Journal of Econometrics 158 (2010) 7–24

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Testing for co-integration in vector autoregressions with non-stationary volatility

Giuseppe Cavaliere [a,*], Anders Rahbek [b], A.M. Robert Taylor [c]

[a] *Department of Statistical Sciences, University of Bologna, Italy*

[b] *Department of Economics, University of Copenhagen and CREATES, Denmark*

[c] *School of Economics and Granger Centre for Time Series Econometrics, University of Nottingham, United Kingdom*

ARTICLE INFO

Article history:
Available online 11 March 2010

JEL classification:
C30
C32

Keywords:
Co-integration
Non-stationary volatility
Trace and maximum eigenvalue tests
Wild bootstrap

ABSTRACT

Many key macroeconomic and financial variables are characterized by permanent changes in unconditional volatility. In this paper we analyse vector autoregressions with non-stationary (unconditional) volatility of a very general form, which includes single and multiple volatility breaks as special cases. We show that the conventional rank statistics computed as in Johansen (1988, 1991) are potentially unreliable. In particular, their large sample distributions depend on the integrated covariation of the underlying multivariate volatility process which impacts on both the size and power of the associated co-integration tests, as we demonstrate numerically. A solution to the identified inference problem is provided by considering wild bootstrap-based implementations of the rank tests. These do not require the practitioner to specify a parametric model for volatility, or to assume that the pattern of volatility is common to, or independent across, the vector of series under analysis. The bootstrap is shown to perform very well in practice.

1. Introduction

A number of recent applied studies have suggested time-varying behaviour, in particular a general decline, in unconditional volatility in the shocks driving macroeconomic time series over the past twenty years or so is a relatively common phenomenon; see, *inter alia*, Busetti and Taylor (2003), Kim and Nelson (1999), McConnell and Perez Quiros (2000), van Dijk et al. (2002), Sensier and van Dijk (2004) and references therein. For example, Sensier and van Dijk (2004) report that over 80% of the real and price variables in the Stock and Watson (1999) data set reject the null of constant innovation variance against the alternative of a one-off change in variance. Similarly, Loretan and Phillips (1994) report evidence against the constancy of unconditional variances in stock market returns and exchange-rate data, while Hansen (1995) notes that empirical applications of autoregressive stochastic volatility models to financial data generally estimate the dominant root in the stochastic volatility process to be close to the non-stationarity boundary at unity. van Dijk et al. (2002) find evidence that volatility changes smoothly over time, while Watson (1999) argues that multiple changes in volatility are commonly observed. Cavaliere and Taylor (2007) report evidence of multiple volatility breaks and trending volatility in the monthly producer price inflation series from the well-known Stock and Watson (1999) database.

These findings have helped stimulate an interest amongst econometricians in analysing the effects of non-constant volatility on univariate unit root and stationarity tests; see, *inter alia*, Kim et al. (2002), Busetti and Taylor (2003), Cavaliere (2004), and Cavaliere and Taylor (2005, 2007, 2008). These authors show that standard unit root and stationarity tests based on the assumption of constant volatility can display significant size distortions in the presence of non-constant volatility. Cavaliere and Taylor (2008) develop wild bootstrap-based implementations of standard unit root tests which are shown to yield pivotal inference in the presence of non-stationary volatility. The impact of non-constant volatility on stable autoregressions has also been analysed by Hansen (1995), Phillips and Xu (2006), Xu and Phillips (2008) and Xu (2008), *inter alia*, who show that non-constant volatility can have a large impact on standard estimation and testing procedures.

Given that non-constant volatility has been found to be a common occurrence in univariate macroeconomic and financial time series, and to have a large impact on univariate time series procedures, it is clearly important and practically relevant to investigate the impact that such behaviour has on multivariate non-stationary time series methods. Indeed, using U.S. data, Hansen (1992a) has shown that the regression error in four published co-integrating relations (namely, real per capita consumption upon

* Corresponding address: Department of Statistical Sciences, Via Belle Arti 41, I-40126 Bologna, Italy. Tel.: +39 051 2098230; fax: +39 051 232153.
E-mail addresses: giuseppe.cavaliere@unibo.it (G. Cavaliere), anders.rahbek@econ.ku.dk (A. Rahbek), robert.taylor@nottingham.ac.uk (A.M.R. Taylor).

