Testing for (In)Finite Moments

Lorenzo Trapani

Cass Business School

October 4, 2017

Abstract

This paper proposes a test to verify whether the k -th 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 k -th 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.

JEL codes: C12.

Keywords: finite moments, randomised tests, Chover-type Law of the Iterated Logarithm, Strong Law of Large Numbers.

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 $\mathbf{X}$ be a random variable with distribution $F(\mathbf{x})$, and consider the functional $\frac{k}{X}(t) = \int_{-t}^t \mathbf{x}^k dF(\mathbf{x}) + \int_t^{\infty} \mathbf{x}^k dF(\mathbf{x})$, where $\mathbf{a} \in (-t; t)$ is finite and the two integrals exist for any $\mathbf{a}$. Then the raw absolute moment of order k is defined as

$$E|\mathbf{X}|^k = \lim_{t \rightarrow \infty} \frac{k}{X}(t) : \quad (1)$$

It is well known that, when the support of $\mathbf{X}$ is not bounded, the integral in (1) needs not be finite, which entails that the k -th 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 and 1995) is based on estimating the so-called tail index. This usually requires some assumptions on $F(\mathbf{x})$ - typically, it is assumed that the tails of $F(\mathbf{x})$ can be approximated as $L(\mathbf{x}) \mathbf{x}^{-\alpha}$, where $L(\mathbf{x})$ is a slowly varying function. The parameter α is referred to as the tail index, and

it is related to the highest finite moment of $\mathbf{X}$ - formally, this means that

$$\lim_{t \rightarrow \infty} \frac{E[\|\mathbf{X}\|^k(t)]}{t^k} = \begin{cases} \infty & \text{if } k > k_0 \\ 1 & \text{if } k = k_0 \\ 0 & \text{if } k < k_0 \end{cases} \quad \text{according as } \begin{cases} k > k_0 \\ k = k_0 \\ k < k_0 \end{cases} : \quad (2)$$

Hence, one could use an estimate of k_0 in order to test for the null hypothesis that $k_0 > k$, which is tantamount to testing for $H_0 : E\|\mathbf{X}\|^k < \infty$. A routinely employed technique is the Hill estimator (Hill, 1975), or some variants thereof; we refer to Embrechts, Kluppelberg and Mikosch (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 k_0 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 equation (21). If this is not chosen 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 k -th raw moment of $\mathbf{X}$ does *not* exist; formally, we develop a test for

$$\begin{aligned} & \begin{cases} \infty \\ > \\ < \end{cases} H_0 : \lim_{t \rightarrow \infty} \frac{E[\|\mathbf{X}\|^k(t)]}{t^k} = 1 \\ & \begin{cases} \infty \\ > \\ < \end{cases} H_A : \lim_{t \rightarrow \infty} \frac{E[\|\mathbf{X}\|^k(t)]}{t^k} < 1 \end{aligned} : \quad (3)$$

We base our analysis on the divergent part of the Strong LLN (SLLN). Defining the k -th sample moment, based on the sample $\{\mathbf{x}_i\}_{i=1}^n$, as

$$\hat{m}_k = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i\|^k; \quad (4)$$

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

$$\hat{\alpha}_k \begin{cases} > 1 \\ < 1 \end{cases} \quad \text{according as} \quad \begin{cases} \lim_{t \rightarrow \infty} \frac{k}{X}(t) = 1 \\ \lim_{t \rightarrow \infty} \frac{k}{X}(t) < 1 \end{cases} ; \quad (5)$$

Based on (5), we use $\hat{\alpha}_k$ to test for $H_0 : \lim_{t \rightarrow \infty} \frac{k}{X}(t) = 1$ in (3).

The literature has proposed several contributions that use (5), both for the purpose of estimating α and for conducting hypothesis testing. As far as the former is concerned, Meerschaert and Scheffer (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 α . As far as the latter issue (hypothesis testing) is concerned, Fedotenkov (2013; see also the related papers by Fedotenkov, 2015a and 2015b) 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 α , and a subsample based one. Under the null hypothesis that α is finite, both statistics would converge to α by virtue of (5). Conversely, under the alternative that α 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 α 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 $\hat{\mathbf{k}}$ 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 (2002, 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 $B\alpha$ times. Conversely, as illustrated by Corradi and Swanson (2006), in this context α is interpreted thus: out of J researchers who apply the test, $J\alpha$ 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 $\hat{\mathbf{k}}$ 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 Appendix.

NOTATION We denote the ordinary limits as $\lim$; 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 short-hand for almost surely, i.o. for infinitely often, and $\stackrel{\text{def}}{=}$ for definitional equality. Finite constants that do not depend on the sample size are denoted as $M, M^0, \dots$, 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 $\hat{\mathbf{e}}_k$.

Step 2 Randomly generate an *i.i.d.* $N(0;1)$ sample of size r , say $\mathbf{e}_j^r$, and define the sample $\mathbf{e}_k^{\wedge r} = \frac{1}{r} \sum_{j=1}^r \mathbf{e}_j^r$.

Step 3 Generate the sequence $\mathbf{e}_k^{\wedge r}(\mathbf{u})$ as

$$\mathbf{e}_k^{\wedge r}(\mathbf{u}) = \mathbf{I}^{\mathbf{e}_k^{\wedge r}} \mathbf{e}_k^{\wedge r} + \mathbf{I}^{\mathbf{u}} \mathbf{u}; \quad (6)$$

for all $\mathbf{e}_k^{\wedge r}$, where $\mathbf{u} \in \mathbb{R}$ is any real number and $\mathbf{I}[\cdot]$ is the indicator function. The values of $\mathbf{u}$ can be selected from a density $f(\mathbf{u})$ on a bounded support $U = [\underline{\mathbf{u}}, \bar{\mathbf{u}}]$.

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

$$\#_{nr}(\mathbf{u}) = \frac{2}{r} \sum_{j=1}^r \mathbf{e}_k^{\wedge r}(\mathbf{u}) = \frac{1}{r} \sum_{j=1}^r \mathbf{e}_k^{\wedge r}(\mathbf{u}); \quad (7)$$

and the test statistic

$$Z_{nr}^{\bar{u}} = \frac{\sum_{\mathbf{u}} \#_{nr}^2(\mathbf{u}) \rho'(\mathbf{u})}{\sum_{\mathbf{u}} \#_{nr}(\mathbf{u})} \quad (8)$$

The following remarks contain comments on the specifications of the test, and a heuristic preview of how the test works; the choice of the artificial sample size, r , defined in Step 2, is discussed after Theorems 1 and 2, and in Section 3.1.

Note also that the statistic $\hat{\epsilon}_k$ is bound to be sensitive to the unit of measurement; thus, a scale invariant transformation thereof should be employed instead - for the sake of a concise discussion, we assume henceforth that $\hat{\epsilon}_k$ is scale-free; in Section 3.1, we explore ways in which scale invariance can be obtained.

2.1 A heuristic description of the main idea of Step 3 is the following. Under the null hypothesis H_0 that $EjXj^k$ does not exist, as $n \rightarrow 1$, $\frac{1}{n} \sum_{j=1}^n \#_{j;n}(\mathbf{u})$ should follow a normal distribution with mean zero and finite variance. This entails that, under H_0 with $n \rightarrow 1$, for any real number $\mathbf{u}$, the random variable $\#_{j;n}(\mathbf{u})$ has a Bernoulli distribution with $P(\#_{j;n}(\mathbf{u}) = 1) = \frac{1}{2}$. Therefore, under H_0 as $n \rightarrow 1$, $\#_{j;n}(\mathbf{u})$ has mean $\frac{1}{2}$ and variance $\frac{1}{4}$. Conversely, under the alternative that $EjXj^k < 1$, $\hat{\epsilon}_k$ converges to a finite value. Hence, $\frac{1}{n} \sum_{j=1}^n \#_{j;n}(\mathbf{u})$ should follow a normal distribution with mean zero and finite variance, so that, for any $\mathbf{u} \neq 0$, $E \#_{j;n}(\mathbf{u}) \neq \frac{1}{2}$.

