Journal of Econometrics 191 (2016) 57–68

Contents lists available at ScienceDirect

Journal of Econometrics

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

Testing for (in)finite moments

Lorenzo Trapani *

Cass Business School, Faculty of Finance, 106 Bunhill Row, London EC1Y 8TZ, United Kingdom

ARTICLE INFO

Article history:
Received 29 July 2014
Received in revised form
14 May 2015
Accepted 31 August 2015
Available online 30 October 2015

JEL classification:
C12

Keywords:
Finite moments
Randomised tests
Chover-type Law of the Iterated Logarithm
Strong Law of Large Numbers

ABSTRACT

This paper proposes a test to verify whether the kth moment of a random variable is finite. We use the fact that, under general assumptions, sample moments either converge to a finite number or diverge to infinity according as the corresponding population moment is finite or not. Building on this, we propose a test for the null that the kth moment does not exist. Since, by construction, our test statistic diverges under the null and converges under the alternative, we propose a randomised testing procedure to discern between the two cases. We study the application of the test to raw data, and to regression residuals. Monte Carlo evidence shows that the test has the correct size and good power; the results are further illustrated through an application to financial data.

1. Introduction

An assumption common to virtually all studies in statistics and econometrics is that the moments of a random variable are finite up to a certain order. Existence of population moments is naturally required when computing sample moments. Moment restrictions are also routinely assumed in the various statements of the Law of Large Numbers (LLN) and of the Central Limit Theorem (CLT), thus playing a crucial role in estimation and testing—we refer to Davidson (2002), *inter alia*, for a comprehensive treatment of asymptotic theory. In addition to statistics and econometric theory, several applications in economics and finance require the calculation (and, therefore, the finiteness) of moments. However, a well-known stylised fact, e.g. when using high frequency financial data, is that heavy tails are often encountered (see e.g. Phillips and Loretan, 1994, and a recent contribution by Linton and Xiao, 2013; see also the references therein). Hence the importance of verifying whether assumptions on the finiteness of moments are satisfied.

In order to formally illustrate the problem, let X be a random variable with distribution $F(x)$, and consider the functional $\Xi_X^k(t) \equiv \int_{-t}^{a} |x|^k dF(x) + \int_a^t |x|^k dF(x)$, where $a \in (-t, t)$ is finite and the two integrals exist for any a. Then the raw absolute moment of order k is defined as

$$E|X|^k \equiv \mu_k = \lim_{t\to\infty} \Xi_X^k(t). \quad (1)$$

It is well known that, when the support of X is not bounded, the integral in (1) need not be finite, which entails that the kth moment (and of course also moments of order higher than k) does not exist. Testing procedures to check for the existence of moments are available, although not always employed. A typical approach (see e.g., in the context of testing for covariance stationarity, Phillips and Loretan, 1991, 1994, 1995) is based on estimating the so-called "tail index". This usually requires some assumptions on $F(x)$—typically, it is assumed that the tails of $F(x)$ can be approximated as $L(x)x^{-\gamma}$, where $L(x)$ is a slowly varying function. The parameter γ is referred to as the "tail index", and it is related to the highest finite moment of X—formally, this means that

$$\lim_{t\to\infty} \Xi_X^k(t) \begin{cases} =\infty \\ <\infty \end{cases} \quad \text{according as} \quad \begin{matrix} k \geq \gamma \\ k < \gamma. \end{matrix} \quad (2)$$

Hence, one could use an estimate of γ in order to test for the null hypothesis that $\gamma > k$, which is tantamount to testing for $H_0 : E|X|^k < \infty$. A routinely employed technique is the Hill estimator (Hill, 1975), or some variants thereof; we refer to Embrechts et al. (1997) and de Haan and Ferreira (2006) for excellent reviews which also consider several improvements of the original Hill estimator. In general, however, estimation of γ is fraught with difficulties. Considering the Hill estimator as a leading example, it is well known that its rate of convergence may be relatively slow: indeed, this is a common feature to all tail index estimators. Moreover, the quality of the Hill estimator depends crucially on selecting the appropriate number of order statistics—see Section 3.2, for details, and in particular the discussion after Eq. (21). If this is not chosen

* Tel.: +44 0 207 0405260.
E-mail address: L.Trapani@city.ac.uk.

http://dx.doi.org/10.1016/j.jeconom.2015.08.006
0304-4076/

correctly, the Hill estimator can yield very poor inference; Resnick (1997) provides an insightful discussion of the main pitfalls of the Hill estimator, and also several possible variants to overcome such pitfalls.

Hypotheses of interest and the main result of this paper

In this paper, we propose a test for the null that the kth raw moment of X does *not* exist; formally, we develop a test for

$$\begin{cases} H_0 : \lim_{t\to\infty} \Xi_X^k(t) = \infty \\ H_A : \lim_{t\to\infty} \Xi_X^k(t) < \infty. \end{cases} \tag{3}$$

We base our analysis on the divergent part of the Strong LLN (SLLN). Defining the kth sample moment, based on the sample $\{x_i\}_{i=1}^n$, as

$$\hat{\mu}_k \equiv \frac{1}{n}\sum_{i=1}^{n} |x_i|^k, \tag{4}$$

as $n \to \infty$ it holds that, almost surely

$$\hat{\mu}_k \to \begin{cases} \infty \\ <\infty \end{cases} \quad \text{according as} \quad \begin{matrix} \lim_{t\to\infty} \Xi_X^k(t) = \infty \\ \lim_{t\to\infty} \Xi_X^k(t) < \infty. \end{matrix} \tag{5}$$

Based on (5), we use $\hat{\mu}_k$ to test for $H_0 : \lim_{t\to\infty} \Xi_X^k(t) = \infty$ in (3).