0304-4076/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2010.03.003

real per capita disposable income; aggregate non-durables and services consumption upon disposable income; real stock prices upon real dividends, short term upon long term interest rates) are all affected by non-stationary variances. Cavaliere and Taylor (2006) consider the impact of non-constant volatility on residual-based tests for the null hypothesis of co-integration.

In this paper we analyse the impact of non-stationary volatility in the (vector) innovation process driving a co-integrated vector autoregressive (VAR) model. We allow for innovation processes whose variances evolve over time according to a quite general mechanism which allows, for example, single and multiple abrupt variance breaks, smooth transition variance breaks, and trending variances. We analyse the impact that this has on the conventional trace and maximum eigenvalue statistics of Johansen (1988, 1991), demonstrating that the asymptotic null distributions of these statistics depend upon the (asymptotic) integrated covariation of the underlying volatility process. Simulation results for a one-time change in volatility suggest that this can have a large impact on both the size and power properties of the tests.

In order to solve the inference problem identified, at least within the class of volatility processes considered, we extend the univariate wild bootstrap-based unit root tests of Cavaliere and Taylor (2008) to the multivariate context by developing wild bootstrap-based implementations of Johansen's maximum eigenvalue and trace test statistics. Our proposed wild bootstrap procedure is set up in such a way that the practitioner is not required to specify any parametric model for volatility, or to assume that the pattern of volatility is common to, or independent across, the vector of series under analysis.

In a recent paper, Boswijk and Zu (2007) discuss maximum likelihood (ML) estimation and co-integration rank testing in VAR models when the (possibly non-stationary) spot volatility changes smoothly over time and can be estimated consistently. In contrast to the wild bootstrap approach used in this paper, the ML approach of Boswijk and Zu (2007), although based on a non-pivotal statistic, is asymptotically efficient (under certain conditions on the volatility process), exploiting (in the limit) the potential power gains that can arise from using the true likelihood ratio test in a correctly specified model. The approach of Boswijk and Zu (2007) might therefore be expected to deliver more powerful tests than obtain from our bootstrap approach. On the other hand, since estimation of the (spot) volatility process is not required for the bootstrap tests discussed here, they are likely to have much better finite sample size properties in the presence of non-stationary volatility than the corresponding tests of Boswijk and Zu (2007).[1] In this respect, the test proposed in Boswijk and Zu (2007) represents an important complement to the wild bootstrap method proposed in this paper. However, it is important to note that we adopt a different assumption to Boswijk and Zu (2007) regarding the class of non-stationary volatility processes allowed. In particular, while we allow for processes which display abrupt volatility shifts, Boswijk and Zu (2007) require the volatility process to be continuous. Moreover, our analysis does not require the existence of a consistent estimator of the underlying spot volatility.

Seo (2007) considers ML estimation (but does not discuss co-integration rank testing) of a co-integrated system when the errors are conditionally heteroskedastic. However, he imposes weak stationarity and so does not allow for time-varying behaviour in the unconditional volatility process. Other related work is considered by Hansen (2003), who considers estimation and testing in a co-integrated VAR model that allows for a finite number of deterministic breaks in the covariance matrix of the system. In contrast to the wild bootstrap approach outlined in this paper, Hansen (2003) adopts a parametric approach to structural change, requiring that the location of the breaks in the parameters of the covariance matrix and the number of co-integrating relations present in the system are known. A further difference is that the innovations in Hansen (2003) are assumed to be homoskedastic within each regime, such that the moving average representation of the system within each regime is identical to that given in Johansen (1996). In particular, this entails that both the co-integrating relations and the common trends are homoskedastic within each regime.