2.2 Step 4 follows directly from Step 3, and it is an application of the CLT. It can be expected that, under the null with $n \rightarrow 1$ and $r \rightarrow 1$, $\#_{nr}(\mathbf{u}) \xrightarrow{d} N(0;1)$ for every choice of $\mathbf{u}$. Conversely, under the alternative that $EjXj^k < 1$, the $\#_{j;n}(\mathbf{u})$ s do not have mean $\frac{1}{2}$ and therefore a CLT does not apply to (7). Given that in (7) there is a sum involving a sequence with non-zero mean, it can be expected that $\#_{nr}(\mathbf{u})$ diverges at a rate $P_{\bar{r}}$. This ensures the consistency of the test under H_A .

2.3 In Step 3, we consider the possibility that several values of $\mathbf{u}$ could be tried. The definition of $\#_{nr}$ in Step 4 is based on combining a continuous set of values of $\#_{nr}(\mathbf{u})$, attaching a different weight to each $\mathbf{u}$ according to some density $\rho'(\mathbf{u})$. Monte Carlo

evidence (Section 3.2) shows that choosing $U = f^{-1}; 1g$ with equal probability works well under any scenario. From a theoretical point of view, it can be expected that the width of U is positively related to both power and size: as it increases, the power versus the alternative that $EjXj^k < 1$ will increase, but there will also be some size distortion under the null.

We now lay out the main assumptions on dependence and tail behaviour. Prior to that, recall the definition of *uniform mixing* (see e.g. Davidson, 2002, p. 209). Let $(\Omega; F; P)$ denote the probability space on which $f_{X_i} g_{i=1}^n$ is defined, and let $F_j^{j+k} = (X_i; j = i, j+k)$ and the sets $A \in F_1^k$ and $B \in F_{k+m}^l$; finally, define

$$m = \sup_{k \in \mathbb{N}} jP(B|A) - P(A)j; \quad (9)$$

Then, $f_{X_i} g_{i=1}^n$ is said to be *uniformly mixing* if $m = 0$ as $m \rightarrow 1$.

Consider the following assumption.

Assumption 1. (i) $f_{X_i} g_{i=1}^n$ is a uniformly mixing sequence with mixing numbers $m = Mm$ for some $M > 0$; (ii) the X_i s have common distribution $F(x)$ which belongs in the domain of attraction of a stable distribution with tail exponent $\alpha \in (0; k]$, viz.

$$1 - F(x) = \frac{c_1(x) + o(1)}{x} L(x) \text{ and } F(-x) = \frac{c_2(x) + o(1)}{x} L(x); \quad (10)$$

as $x \rightarrow 1$, with: $L(x) > 0$ slowly varying at infinity in the Karamata sense, $c_i(x) > 0$ and $c_1 + c_2 > 0$ where $c_i = \lim_{x \rightarrow 1} c_i(x)$ for $i = 1, 2$.

Part (i) of the assumption imposes some structure on the dependence of the data - having $M > 0$ is a very mild requirement on the amount of memory allowed in the data. As mentioned in the Introduction, the main tool employed in the proofs - which is also one of the contributions of this paper - is a Chover-type LIL for uniformly mixing sequences. In addition to considering dependence (with a quite flexible amount of memory, since only $M > 0$ is required) uniform mixing affords great analytical tractability. Other forms of

dependence could also be considered, as long as Chover's LIL holds - see e.g. the results in Trapani (2014) for the case of strongly mixing data.

Part (ii) of the assumption contains the null hypothesis, represented by having k . The exact specification of the slowly varying function $L(x)$ is not required, and thus, basically, part (ii) is needed only to ensure that some moments of X exist.

Defining P as the probability law of $\sum_{j=1}^r \mathbf{J}_j(\mathbf{u})$ conditional on the sample, let $\xrightarrow{P}$ denote convergence in distribution according to P .

The limiting distribution of $\mathbf{nr}$ under the null is given in the following Theorem.

Theorem 1 *Let Assumption 1 hold. As $(n; r) \rightarrow \infty$ with*

$$\frac{r}{\exp(n^{\frac{k}{k-1}} \ln^{\frac{k}{k-1}} n)} \rightarrow 0; \quad (11)$$

it holds that $\mathbf{nr} \xrightarrow{P} \chi^2_1$ a.s. conditionally on the sample.

Theorem 1 stipulates that $\mathbf{nr}$ has, under the null, a chi-squared distribution with one degree of freedom, as can be expected from the discussion above. Convergence to the null limiting distribution requires both $n \rightarrow \infty$ and $r \rightarrow \infty$. As far as the latter is concerned, r needs to pass to infinity subject to (11), in order to ensure that a CLT holds. Since the test statistic is based on an *i.i.d.* sequence of uniformly distributed random variables, it can be expected that convergence should be quite fast, and therefore in practice r is not needed to be too large.

We now consider the consistency of the test versus the alternative that the k -th moment exists, i.e. $H_A : \lim_{t \rightarrow \infty} t^{-k} E(X^k) < 1$.

Theorem 2 *Let Assumption 1 hold with $k = k + \epsilon$ for some $\epsilon > 0$ in part (ii). Define c as $P[\mathbf{nr} \leq c] = \alpha$ under H_0 . Under H_A , as $(n; r) \rightarrow \infty$ with the restriction (11), it holds that $P[\mathbf{nr} > c] = 1$ a.s. conditionally on the sample if*

$$\lim_{(n;r) \rightarrow \infty} \frac{r}{\exp(n^{\frac{k}{k-1}} \ln^{\frac{k}{k-1}} n)} = 1; \quad (12)$$

Theorem 2 states that tests based on $\mathbf{r}$ have non trivial power versus the alternative that $Ej\mathbf{X}j^k$ exists. In the proof, we show that, under H_A , $\#_{\mathbf{r}}(\mathbf{u})$ has a non-centrality parameter proportional to $Q \frac{\mathbf{r}}{\exp(-k)}$, whence (12).

2.1 Application to regression residuals

We now turn to discussing the application of the test to regression residuals. Indeed, this is a natural application, since the existence of moments up to a certain order is routinely assumed for the error term in the regression model

$$\mathbf{y}_i = \beta_0 + \beta_1 \mathbf{x}_i + \epsilon_i \quad (13)$$

for example, a typical assumption in (13) is that the ϵ_i s have finite second moment, i.e. $Ej\epsilon_i^2 < \infty$. In order to verify the validity of this (and similar) assumptions, we show under which conditions the test developed above can be applied to the OLS residuals computed from (13).

Henceforth, we denote the distribution of ϵ_i as $F(\mathbf{x})$, and define the OLS residuals $\hat{\epsilon}_i$. Also, let $k \geq 1$ and $k(t) = \int_{-\infty}^t j\mathbf{x}j^k dF(\mathbf{x}) + \int_{-\infty}^t j\mathbf{x}j^k dF(\mathbf{x})$ with both integrals existing for any finite $a \in (-t; t)$; we propose to test for

$$\begin{aligned} & \geq H_0 : \lim_{t \rightarrow \infty} k(t) = \infty \\ & > H_A : \lim_{t \rightarrow \infty} k(t) < \infty \end{aligned} \quad (14)$$

by using

$$\hat{k} = \frac{1}{n} \sum_{i=1}^n j\hat{\epsilon}_i j^k;$$

or a scale invariant transformation. Using $P_{\hat{k}}$, we define the corresponding test statistics as $\mathbf{r}$.

Consider the following extension of Assumption 1.

Assumption 1*. (i) Assumption 1 holds for $f_{ig_{i=1}^n}$, with tail index α ; (ii) Assumption 1 holds for $f_{\mathbf{x}ig_{i=1}^n}$, with tail index $\alpha_{\mathbf{x}} = k + \alpha_{\mathbf{x}}$ for some $\alpha_{\mathbf{x}} > 0$; (iii) $f_{\mathbf{x}ig_{i=1}^n}$ and

$f_{i \in \mathcal{G}_{i=1}^n}$ are two mutually independent groups; (iv) $\mathbf{x}_i \notin 0$ a.s. for $1 \leq i \leq n$.

