Journal of Econometrics 184 (2015) 62–80

Contents lists available at ScienceDirect

Journal of Econometrics

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

Multi-scale tests for serial correlation

Ramazan Gençay *, Daniele Signori

Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, British Columbia, V5A 1S6, Canada

ARTICLE INFO

Article history:
Received 8 January 2013
Received in revised form
5 August 2014
Accepted 18 August 2014
Available online 16 September 2014

JEL classification:
C1
C2
C12
C22
C58
F31
G0
G1

Keywords:
Serial correlation
Wavelets
Independence
Discrete wavelet transformation
Maximum overlap wavelet transformation
Variance ratio test
Variance decomposition

ABSTRACT

This paper introduces a new family of portmanteau tests for serial correlation. Using the wavelet transform, we decompose the variance of the underlying process into the variance of its low frequency and of its high frequency components and we design a variance ratio test of no serial correlation in the presence of dependence. Such decomposition can be carried out iteratively, each wavelet filter leading to a rich family of tests whose joint limiting null distribution is a multivariate normal. We illustrate the size and power properties of the proposed tests through Monte Carlo simulations.

1. Introduction

This paper proposes a new family of frequency-domain tests for the white noise hypothesis, the assumption that a process is uncorrelated. Frequency-domain tests take as their starting point the result that, under stationarity conditions, the linear dependence structure of a process $\{y_t\}$ is fully captured by its spectral density function $S_y(f)$. We focus our attention on the relation between the spectral density function and the variance,

$$\mathrm{var}(y) = 2\int_0^{1/2} S_y(f)\, df,$$

which, paraphrasing, says that the contribution of the frequencies in a small interval Δf containing f is approximately $S_y(f)\Delta f$. It is an elementary result that – when defined – the spectral density function of an uncorrelated process is constant or, in other words, that each frequency contributes equally to the variance of a white noise process; instead, when a process is serially correlated, each frequency generally contributes in different amounts and the spectral density function is non-constant.

Such contrast is the basis for the tests developed in this paper. Imagine that $\{y_t\}$ is a Gaussian white noise process (Fig. 1, left panel). Then high frequencies, say those in the band $[1/4, 1/2]$, will contribute exactly half of the total variance of $\{y_t\}$. On the other hand, if $\{y_t\}$ is an autoregressive process of order 1 with a positive coefficient (right panel), high frequencies will account for less than half of the total variance. This example motivates the introduction of the variance ratio $\mathcal{E}(a, b)$, defined as the ratio of the total variance contributed by the frequency band (a, b). Under the null of no serial correlation, $\mathcal{E}(a, b)$ is equal to the length of the interval (a, b) and any departure from this benchmark provides the means to detect serial correlation.

Although the variance ratio can be defined for an arbitrary frequency domain, the need to estimate the corresponding integral of the spectral density function – the numerator of $\mathcal{E}$ – imposes practical limitations. We resort to wavelet analysis to address this

* Corresponding author.
E-mail addresses: rgencay@sfu.ca (R. Gençay), dsignori@sfu.ca (D. Signori).

http://dx.doi.org/10.1016/j.jeconom.2014.08.002
0304-4076/

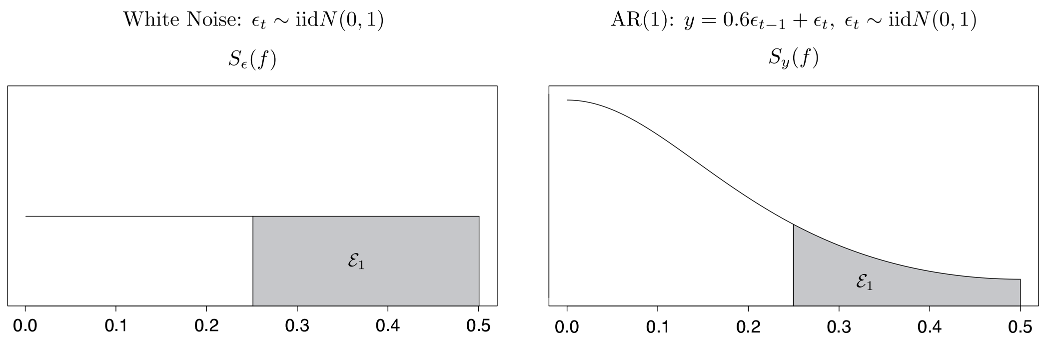

Fig. 1. High frequency contribution (in gray) to the total variance of a white noise process (left) and an AR(1) process (right).

need. For frequency bands of a particular form, the numerator of the statistic $\mathcal{E}$ is a well known quantity, the wavelet variance,[1] which can be estimated efficiently using the maximum-overlap discrete wavelet transformation estimator. In this light, given the temporal resolution properties of the wavelet transform, it is appropriate to refer to $\mathcal{E}(a, b)$ as a multiscale variance ratio. The recursive application of this procedure generates a family of tests whose joint limit distribution is multivariate normal under mild restrictions.