The remainder of the paper is organized as follows. Section 2 outlines our heteroskedastic co-integrated VAR model, giving both error correction and common trend representations for the model. Here we also discuss the form of the co-integrating relationships in the context of this model. In Section 3 the impact of non-stationary volatility on the large sample properties of Johansen's maximal eigenvalue and trace statistics is detailed. Here we also demonstrate the important result that the MLE of the parameters from our co-integrated VAR model remains consistent. Our wild bootstrap-based approach, which also incorporates a sieve procedure using the (consistently) estimated coefficient matrices from the co-integrated VAR model, is outlined in Section 4 and it is shown that this solves the inference problem caused by non-stationary volatility, yielding asymptotically pivotal co-integration tests. Monte Carlo experiments illustrating the effects of one-time (co)variance shifts on both standard and bootstrap co-integration tests are presented in Section 5. Here it is shown that the proposed bootstrap tests perform very well for finite samples. Section 6 concludes. All proofs are contained in the Appendix.

In the following '$\stackrel{w}{\to}$' denotes weak convergence, '$\stackrel{p}{\to}$' convergence in probability, and '$\stackrel{w}{\to}_p$' weak convergence in probability (Giné and Zinn, 1990; Hansen, 1996); $\mathbb{I}(\cdot)$ denotes the indicator function and '$x := y$' ('$x =: y$') indicates that x is defined by y (y is defined by x); $\lfloor \cdot \rfloor$ denotes the integer part of its argument. The notation $\mathcal{C}_{\mathbb{R}^{m\times n}}[0, 1]$ is used to denote the space of $m \times n$ matrices of continuous functions on [0, 1]; $\mathcal{D}_{\mathbb{R}^{m\times n}}[0, 1]$ denotes the space of $m \times n$ matrices of càdlàg functions on [0, 1]. The space spanned by the columns of any $m \times n$ matrix A is denoted as $\text{col}(A)$; if A is of full column rank $n < m$, then $A_\perp$ denotes an $m \times (m - n)$ matrix of full column rank satisfying $A'_\perp A = 0$. For any square matrix, A, $|A|$ is used to denote the determinant of A, $\|A\|$ the norm $\|A\|^2 := \text{tr}\left\{A'A\right\}$, and $\rho(A)$ its spectral radius (that is, the maximal modulus of the eigenvalues of A). For any vector, x, $\|x\|$ denotes the usual Euclidean norm, $\|x\| := (x'x)^{1/2}$.

2. The heteroskedastic co-integration model

We consider the following VAR(k) model in error correction format:

$$\Delta X_t = \alpha\beta' X_{t-1} + \Psi U_t + \mu D_t + \varepsilon_t, \quad t = 1, \ldots, T \tag{1}$$

$$\varepsilon_t = \sigma_t z_t \tag{2}$$

where: X_t and ε_t are $p \times 1$, σ_t is $p \times p$, $U_t := (\Delta X'_{t-1}, \ldots, \Delta X'_{t-k+1})'$ is $p(k-1) \times 1$ and $\Psi := (\Gamma_1, \ldots, \Gamma_{k-1})$, where $\{\Gamma_i\}_{i=1}^{k-1}$ are $p \times p$ lag coefficient matrices, D_t is a vector of deterministic terms, z_t is p-variate i.i.d., $z_t \sim (0, I_p)$, where I_p denotes the $p \times p$ identity matrix, and α and β are full column rank $p \times r$ matrices, $r \le p$. The initial values $\mathbb{X}_0 := (X'_0, \ldots, X'_{-k+1})'$ are assumed to be fixed

[1] This trade-off of size against power is well documented in the univariate case. Specifically, Boswijk (2005) shows that in the presence of time-varying unconditional volatility, a ML test combined with a consistent estimator of the spot volatility leads to unit root tests with power almost indistinguishable from that of the (asymptotic) local power envelope. This approach, however, can suffer from quite serious size distortions in small samples. Conversely, wild bootstrap unit root tests as in Cavaliere and Taylor (2008), although not optimal, lead to tests with very good size properties.