It holds that:

Corollary 1 *Let $k \geq 1$. Under H_0 , if Assumption 1*(i) holds, then as $(n; r) \rightarrow \infty$ with (11), it holds that $\sqrt{n} \hat{\beta} \xrightarrow{d} N(0, \Sigma)$. Under the alternative $\beta = \beta_0 + \delta$ with $\delta \neq 0$, if Assumption 1* holds, then, as $(n; r) \rightarrow \infty$ with (11) and (12), it holds that $P[\sqrt{n} \hat{\beta} > c] = 1$.*

Corollary 1 shows that the test can be extended to regression residuals, with the same null distribution and consistency properties as above. The only proviso is that the regressors have enough moments: in essence, the corollary requires that $Ej_{\mathbf{x}_i}^k < \infty$. As shown in the proof, this is needed in order for the test to have power: the regression residuals $\hat{\epsilon}_i$ contain $\mathbf{x}_i$, and therefore even when $Ej_{\mathbf{x}_i}^k < \infty$, if $Ej_{\mathbf{x}_i}^k = \infty$ this may entail that $\hat{\epsilon}_i$ diverges, thereby yielding zero power. From a methodological point of view, therefore, the test ought to be applied to residuals after checking whether $Ej_{\mathbf{x}_i}^k < \infty$ or not.

Note that the test is designed for $k \geq 1$ only. When $k < 1$, the test may still work, but this depends on the maximum moments of both $\mathbf{x}_i$ and ϵ_i ; the reason why the test may fail is that, under the alternative, the OLS estimator of β in (13) may be inconsistent, and even diverge to positive infinity - this is e.g. the case if $Ej_{\mathbf{x}_i}^{2+\delta} < \infty$ for some $\delta > 0$; see also Cline (1989).

3 Applying the test: guidelines, simulations and empirical applications

In this section, we study three inter-related issues concerning how to apply the test. Firstly, we discuss how to make the test statistic scale invariant - Section 3.1. Secondly, we evaluate the properties of our test through a simulation exercise, also proposing some guidelines as to how to choose some of the test specifications - Section 3.2. Finally, we illustrate the implementation of the test and the interpretation of its results by means of an application to financial data - Section 3.3.

3.1 Scale invariance

As pointed out in Section 2, the raw absolute moment of order k , $\hat{\mu}_k$, cannot be employed directly as it is not scale invariant. Hence, the need for a re-scaling of the statistic $\hat{\mu}_k$: although several alternatives can be proposed, as a general rule it is necessary to have a scaling factor that does not diverge (or diverges at a slower rate than $\hat{\mu}_k$) under the null, and is bounded under the alternative.

Several choices are possible; in particular, based on Lemma 1 in Appendix, we propose the following family of scaling factors

$$\hat{\mu}_k = \frac{1}{n} \sum_{i=1}^n |x_i|^k ; \quad (15)$$

where k is chosen such that $k > 0$; then, the re-scaled test statistic is

$$\hat{\mu}_k = \frac{\hat{\mu}_k}{\hat{\mu}_k} ; \quad (16)$$

The test could then be carried out exactly as in Section 2, using $\hat{\mu}_k$ instead of $\hat{\mu}_k$. Although in principle any choice of $k > 0$ can be employed, natural choices are $k = 1$ or $k = 2$ (i.e., using the variance of the data) - in the simulations, we employ $k = 2$, which proves to be a good choice in terms of empirical rejection frequencies; only minor changes are noted, anyway, when setting $k = 1$. Whilst $\hat{\mu}_k$ fulfills the purpose of making $\hat{\mu}_k$ scale invariant, in the simulations we also found that, at least in finite samples, better results are achieved by further rescaling $\hat{\mu}_k$ by the value it would have if $\mathbf{X}$ were normally distributed. Letting $\mu_k^{(N)}$ be the k -th absolute moment of a standard normal, the test statistic is thus based on

$$\hat{\mu}_k = \hat{\mu}_k \frac{\mu_k^{(N)}}{\mu_k^{(N)}} ; \quad (17)$$

The rationale for employing $\hat{\mu}_k$ can be illustrated as follows. Based on Lemma 1 in

Appendix, and modulo a slowly varying function, it holds that, under $\mathbf{H}_0$

$$\hat{\alpha} = \begin{matrix} \gtrsim \\ \lesssim \\ \gtrsim \end{matrix} \begin{matrix} \mathbf{O}_{a:s}(1) \\ \mathbf{O}_{a:s}(\mathbf{n}^{-1} \ln^{-1} \mathbf{n}) \end{matrix} \quad \text{according as} \quad \begin{matrix} < \\ > \end{matrix};$$

therefore

$$\hat{\alpha}^{k=} = \begin{matrix} \gtrsim \\ \lesssim \\ \gtrsim \end{matrix} \begin{matrix} \mathbf{O}_{a:s}(1) \\ \mathbf{O}_{a:s}(\mathbf{n}^{\frac{k}{k-1}} \ln^{\frac{k}{k-1}} \mathbf{n}) \end{matrix} \quad \text{according as} \quad \begin{matrix} < \\ > \end{matrix};$$

As can be seen, in either case the rate of divergence of the scaling factor $\hat{\alpha}^{k=}$ is always smaller than that of $\hat{\alpha}_k$. Since, under $\mathbf{H}_0$, $\hat{\alpha}_k = \mathbf{O}_{a:s}(\mathbf{n}^{\frac{k}{k-1}} \ln^{\frac{k}{k-1}} \mathbf{n})$, we have

$$\hat{\alpha}_k = \begin{matrix} \gtrsim \\ \lesssim \\ \gtrsim \end{matrix} \begin{matrix} \mathbf{O}_{a:s}(\mathbf{n}^{\frac{k}{k-1}} \ln^{\frac{k}{k-1}} \mathbf{n}) \\ \mathbf{O}_{a:s}(\mathbf{n}^{\frac{k}{k-1}} \ln^{\frac{k}{k-1}} \mathbf{n}) \end{matrix} \quad \text{according as} \quad \begin{matrix} < \\ > \end{matrix}; \quad (18)$$

and the same holds for $\hat{\alpha}_k$. Under the alternative, $\hat{\alpha}$ is bounded by the SLLN, so that $\hat{\alpha}_k$ is also bounded.

Finally, we return to the issue of selecting the size of the artificial sample, $\mathbf{r}$. Equations (11) and (18) can be combined in order to choose $\mathbf{r}$; after standard algebra, it can be verified that it must hold

$$\begin{matrix} \frac{\mathbf{r}}{\exp(\mathbf{n}^{\frac{k}{k-1}} \ln^{\frac{k}{k-1}} \mathbf{n})} \neq 0 \\ \frac{\mathbf{r}}{\exp(\mathbf{n}^{\frac{k}{k-1}} \ln^{\frac{k}{k-1}} \mathbf{n})} \neq 0 \end{matrix} \quad \text{according as} \quad \begin{matrix} < \\ > \end{matrix}; \quad (19)$$

When $k > 1$, any choice such that $\mathbf{r}$ is a polynomial function of $\mathbf{n}$ will be appropriate. The case $k = 1$ is of special interest: in such case, (19) boils down to requiring $\frac{\mathbf{r}}{\mathbf{n}} \neq 0$. Thus, choosing $\mathbf{r} = \mathbf{o}(\mathbf{n})$ always satisfies (19).