While the main intuition behind multiscale variance ratios originates under covariance stationarity assumptions, the corresponding test statistics are informative in more general scenarios. Indeed, the null hypothesis can be relaxed to allow for a degree of non-stationarity, specifically, for heteroskedastic white noise. Heteroskedastic white noise is an uncorrelated process with varying variance. We develop the asymptotic theory of multiscale variance ratios for uncorrelated but possibly dependent processes within the framework of *near-epoch dependence* (NED). Besides accommodating heterogeneity, there are three further benefits of this approach. Firstly, the asymptotic results originate from one of the most general Gaussian central limit theorems for dependent processes (De Jong, 1997). Secondly, it permits trending higher moments (see Assumptions A and B1). Finally, it leverages a rich literature devoted to the derivation of the NED property for many nonlinear time series models and, thus, parametric restrictions for the validity of our test can be obtained in several typical cases.[2]

We contribute to the literature on tests for serial correlations in several ways. First, the design we propose leads to serial correlation tests with desirable empirical size and power in small samples. Second, as argued in the previous paragraph, our test is robust to the presence of higher order dependence, heteroskedasticity, and trending moments, while at the same time the asymptotic theory is developed in great generality. Third, ours is the first test of serial correlation that directly utilizes the wavelet coefficients of the observed time series to construct the wavelet-based test statistics.[3] The tests we design generalize, on the one hand, variance ratios tests (Lo and MacKinlay, 1988), and on the other, they are related to ratios of quadratic forms and Von Neumann ratios (1941). In addition, since the proposed test statistic does not rely on a point estimate of the spectral density, the rate of convergence issues relating to the nonparametric spectral density are not of the first order of importance.

One of the well-known time-domain portmanteau tests for serial correlation is the Box and Pierce test Q_K (BP). Given independent and identically distributed observations, Box and Pierce (1970) show that the sum of K sample autocovariances times the number of observations is approximately distributed as a Chi-squared distribution with K degrees of freedom; statistically large values of Q_K indicate a likely serial correlation among the data. In practice, the strict restriction of independence and homogeneity is violated, possibly leading to a very inaccurate inference. There is a long streak of papers that address these limitations, starting from the small sample improvements of Ljung and Box (1978), to the more recent robustification program of Lobato (2001) and Lobato et al. (2002). Robust inference can also be achieved using bootstrapping methods. Building on the block bootstrap inference for autocorrelations of Romano and Thombs (1996), Horowitz et al. (2006) develop a blocks-of-blocks bootstrap that reduces the error rejection probability to nearly zero for samples with at least 500 observations. Finally, Escanciano and Lobato (2009) (EL) combine robustification techniques with a data-driven approach for automatic lag selection. The resulting adaptive test has particularly high empirical power in finite samples.

Frequency-domain tests provide an alternative framework for the tests of serial correlation. Hong (1996) uses a kernel estimator of the spectral density for testing serial correlation of arbitrary form. His procedure relies on a distance measure between two spectral densities of the data and the one under the null hypothesis of no serial correlation. Paparoditis (2000) proposes a test statistic based on the distance between a kernel estimator of the ratio between the true and the hypothesized spectral density and the expected value of the estimator under the null. Wavelet methods are particularly suitable in such situations where the data has jumps, kinks, seasonality and nonstationary features. The framework established by Lee and Hong (2001) is a wavelet-based test for serial correlation of unknown form that effectively takes into account local features, such as peaks and spikes in a spectral density. Duchesne (2006) extends the Lee and Hong (2001) framework to a multivariate time series setting. Hong and Kao (2004) extend the wavelet spectral framework to the panel regression. The simulation results of Lee and Hong (2001) and Duchesne (2006) indicate size over-rejections and modest power in small samples. Reliance on the estimation of the nonparametric spectral density together with the choice of the smoothing parameter affects their small sample performance. Recently, Duchesne et al. (2010) have made use of wavelet shrinkage (noise suppression) estimators to alleviate the sensitivity of the wavelet spectral tests to the choice of the resolution parameter. This framework requires a data-driven threshold choice and the empirical size may remain relatively far from the nominal size. Therefore, although a shrinkage framework provides some refinement, the reliance on the estimation of the

[1] The wavelet variance was studied, among others, by Allan (1966), Percival (1983), Percival and Guttorp (1994), Percival and Percival (1983), and Howe and Percival (1995).

[2] These results include GARCH, IGARCH, FIGARCH, ARCH(∞) (Davidson, 2004), ARMA, Bilinear models, switching and threshold autoregressive models, and smooth nonlinear autoregressions (Davidson, 2002).

[3] This approach was originally put forth by Fan and Gençay (2010) in unit root testing. Within a similar framework, Xue et al. (2014) propose discrete wavelet-based jump tests to detect jump arrival times in high frequency financial time series data.