are currently under investigation by the authors.

References

Andrews, D.W.K. and M. Buchinsky (2000): "A Three-Step Method for Choosing the Number of Bootstrap Repetitions", *Econometrica*, 68, 23–51.

Basawa, I.V., A.K. Mallik, W.P. McCormick, J.H. Reeves, and R.L. Taylor (1991): "Bootstrapping Unstable First-Order Autoregressive Processes", *The Annals of Statistics,* 19, 1098–1101.

Boswijk, H.P. and J.A. Doornik (2004): "Identifying, Estimating and Testing Restricted Cointegrated Systems: An Overview", *Statistica Neerlandica*, 58, 440–465.

Cavaliere, G., H.B. Nielsen and A. Rahbek (2014): "Supplement to 'Bootstrap Testing of Hypotheses on Co-Integration Relations in VAR Models'", *Econometrica Supplemental Material.*

Cavaliere, G., A. Rahbek and A.M.R. Taylor (2010*a*): "Co-Integration Rank Testing Under Conditional Heteroskedasticity", *Econometric Theory*, 26, 1719–1760.

—— (2010*b*): "Testing for Co-Integration in Vector Autoregressions with Non-Stationary Volatility", *Journal of Econometrics*, 158, 7–24.

—— (2012): "Bootstrap Determination of the Co-Integration Rank in Vector Autoregressive Models", *Econometrica,* 80, 1721–1740.

Davidson R. and J. MacKinnon (2000): "Bootstrap Tests: How Many Bootstraps?", *Econometric Reviews,* 19, 55–68.

—— (2006): "The Power of Bootstrap and Asymptotic Tests", *Journal of Econometrics*, 133, 421–441.

Elliott G. (1998): "On the Robustness of Cointegration Methods When Regressors Almost Have Unit Roots", *Econometrica*, 66, 149–158.

Fachin, S. (2000): "Bootstrap and Asymptotic Tests of Long-Run Relationships in Cointegrated Systems", *Oxford Bulletin of Economics and Statistics*, 62, 543–551.

Gredenhoff, M. and T. Jacobson (2001): "Bootstrap Testing Linear Restrictions on Cointegrating Vectors", *Journal of Business & Economic Statistics*, 19, 63–72.

Hansen, B.E. (1996): "Inference When a Nuisance Parameter Is Not Identified Under the Null Hypothesis", *Econometrica,* 64, 413–430.

Johansen, S. (1991): "Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models", *Econometrica,* 59, 1551–1580.

—— (1996): *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*, Second Edition, Oxford: Oxford University Press.

—— (2000): "A Bartlett Correction Factor for Tests on the Cointegrating Relations", *Econometric Theory*, 16, 740–778.

—— (2002*a*): "A Small Sample Correction for Tests of Hypotheses on the Cointegrating Vectors", *Journal of Econometrics,* 111, 195–221.

—— (2002*b*): "A Small Sample Correction of the Test for Cointegrating Rank in the Vector Autoregressive Model", *Econometrica,* 70, 1929–1961.

Kessler, M. and A. Rahbek (2001): "Asymptotic Likelihood Based Inference for Cointegrated Homogenous Gaussian Diffusions", *Scandinavian Journal of Statistics*, 28, 455–470.

—— (2004): "Identification and Inference for Multivariate Cointegrated and Ergodic Gaussian Diffusions", *Statistical Inference for Stochastic Processes,* 7, 137–151.

Li, H. and G. S. Maddala (1996): "Bootstrapping Time Series Models", *Econometric Reviews*, 15, 115–158.

—— (1997): "Bootstrapping Cointegrating Regressions", *Journal of Econometrics*, 80, 297–318.

Omtzigt, P. and S. Fachin (2006): "The Size and Power of Bootstrap and Bartlett-Corrected Tests of Hypothesis on the Cointegrating Vectors", *Econometric Reviews,* 25, 41–60.

Swensen, A.R. (2006): "Bootstrap Algorithms for Testing and Determining the Cointegration Rank in VAR Models", *Econometrica,* 74, 1699–1714.

—— (2009): "Corrigendum to "Bootstrap Algorithms for Testing and Determining the Cointegration Rank in VAR Models"", *Econometrica*, 77, 1703–1704.

Trenkler, C. (2009): "Bootstrapping Systems Cointegration Tests With a Prior Adjustment for Deterministic Terms", *Econometric Theory*, 25, 243–269.