3.2 Monte Carlo simulations

We generate $n + 1000$ datapoints from an $\mathbf{ARMA}(1;1)$ process, discarding the first 1000 observations to avoid dependence on initial conditions:

$$\mathbf{y}_t = \mathbf{y}_{t-1} + \mathbf{a}_t + \mathbf{a}_{t-1}; \quad (20)$$

with $(\mathbf{a}; \mathbf{g}) = f(0;0);(0.5;0);(0;0.5);(0;-0.5)g$; the sample sizes considered are $n \in \{100;250;500;1000;10000;100000\}g$.

The innovation $\mathbf{a}_t$ is generated according to two schemes. In the first set of experiments, we generate data according to a Student t distribution with ν degrees of freedom ($\mathbf{t}$); such distribution is often found to be good at capturing the features of financial data (see e.g. Hurst and Platen, 1997; and Markowitz and Usmen, 1996a, 1996b), and therefore the results from this experiment should provide a set of guidelines for the applied user; indeed, we analyse the impact, on the test, of several specifications under this distributional design. We also consider a second set of experiments, where $\mathbf{a}_t$ is generated as having a power law, as a robustness check to assess how the test responds to a different distribution; data are generated according to standard procedures, and we refer to e.g. Clauset, Shalizi and Newman (2009).

As far as the testing problem is concerned, without loss of generality, we consider testing for the existence of the fourth moment, viz.

$$\begin{aligned} & \geq \mathbf{H}_0 : \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^t \mathbf{a}_i^4 = 1 \\ & \leq \mathbf{H}_A : \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^t \mathbf{a}_i^4 < 1 \end{aligned} :$$

Let ν denote the degree of freedom of the Student t distribution data are drawn from in the first set of experiments; and the tail index of the power law from which data are generated in the second set of experiments. We set $\nu \in \{2;3;4;5;6\}g$; the first three values are used to assess the size of the test, and the last two values are used in order to assess the power.

Based on Corollary 1, we apply the test to pre-whitened data, by fitting an $\mathbf{AR}(7)$

model to $\mathbf{y}_i$ and then applying the test to the residuals; unreported simulations show that, when applying the test directly to the raw data, this results in a massive oversize when there is dependence - hence, a guideline is that the test ought to be applied to pre-whitened data. In order to make the test statistic scale invariant, we use, as suggested in Section 3.1

$$\hat{\lambda}_4 = \frac{1}{3} \frac{\hat{\lambda}_4^4}{\hat{\lambda}_2^2}.$$

As far as the test specifications are concerned, on account of Theorems 1 and 2, it can be expected that the empirical rejection frequencies will be affected by the size of the artificial sample, $\mathbf{r}$, and by the values of $\mathbf{u}$ employed (that is, by the set $\mathbf{U}$). As regards the former, in the proof of Theorem 2 we show that, under the alternative, the test statistic has a noncentrality parameter proportional to $\frac{1}{\exp(\frac{\mathbf{r}}{k})}$: thus, a large $\mathbf{r}$ is bound to increase the power of the test. This is in a trade-off with (11), which indicates that a large value of $\mathbf{r}$ will yield size distortion; of course, this is valid for finite samples, since asymptotically the test will have the correct size as long as (11) is satisfied. Similarly, the width of $\mathbf{U}$ also has an impact on power and size. Indeed, as shown in the proof of Theorem 2, under the alternative the test has a noncentrality parameter which increases as the width of $\mathbf{U}$ increases: hence, using large values of $\mathbf{u}$ should boost the power, at the expense of size. In order to analyse the impact of $\mathbf{r}$ and $\mathbf{U}$ on the size and power of the test, we run four different experiments for the leading case of Student t data. In the first, benchmark case, we set $\mathbf{r} = \mathbf{n}^{\frac{4}{5}}$ and $\mathbf{U} = f(1;1g$, with each value drawn with equal probability of $\frac{1}{2}$; this proved to be the best choice for all cases considered. We also consider the cases $\mathbf{r} = \mathbf{n}^{\frac{1}{2}}$ and $\mathbf{U} = f(1;1g$, and $\mathbf{r} = \mathbf{n}^{\frac{4}{5}}$ and $\mathbf{U} = f(2;2g$. In the former case, the test is expected to be less powerful, whereas in the latter higher power should be observed, in presence of size distortion, at least for small samples. Finally, we also report the empirical rejection frequencies for the intermediate case $\mathbf{r} = \mathbf{n}^{\frac{1}{2}}$ and $\mathbf{U} = f(2;2g$. Results should look similar as $\mathbf{n} \rightarrow \infty$, and in this respect having $\mathbf{n} = 100000$ in the simulations should shed some light on the asymptotic performance of the test.

By way of comparison, we also report a set of experiments to determine the size and power of a test for $\mathbf{H}_0 : \lambda = 4$ based on a direct estimate of the tail index λ . Based on Hill

(1975), the estimator of α , say $\hat{\alpha}_{Hill}$, is calculated as

$$\hat{\alpha}_{Hill} = \frac{1}{h} \sum_{i=1}^h \ln \frac{\hat{\alpha}_{(n-h+i)}}{\hat{\alpha}_{(n-h)}}; \quad (21)$$

where $\hat{\alpha}_{(s)}$ denotes the s -th order statistic (in descending order) of the sample of the absolute values of the residuals from fitting an $AR(7)$ model to the y_s , i.e. from $\hat{f}^n_{ij} g^n_{i=1}$. The estimator is applied to the residuals of the $AR(7)$ model, rather than directly to the data: according to Embrechts, Kluppelberg and Mikosch (1997; see in particular Figure 5.5.4 on p. 270), this is an effective way of attenuating the impact of serial dependence in the data. As far as the threshold h is concerned, we use $h = n^{3/4} = \ln n$. We found this to be the best choice for h ; note that, in the second set of experiments, data are generated according to a strict Pareto model for the tails, and therefore in that case we can expect the Hill estimator to be unbiased. Alternatively, one could employ data-driven rules to select the optimal h - we refer to e.g. Drees and Kaufmann (1998) and the references therein. Since the data can be expected to be still dependent even after pre-whitening, we base the test on the statistic $P_{\hat{h}_{Hill}/s}^4$, where s^2 is the inverse of the long-run variance estimator proposed in Hill (2010); we refer thereto for details on the implementation, and point out here that we have employed the Bartlett kernel, and set the bandwidth to $b = n^{1/5}$, based on the condition $b = o(hn)^{1/2}$. Under the null, it holds that

$$P_{\hat{h}_{Hill}/s}^4 \xrightarrow{d} N(0;1); \quad (22)$$

We point out that the test based on $\hat{\alpha}_{Hill}$ is not meant to be the only possible alternative to our approach. Indeed, the performance of the Hill estimator can be quite poor, and various improvements have been suggested - see de Haan and Ferreira (2006). Rather, we would suggest to interpret the test based on $\hat{\alpha}_{Hill}$ as a naive benchmark. It is however worth noting that the applied literature customarily uses this approach to verify whether the data have finite moments of a certain order or not (e.g. Phillips and Loretan, 1991, 1994 and 1995).

Tables 1a-1c contains empirical rejection frequencies for the cases of data having a

Student t and a power law distribution, respectively. The number of simulations has been set equal to 1000, so that, when evaluating the size of the test, the empirical rejection frequencies should lie in the interval [0.036;0.064].

[Insert Tables 1a-1c somewhere here]

The effect of pre-whitening (in both cases: Student t and, as shown in Table 1c below, power law) is that the test is nearly unaffected by the presence of serial correlation: size and power almost do not change across the various combinations of $(\rho; \gamma)$.