The literature has proposed several contributions that use (5), both for the purpose of estimating γ and for conducting hypothesis testing. As far as the former is concerned, Meerschaert and Scheffler (1998; see also the related contribution by McElroy and Politis, 2007, and the references therein) exploit the generalised version of the CLT to propose a moment-based estimator of γ. As far as the latter issue (hypothesis testing) is concerned, Fedotenkov (2013; see also the related papers by Fedotenkov, 2015a,b) develops a bootstrap-based methodology whose main idea is closely related to the contribution of the present paper. In particular, Fedotenkov (2013) proposes comparing two statistics: the full-sample estimator of μ_k, and a subsample based one. Under the null hypothesis that μ_k is finite, both statistics would converge to μ_k by virtue of (5). Conversely, under the alternative that μ_k is not finite, the two statistics diverge at a different rate. Building on this, the test proposed by Fedotenkov (2013) is essentially based on comparing (by means of the bootstrap) the two statistics, checking whether their difference is bounded or diverges.

In the context of this paper, (5) is employed in order to test for the null hypothesis that μ_k does *not* exist. From a technical point of view, however, (5) is not used directly; rather, the main results in the paper hinge on a version of the Law of the Iterated Logarithm (LIL) for random variables that do not admit a finite first absolute moment, known in the literature as the "Chover-type LIL" (Chover, 1966). Thus, an ancillary contribution of this paper is the development of a Chover-type LIL for dependent data. From a methodological point of view, the results in this paper share, with the works cited above, the (desirable) feature of not having to determine an optimal number of order statistics to carry out inference, which is one of the main problems of the Hill estimator. However, note that, under the null hypothesis of an infinite k-order moment, there is no randomness in (5): the statistic $\hat{\mu}_k$ does not converge to any distribution (it diverges to positive infinity), and it cannot be used directly in order to conduct the test. Consequently, we employ a randomised testing procedure, which builds on a contribution by Pearson (1950). From a conceptual point of view, such approach is based on the idea that, when a statistic does not have randomness under the null (e.g. because it diverges) or when it has a non standard limiting distribution, randomness can be added by the researcher. Corradi and Swanson (2006) and Bandi and Corradi (2014) have recently employed randomised testing procedures. In particular, Bandi and Corradi (2014) propose a test to evaluate rates of divergence, which, albeit in a very different context, is essentially the same problem investigated in this paper. As far as conducting inference is concerned, we follow the approach used in Corradi and Swanson (2006), where randomisation is employed in conjunction with sample conditioning. This entails adding randomness to the basic statistic, and then deriving the asymptotics conditional on the sample, showing that limiting distribution and consistency hold for all samples save for a set of zero measure. Such approach is somehow akin to bootstrap based inference, which is also carried out conditional on the sample—although using bootstrap in this context would be problematic, e.g. due to the difficulties in extending the theory to the case of data with infinite first moment (see Cornea-Madeira and Davidson, 2014). A key difference with bootstrap-based inference is the interpretation that the notion of test size has in this context. Indeed, it is well known that, in a classical hypothesis testing context, the level α of a test means that, if a researcher applies the test B times and the null is valid, then (s)he will reject the null with frequency α—that is, (s)he will be wrong αB times. Conversely, as illustrated by Corradi and Swanson (2006), in this context α is interpreted thus: out of J researchers who apply the test, αJ of them will reject the null when this is true. Despite such interpretational difference, as we show in Section 2, using this approach we overcome the issue of $\hat{\mu}_k$ diverging under the null, and we obtain a test statistic which, for a given level α, rejects the null with probability α when true, and with probability 1 when false.

The remainder of the paper is organised as follows. In Section 2, we discuss the test, its theoretical properties (null distribution and consistency), and possible extensions to regression residuals (Section 2.1). Section 3 contains, in addition to a set of guidelines on how to use the test (Section 3.1), a Monte Carlo exercise (Section 3.2), and an application (Section 3.3). Section 4 concludes. Proofs are in the Appendix.

NOTATION We denote the ordinary limits as "$\to$"; convergence in distribution as "$\xrightarrow{d}$"; convergence in probability and almost surely as "$\xrightarrow{p}$" and "$\xrightarrow{a.s.}$" respectively. We use "a.s." as shorthand for "almost surely", "i.o." for "infinitely often", and "$\equiv$" for definitional equality. Finite constants that do not depend on the sample size are denoted as $M, M', \ldots$, etc. Other relevant notation is introduced in the remainder of the paper.

2. The test

This section contains a description of how the test statistic is constructed, and its theoretical properties (reported in Theorems 1 and 2). In Section 2.1, we study the application of the test to regression residuals.

We start by reporting the testing procedure as a four step algorithm.

Step 1 Compute $\hat{\mu}_k$.

Step 2 Randomly generate an *i.i.d.* $N(0,1)$ sample of size r, say $\{\xi_j\}_{j=1}^r$, and define the sample $\left\{\sqrt{e^{\hat{\mu}_k}} \times \xi_j\right\}_{j=1}^r$.

Step 3 Generate the sequence $\{\zeta_{j,n}(u)\}_{j=1}^r$ as

$$\zeta_{j,n}(u) \equiv I\left[\sqrt{e^{\hat{\mu}_k}} \times \xi_j \le u\right], \tag{6}$$

for all j, where $u \neq 0$ is any real number and $I[\cdot]$ is the indicator function. The values of u can be selected from a density $\varphi(u)$ on a bounded support $U = [\underline{u}, \bar{u}]$.

Step 4 For each $u \in U \setminus \{0\}$, define

$$\vartheta_{nr}(u) \equiv \frac{2}{\sqrt{r}} \sum_{j=1}^{r} \left[\zeta_{j,n}(u) - \frac{1}{2}\right], \tag{7}$$