Dept. of Statistical Sciences, University of Bologna, via delle Belle Arti 41, I-40126 Bologna, Italy and University of Copenhagen; giuseppe.cavaliere@unibo.it,

Dept. of Economics, University of Copenhagen, Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K, Denmark; heino.bohn.nielsen@econ.ku.dk

and

Dept. of Economics, University of Copenhagen, Øster Farimagsgade 5, Building 26, DK-1353 Copenhagen K, Denmark; anders.rahbek@econ.ku.dk

Manuscript received October, 2013; final revision received November, 2014

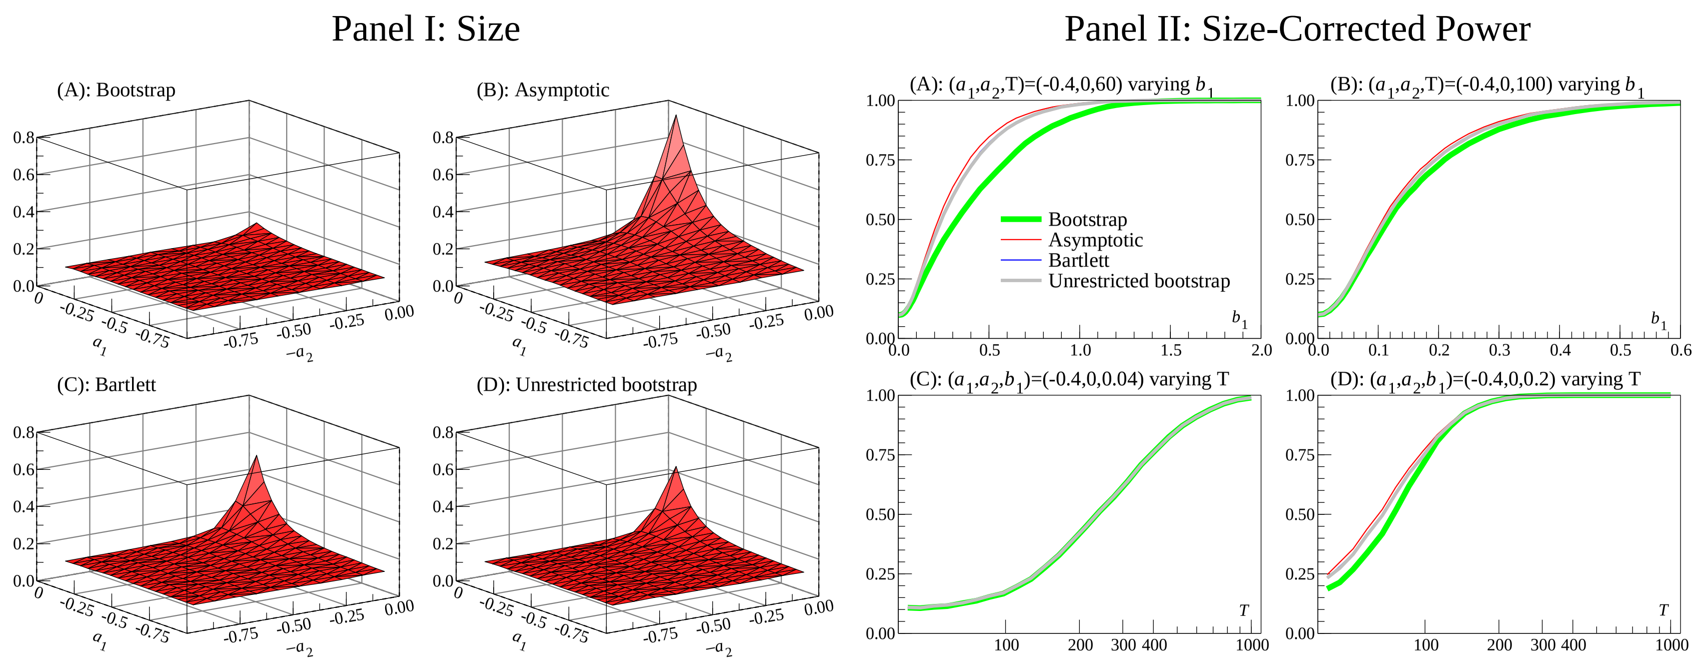

FIGURE 1: *Panel I shows empirical rejection frequencies under the null hypothesis for the different tests with $T = 100$ observations. Panel II shows pointwise size-corrected rejection frequencies for the hypothesis $\mathsf{H}_0 : \beta = \tau := (1, 0, 0, 0)'$, for a sequence of DGPs defined by $\beta = (1, b_1, 0, 0)'$, with $b_1 > 0$. Results are based on a 10% nominal level.*