As far as size is concerned, consider first the benchmark case where $\mathbf{r} = \mathbf{n}^{\frac{4}{5}}$ and $\mathbf{U} = f^{-1}; 1g$ - Table 1a. The test has the correct size in both non-boundary cases $\gamma = 2$ and $\gamma = 3$. As could be expected, the test exhibits higher empirical rejection frequencies when $\gamma = 4$; this, however, attenuates when $\mathbf{n} = 500$, with the test having the correct size again in all cases considered. Turning to the power (cases $\gamma = 5$ and $\gamma = 6$), this is higher than 50% for all cases considered when $\mathbf{n} = 500$, and anyway higher than the 50% threshold for $\mathbf{n} = 250$ in the non-boundary case $\gamma = 6$: the power increases monotonically with $\mathbf{n}$ and γ , both features being in line with what can be expected. As mentioned above, these specifications (for $\mathbf{r}$ and $\mathbf{U}$) correspond to the best results under all scenarios considered: thus, a guideline from Table 1a and, in general, from this section is to choose $\mathbf{r}$ quite close to $\mathbf{n}$, and use $\mathbf{U} = f^{-1}; 1g$.

The other cases, displayed in Table 1b, complement the conclusions above. Results were rather similar across the various combinations of $(\rho; \gamma)$, and therefore only the case with no serial dependence is reported for all three set-ups. The results in the table are in line with the theory: chiefly in the case $\gamma = 4$, decreasing $\mathbf{r}$ reduces the empirical rejection frequency. The power, on the other hand, is reduced quite substantially, and the test has power higher than 50% only for $\mathbf{n} = 10000$. Conversely, increasing the width of $\mathbf{U}$ boosts the power, but the test appears massively oversized, especially for the case of $\gamma = 4$, where the correct size is attained only when $\mathbf{n} = 100000$. The intermediate case $\mathbf{r} = \mathbf{n}^{\frac{1}{2}}$ and $\mathbf{U} = f^{-2}; 2g$ is the most similar to the case of $\mathbf{r} = \mathbf{n}^{\frac{4}{5}}$ and $\mathbf{U} \geq f^{-1}; 1g$, although the power is slightly lower and the size, when $\gamma = 4$, is never correct unless $\mathbf{n} = 10000$.

Turning to the case of data following a power law (Table 1c), there are few instances of oversize when n is small ($n = 100$) and α is equal to 3, but such tendency is relatively infrequent and it disappears for larger sample sizes. When $\alpha = 4$, the empirical rejection frequencies are higher than in the Student t case, with oversize attenuating when $n = 1000$. Further, as far as power is concerned, the test is less powerful, for large samples, than with Student t data. This can be further considered in conjunction with the high rejection rates when $\alpha = 4$.

Note, finally, that tests based on the Hill estimator have lower power, even for relatively large n - the power does increase above 50% for $n = 10000$ when data follow a Student t distribution, whereas it tends to be lower when the data follow a power law.

3.3 Application

In this section we illustrate how the test works through an empirical application to financial data. We consider daily returns from 3rd January, 2008, until 30th September, 2013, which corresponds to a sample size of $n = 1499$. We consider two groups of stocks from the FTSE 100: the banking sector (5 stocks) and the financial services sector (4 stocks), for a total of 9 stocks. A list of the constituents is in Tables 2a-2b.

We test for the existence of the first, second, third and fourth moment. In particular, letting $y_{j,i}$ be the return on stock j at day i , we use the following test statistics to verify the existence of the first four moments

$$\begin{aligned}\hat{\mu}_{1j} &= \frac{1}{n} \sum_{i=1}^n y_{j,i}; \\ \hat{\mu}_{2j} &= \frac{1}{n} \sum_{i=1}^n y_{j,i}^2 - \left(\hat{\mu}_{1j} \right)^2; \\ \hat{\mu}_{kj} &= \frac{1}{n} \sum_{i=1}^n y_{j,i}^k - \left(\hat{\mu}_{1j} \right)^k; \text{ for } k = 3, 4;\end{aligned}$$

Based on the results from the simulations, the test is applied to pre-whitened data, using $U = f^{-1}(g)$.

Results are reported in Tables 2a-2b:

[Insert Tables 2a-2b somewhere here]

Tables 2a-2b contain, in addition to the tests, the first four sample moments for each stock, showing large values of the kurtosis for all the stocks considered. We also report an estimate of the tail index based on (21). The point estimate, $\hat{Hill;j}$ is generally around 3, with some exceptions (Ashmore and Standard Chartered) for which it is higher. In view of the poor performance of tests based on $\hat{Hill;j}$ (as appears from the Monte Carlo exercise), we report $\hat{Hill;j}$ only as a preliminary indicator: its computation is conducted in the same way as in the Monte Carlo section. Based on the $\hat{Hill;j}$ s, it can be expected that all series have finite variances, whilst further testing is required for the third and fourth moments.

Tests are reported in the lower halves of Tables 2a and 2b. The tables show that, for all stocks considered, mean and variance exist, which is consistent with other studies in the literature (see e.g. Phillips and Loretan, 1991). As far as higher moments are concerned, all stocks in the banking sector have finite third (and, therefore, fourth) moment. Results are less clear-cut for the financial sector. Two stocks appear to have finite third moment (Aberdeen and Ashmore), although the null that the fourth moment is infinite is accepted. Indeed, in both cases there is not an overwhelming amount of evidence in favour of the null (e.g., in the case of Aberdeen, the null is accepted at 5% level, but it would be rejected at 10%). Heuristically, this is in line with the descriptive statistics: both stocks have kurtosis around 9, which (were one to assume a Student t distribution for the data) would correspond to a degree of freedom of 5, thus admitting finite fourth moments. The other stocks in the financial sector have the same behaviour as observed for the banking sector, namely infinite third moments. Again, these results are reinforced by the estimated values of the kurtosis, which are similar to the ones found in the banking sector.

4 Conclusions

This paper proposes a test for the null that the k -th moment of a random variable does not exist. The test uses the SLLN, which stipulates that sample moments diverge or converge according as their population counterparts are infinite or finite. Since, under the null, sample moments diverge to infinity and therefore have no randomness, we propose a randomised testing procedure. From a methodological point of view, this approach to testing for the finiteness of moments avoids having to estimate the tail index, which is known to be fraught with difficulties. Our simulations show that the test has the correct size and good power.

A natural question that stems from this paper is whether it is possible to derive an estimate of α . This contribution is focused on providing a test for the null that $Ej\mathbf{X}^k$ does not exist - this can be of relevance e.g. when computing descriptive statistics; or when employing a theory that requires the existence of moments up to a certain order. It would be possible to use a sequential approach, based on testing for the existence of consecutive, (possibly) integer values of k , although in general this methodology would have to take into account the risk of a high procedure-wise rejection frequency (see in particular a very insightful paper by Fedotenkov, 2015b). On the other hand, the approach proposed in this paper complements the estimator of α suggested by Meerschaert and Scheffer (1998), who propose

$$\frac{1}{\alpha} = \frac{\ln \prod_{i=1}^n \mathbf{x}_i^2}{2 \ln n},$$

showing that it is consistent, albeit at the slow rate $O_p \left(\frac{1}{\ln n} \right)$.

Finally, a word of warning on the meaning of the hypothesis testing framework is in order. Indeed, testing whether a quantity is passing to infinity, when samples are naturally finite, is bound to be conceptually unclear. In order to understand the rationale of the test, note that the null hypothesis is tested for by evaluating the rate of divergence of a sample moment (or, rather, of a scale invariant transformation thereof) - this lends itself to being put into an asymptotic setup. However, despite such asymptotic characterisation, the purpose of the analysis is to test for the *magnitude* of a sample moment, rather than

for its actual behaviour at infinity. In this respect, the approach suggested in this paper is strongly related to the contribution of Bandi and Corradi (2014), who also propose a test for rates of divergence: as the authors put it, evaluating magnitudes is essential to a variety of econometric problems. Thus, the purpose of this paper is to propose a procedure to allow the researcher to decide whether the moments of a random variable are small enough to be able to assume that the underlying distribution admits such moments, or not.

Acknowledgement

I am grateful to the Editor (Oliver Linton), one anonymous Associate Editor and two anonymous Referees for very constructive feedback which has greatly improved the generality of this paper. I also wish to thank the participants to the 5th International conference on Computational and Financial Econometrics (CFE 11; London, December 17-19, 2011), to the 12th OxMetrics User Conference (London, September 3-4, 2012), to the New York Camp Econometrics IX (Syracuse University, April 4-6, 2014), to the 6th ICEEE conference (University of Salerno, January 21-23, 2015) and to Lajos Horvath. The usual disclaimer applies.

Appendix: Proofs and Derivations

Recall that $\mathbf{P}$ is the probability law of the $\sum_{j=1}^r$ conditional on the sample; henceforth, we let $\mathbf{E}$ and $\mathbf{V}$ denote the expected value and the variance calculated with respect to $\mathbf{P}$.

We start with a preliminary Lemma, which contains a Law of the Iterated Logarithm (LIL).

Lemma 1. *Under Assumption 1, it holds that*

$$\limsup_{n \rightarrow \infty} \frac{1}{n^{1/k}} \sum_{i=1}^n j_{\mathbf{x}_i}^k = e^{1/k} \text{ a.s.}$$

Proof. The result in the lemma is a LIL for α -stable processes, and it is also known as a Chover-type LIL (Chover, 1966). Several results on Chover-type LILs are available in the literature (see e.g. Cai, 2006; Wu and Jiang, 2010; Trapani, 2014); thus, when possible, passages in the proof are omitted to save space.

Let $\mathbf{y}_i = j_{\mathbf{x}_i}^k$; on account of Assumption 1 (i) and of Theorem 14.1 in Davidson (2002, p. 210), the non negative sequence $\mathbf{y}_i$ is also uniformly mixing with mixing numbers of the same size. Define also $\mathbf{y}_i^{(n)} = \mathbf{y}_i \mathbf{I}[j_{\mathbf{x}_i} < \mathbf{a}_n]$, and $\mathbf{a}_n = [nL(n)f(n)]^{k=}$ with $f(n)$ a function such that

$$\limsup_{t \rightarrow \infty} \sup_{t^0 \leq t} \frac{f(t^0)}{f(t)} < 1;$$

Further, define $\mathbf{S}_j = \sum_{i=1}^j \mathbf{y}_i$ and $\mathbf{S}_j^{(n)} = \sum_{i=1}^j \mathbf{y}_i^{(n)}$ $\mathbf{E} \mathbf{y}_i^{(n)} = \mathbf{a}_n$.

We start by showing the upper half of the LIL, i.e.

$$\limsup_{n \rightarrow \infty} \frac{\sum_{i=1}^n \mathbf{y}_i}{n^{1/k}} = e^{(1+ \epsilon)/k} \text{ a.s.;} \quad (23)$$

for every $\epsilon > 0$. This requires showing that

$$\sum_{n=1}^{\infty} \frac{1}{n} \mathbf{P} \left(\max_{1 \leq j \leq n} j \mathbf{S}_j > \mathbf{a}_n \right) < 1; \quad (24)$$

for some $\epsilon > 0$. The passages are very similar e.g. to those in the proof of Theorem 2.3 in Cai (2006). Indeed

$$\mathbb{P} \left(\max_{1 \leq j \leq n} j S_j > \epsilon a_n \right) \leq \mathbb{P} \left(\max_{1 \leq j \leq n} y_j > \epsilon a_n \right) + \mathbb{P} \left(\max_{1 \leq j \leq n} S_j^{(n)} > \epsilon a_n \right) \leq \max_{1 \leq j \leq n} \frac{1}{n} \sum_{i=1}^n \mathbb{E} y_i^{(n)} : \quad (25)$$

Under Assumption 1, as $n \rightarrow \infty$ we have

$$\frac{1}{a_n} \max_{1 \leq j \leq n} \frac{1}{n} \sum_{i=1}^n \mathbb{E} y_i^{(n)} = 0; \quad (26)$$

the proof is in Cai (2006). Combining (25) and (26)

$$\sum_{n=1}^{\infty} \frac{1}{n} \mathbb{P} \left(\max_{1 \leq j \leq n} j S_j > \epsilon a_n \right) \leq \sum_{n=1}^{\infty} \frac{1}{n} \mathbb{P} [y_j > \epsilon a_n] + \sum_{n=1}^{\infty} \frac{1}{n} \mathbb{P} \left(\max_{1 \leq j \leq n} S_j^{(n)} > \epsilon a_n \right);$$

for some $0 < \epsilon < \epsilon$ and n large enough. Assumption 1(ii) entails that $\sum_{j=1}^n \frac{1}{n} \mathbb{P} [y_j > \epsilon a_n] < 1$. Finally, consider

$$\sum_{n=1}^{\infty} \frac{1}{n} \mathbb{P} \left(\max_{1 \leq j \leq n} S_j^{(n)} > \epsilon a_n \right) \leq \mathbb{E} \sum_{n=1}^{\infty} \frac{1}{n a_n^2} \mathbb{E} \left(\max_{1 \leq j \leq n} S_j^{(n)} \right)^2 \leq \frac{1}{a_n^2} \mathbb{E} \sum_{n=1}^{\infty} \frac{1}{n a_n^2} \sum_{i=1}^n \mathbb{E} y_i^{(n)2};$$

which follows from Markov's inequality and Rosenthal's inequality (see Lemma 1.3 and Theorem 2.1 in Xuejun, Shuhe, Yan and Wenzhi, 2009). Hence, exploiting the definition of a_n

$$\begin{aligned} & \sum_{n=1}^{\infty} \frac{1}{n} \mathbb{P} \left(\max_{1 \leq j \leq n} S_j^{(n)} > \epsilon a_n \right) \\ & \leq \mathbb{E} \sum_{n=1}^{\infty} \frac{1}{n a_n^2} \mathbb{E} y_i^{(n)2} = \mathbb{E} \sum_{n=1}^{\infty} \frac{1}{n a_n^2} \int_{j x_j}^{\infty} \int_{a_n^{1=k}}^{\infty} j x_j^{2k} dF(x) \\ & = \mathbb{E} \sum_{n=1}^{\infty} \sum_{h=1}^{\infty} \frac{1}{n a_n^2} \int_{j x_j}^{\infty} \int_{a_h^{1=k} < j x_j}^{\infty} j x_j^{2k} dF(x) \leq \mathbb{E} \sum_{h=1}^{\infty} \frac{1}{a_h^2} \int_{a_h^{1=k} < j x_j}^{\infty} \int_{a_h^{1=k}}^{\infty} dF(x) \leq \sum_{h=1}^{\infty} \frac{1}{a_h^2} \\ & \leq \sum_{h=1}^{\infty} \frac{1}{a_h^2} \int_{a_h^{1=k} < j x_j}^{\infty} \int_{a_h^{1=k}}^{\infty} dF(x) \leq \sum_{h=1}^{\infty} \frac{1}{a_h^2} \int_{a_h^{1=k}}^{\infty} \frac{dx}{x f(x)} < 1; \end{aligned}$$

This completes the proof of (24). Equation (23) follows from (24) upon choosing $\mathbf{f}_n = \ln^{1+\alpha} n$ - see e.g. the proof of Corollary 2.4 in Cai (2006).

We now turn to the lower half of the LIL, i.e.

$$\limsup_{n \rightarrow \infty} \frac{\sum_{i=1}^n \mathbf{y}_i}{n^{k/2}} = e^{(1-\alpha)k/2} \text{ a.s.}, \quad (27)$$

for any $\alpha \in (0; 1)$. This requires showing that if, for some sequence $\mathbf{h}_n$ of positive numbers, $\sum_{n=1}^{\infty} \mathbf{P}[\mathbf{y}_n > \mathbf{h}_n] < \infty$, then

$$\limsup_{n \rightarrow \infty} \frac{\sum_{i=1}^n \mathbf{y}_i}{\mathbf{h}_n} = 1 \text{ a.s.}; \quad (28)$$

for any non decreasing sequence $\mathbf{h}_n$ of positive numbers. To show this, recall that the sequence $\mathbf{f}_{\mathbf{y}_i}$ is uniformly mixing of the same size as $\mathbf{x}_i$. Hence, the second Borel-Cantelli Lemma holds - see Lemma 1.1.2 in Iosifescu and Theodorescu (1969). Thus, a 0/1 law can be shown, yielding $\mathbf{P}[\mathbf{y}_n > \mathbf{h}_n \text{ i.o.}] = 0$ or 1 according as $\sum_{n=1}^{\infty} \mathbf{P}[\mathbf{y}_n > \mathbf{h}_n] < \infty$ or $= \infty$; see e.g. Lemma 4(ii) in Wu and Jiang (2010). Hence, equation (28) can be shown by contradiction. Assuming that (28) does not hold when $\sum_{n=1}^{\infty} \mathbf{P}[\mathbf{y}_n > \mathbf{h}_n] < \infty$, this entails stating that there is a $\mathbf{d}_0 \in (0; +\infty)$ such that $\limsup_{n \rightarrow \infty} \mathbf{h}_n^{-1} \sum_{i=1}^n \mathbf{y}_i = \mathbf{d}_0$ almost surely. Clearly $\limsup_{n \rightarrow \infty} \mathbf{h}_n^{-1} \sum_{i=1}^n \mathbf{y}_i = (\mathbf{d}_n - \mathbf{d}_{n-1}) \mathbf{h}_n + \limsup_{n \rightarrow \infty} \mathbf{h}_n^{-1} \sum_{i=1}^{n-1} \mathbf{y}_i = \mathbf{d}_0$ almost surely. However, since $\mathbf{h}_n^{-1} \mathbf{y}_n \xrightarrow{P} 0$, we also have $\limsup_{n \rightarrow \infty} \mathbf{h}_n^{-1} \sum_{i=1}^n \mathbf{y}_i = 0$, so that $\limsup_{n \rightarrow \infty} \mathbf{h}_n^{-1} \sum_{i=1}^n \mathbf{y}_i = 2\mathbf{d}_0$ a.s.; hence, $\mathbf{P}[\mathbf{y}_n > M\mathbf{h}_n \text{ i.o.}] = 0$ for some $M > 2\mathbf{d}_0 < \infty$. By the Borel-Cantelli Lemma, this entails that $\sum_{n=1}^{\infty} \mathbf{P}[\mathbf{y}_n > M\mathbf{h}_n] < \infty$. But this contradicts the initial statement. Thus, (28) holds for any sequence $\mathbf{h}_n$. Now, given a non decreasing sequence of positive numbers $\mathbf{f}_n$, by Assumption 1(ii) we have $\sum_{n=1}^{\infty} \mathbf{P}[\mathbf{y}_n > (\mathbf{nf}_n)^{k/2}] < \infty$, as long as $\sum_{n=1}^{\infty} (\mathbf{nf}_n)^{-1} = \infty$. In such case, it follows that $\limsup_{n \rightarrow \infty} (\mathbf{nf}_n)^{-k/2} \sum_{i=1}^n \mathbf{y}_i = 1$ a.s.; hence, equation (27) can be proved following exactly the same passages as in the proof of Theorem 2 in Wu and Jiang (2010), setting $\mathbf{f}_n = \ln^{1+\alpha} n$. Combining (23) and (27), the Lemma follows. QED

Proof of Theorem 1. We prove the theorem for a more general set-up, defining $Q(\wedge_k)$ be a continuous, positive, monotonically increasing transformation of $\wedge_k$, with

$\lim_{t \rightarrow 1} Q(t) = +1$. By Lemma 1, under H_0 we have $\hat{\alpha}_k = O_p(n^{-1/2} \ln^{k-1} n)$; thus, by continuity, $Q(\hat{\alpha}_k) = O_p(n^{-1/2} \ln^{k-1} n)$. Henceforth, the proof of Theorem 1 is based on a similar logic to the proof of Theorem 5 in Bandi and Corradi (2014), and we only report the main passages. Let $\mathbf{g}_{j;n} = Q^{1/2}(\hat{\alpha}_k) \mathbf{g}_j$; conditional on the sample, $\mathbf{g}_{j;n} \sim N[0; Q(\hat{\alpha}_k)]$. Define now the set

$$A_n(\eta) = \left\{ \mathbf{g}_{j;n} : \frac{Q(\hat{\alpha}_k)}{n^{k-1}} > \eta > 0 \right\};$$

for any $\eta > 0$. Lemma 1 entails that, under H_0 , $P[\lim_{n \rightarrow \infty} A_n(\eta) = 1] = 1$. All the passages below are reported conditional on $A_n(\eta)$. For each $\mathbf{u}$ we have

$$\begin{aligned} \#_{nr}(\mathbf{u}) &= \sum_{\mathbf{r}} \sum_{j=1}^K \mathbf{g}_{j;n}(\mathbf{u})' \mathbf{E} \mathbf{g}_{j;n}(\mathbf{u}) + \sum_{\mathbf{r}} \sum_{j=1}^K \mathbf{E} \mathbf{g}_{j;n}(\mathbf{u})' \mathbf{g}_{j;n}(\mathbf{u}) \\ &= \mathbf{I} + \mathbf{II}; \end{aligned} \quad (29)$$

with

$$\mathbf{E} \mathbf{g}_{j;n}(\mathbf{u}) = \frac{1}{2} + \mathbb{P} \frac{1}{2 \overline{T(\hat{\alpha}_k)}} \int_0^{Z^u} \exp \left\{ -\frac{1}{2} \frac{t^2}{Q(\hat{\alpha}_k)} \right\} dt; \quad (30)$$

where we consider the case of $\mathbf{u} > 0$ without loss of generality. Consider first $\mathbf{II}$ in (29); based on (30), we have

$$\begin{aligned} \mathbf{II} &= \sum_{\mathbf{r}} \sum_{j=1}^K \mathbb{P} \frac{1}{2 \overline{T(\hat{\alpha}_k)}} \int_0^{Z^u} \exp \left\{ -\frac{1}{2} \frac{t^2}{Q(\hat{\alpha}_k)} \right\} dt \\ &\quad - \sum_{\mathbf{r}} \sum_{j=1}^K \mathbb{P} \frac{1}{2 \overline{T(\hat{\alpha}_k)}} \int_0^{Z^u} \exp \left\{ -\frac{1}{2} \frac{s^2}{Q(\hat{\alpha}_k)} \right\} ds - \mathbf{M} \frac{1}{2 \overline{T(\hat{\alpha}_k)}} \frac{\mathbf{r}}{Q(\hat{\alpha}_k)}; \end{aligned} \quad (31)$$

where the third inequality is based on Taylor's expansion around $\mathbb{P} \frac{1}{2 \overline{T(\hat{\alpha}_k)}}$. By Lemma 1, it follows that $\mathbf{II}$ in (29) is $O_p(n^{-1/2} \ln^{k-1} n)$; this is $o_p(1)$ uniformly in $\mathbf{u}$ if (11) holds.

Turning to $\mathbf{I}$ in (29), note that $\mathbf{g}_{j;n}(\mathbf{u})' \mathbf{g}_{j;n}(\mathbf{u}) = \mathbf{E} \mathbf{g}_{j;n}(\mathbf{u})' \mathbf{g}_{j;n}(\mathbf{u})$ is an *i.i.d.* sequence

with mean zero; it holds that

$$\begin{aligned} V_{\mathbf{r}}^{-2} \sum_{j=1}^2 \mathbf{X}_{j;n}(\mathbf{u})^2 &= E \left[\sum_{j=1}^2 \mathbf{X}_{j;n}(\mathbf{u})^2 \right] = \sum_{j=1}^2 E \left[\mathbf{X}_{j;n}(\mathbf{u})^2 \right] \\ &= 1 + O_p \left(\frac{1}{n^{k-1}} \right); \end{aligned}$$

where the second equality comes from the fact that $\mathbf{X}_{j;n}(\mathbf{u})$ is generated independently across j , and the last equality comes from (30) and the passages thereafter. This holds uniformly in $\mathbf{u}$ by the same passages as above. Thus, a CLT can be applied to $\mathbf{I}$, so that, as $(n; \mathbf{r}) \rightarrow 1$, $\mathbf{I} \xrightarrow{d} N(0; 1)$. Putting everything together, as $(n; \mathbf{r}) \rightarrow 1$ with (11), $\#_{nr}(\mathbf{u}) \xrightarrow{d} N(0; 1)$ uniformly in $\mathbf{u}$. This entails that $\lim_{n; \mathbf{r} \rightarrow 1} \int_{\mathbb{R}^u} \#_{nr}(\mathbf{u})^2 d\mathbf{u} = \int_{\mathbb{R}^u} 1 d\mathbf{u} = 1$. QED

Proof of Theorem 2. Under Assumption 1, with part (ii) modified so as to allow for $k > 0$, it holds that $E[jY_j \ln jY_j] < 1$. Since uniform mixing implies strong mixing, we can apply a SLLN for strong mixing sequences (e.g. Rio, 1995): thus, $\frac{1}{n} \sum_{j=1}^n jY_j \xrightarrow{a.s.} k < 1$, so that $Q(\hat{\lambda}_k) \xrightarrow{a.s.} Q(k) < 1$. Similarly to the proof of Theorem 1, define

$$\tau_n^+ = \inf \{ t : Q(\hat{\lambda}_k) < M < 1 \};$$

such that under H_A we have $P[\lim_{n \rightarrow \infty} \tau_n^+ = 1] = 1$. All the passages below are reported conditional on $\tau_n^+ \geq \tau_n^+$.

Consider (29). Term $\mathbf{I}$ still satisfies a CLT by construction, so that, under H_A , $\mathbf{I} \xrightarrow{d} N(0; 1)$. As far as $\mathbf{II}$ in (29) is concerned, by (30) we have (considering the case of $\mathbf{u} > 0$)

$$\mathbf{II} = 2^{P_{\mathbf{r}}} \frac{1}{\Gamma(\frac{1}{2})} \exp \left(-\frac{1}{2} \frac{\mathbf{t}^2}{Q(\hat{\lambda}_k)} \right) \frac{d\mathbf{t}}{d\mathbf{t}} \frac{\mathbf{r}}{Q(\hat{\lambda}_k)}; \quad (32)$$

Note that the non-centrality parameter $\mathbf{II}$ increases with the width of $\mathbf{U}$. These passages entail that $\#_{nr}(\mathbf{u})$ has a non-centrality parameter proportional to $\frac{1}{Q(\hat{\lambda}_k)}$; therefore, $\#_{nr}$ has a noncentrality parameter that diverges under (12), giving the desired result. QED

Proof of Corollary 1. The proof is based on combining the proofs of Theorems 1 and 2. Define the n -dimensional vectors $[1; \dots; n]^0$ and $[\wedge_1; \dots; \wedge_n]^0$. By construction, $\wedge = M$, with $M = I_n - W(W^0 W)^{-1} W^0$, with I_n the n -dimensional identity matrix and W the design matrix.

In order to study the behaviour under the null, it suffices to show that when $Ej_1^k = 1$, Lemma 1 holds for $\wedge_k$. Let $\|h\|_k$ denote the L_k -norm of an n -dimensional vector h i.e. $\|h\|_k = \left(\sum_{j=1}^n |h_j|^k \right)^{1/k}$; we can write $\wedge_k = n^{-1} k \wedge_k^k = n^{-1} k M^k k_k^k$. Based on Lemma 2.2 in Grcar (2003), there exists a finite and strictly positive constant, say C_M , such that $n^{-1} k M^k k_k^k \leq C_M n^{-1} k k_k^k$ so that $\wedge_k \leq C_M n^{-1} k k_k^k$. By applying Lemma 1, $n^{-1} k k_k^k = O_{a.s.}(n^{-k-1} \ln^k n)$ under H_0 . The same arguments as in the proof of Theorem 1 yield the asymptotics under the null of $\wedge$.

As far as the behaviour of $\wedge$ under alternatives is concerned, it suffices to show that, under H_A , $\wedge_k$ converges to a finite limit as $n \rightarrow \infty$. Let $\wedge = \sum_{i=1}^n x_i^2 = \sum_{i=1}^n x_i i$; by applying the C_r -inequality we have

$$\frac{1}{n} \sum_{i=1}^n j_i^k \leq 2^{k-1} \frac{1}{n} \sum_{i=1}^n j_i^k + 2^{k-1} \wedge^k \leq \frac{1}{n} \sum_{i=1}^n j_i^k = 2^{k-1} (I + II);$$

As far as I is concerned, using Assumption 1*(i) under the alternative, the SLLN entails that $n^{-1} k k_k^k$ converges a.s. to a finite limit, so that $I = O_{a.s.}(1)$. Further

$$II = \sum_{i=1}^n \frac{1}{n} x_i^2 \sum_{i=1}^n \frac{1}{n} x_i i \sum_{i=1}^n \frac{1}{n} j_i^k = \sum_{i=1}^n \frac{1}{n} j_i^k \sum_{i=1}^n \frac{1}{n} x_i i \sum_{i=1}^n \frac{1}{n} j_i^k = M^0 \sum_{i=1}^n \frac{1}{n} j_i^k \sum_{i=1}^n \frac{1}{n} j_i^k;$$

note that it is not important whether $n^{-1} \sum_{i=1}^n x_i^2$ converges or not in order to prove that II is bounded: what matters is that it is bounded away from zero - part (iii) of Assumption 1* rules this out. By virtue of Assumption 1*(ii), the SLLN yields $n^{-1} \sum_{i=1}^n j_i^k = O_{a.s.}(1)$. Consider the remaining term; by virtue of the independence between x_i and i , the sequence $f_{x_i} g_{i=1}^n$ is also ϕ -mixing (see Theorem 5.2 in Bradley, 2005). Under the alternative that $\wedge > k$, and on account of Assumption 1*(ii), the SLLN can be employed

again, yielding $n^{-1} \sum_{i=1}^n j(\mathbf{x}_i) = O_{a:s}(1)$; this is not necessarily the sharpest bound, but it suffices for our purposes. Hence, $\mathbf{H} = O_{a:s}(1)$; this yields that, under the alternative, $\hat{\mathbf{k}} = O_{a:s}(1)$. The proof henceforth is the same as in Theorem 2. QED

References

- Bandi, F., Corradi, V., 2014. Nonparametric nonstationarity tests. *Econometric Theory*, 30, 127–149.
- Cai, G.-H., 2006. Chover-type laws of the iterated logarithm for weighted sums of α -mixing sequences. *Journal of Applied Mathematics and Stochastic Analysis*, 2006:65023.
- Chover, J., 1966. A law of the iterated logarithm for stable summands. *Proceedings of the American Mathematical Society*, 17, 441–443.
- Clauset, A., Shalizi, C.R., Newman, M.E.J., 2009. Power-law distributions in empirical data. *SIAM Review*, 51, 661–703.
- Cline, D.B.H., 1982. Consistency for least squares regression estimators with infinite variance data. *Journal of Statistical Planning and Inference*, 23, 163–179.
- Cornea-Madeira A., Davidson, R., 2014. A parametric bootstrap for heavy-tailed distributions. *Econometric Theory*, 31, 1–22.
- Corradi, V., Swanson, N.R., 2006. The effects of data transformation on common cycle, cointegration, and unit root tests: Monte Carlo and a simple test. *Journal of Econometrics*, 132, 195–229.
- Davidson, J., 2002. *Stochastic limit theory*. Oxford: Oxford University Press.
- Drees, H., Kaufmann, E., 1998. Selecting the optimal sample fraction in univariate extreme value estimation. *Stochastic Processes and Applications*, 75, 149–172.
- Embrechts, P., Kluppelberg, C., Mikosch, T., 1997. *Modelling extremal events for insurance and finance*. New York: Springer.
- Fedotenkov I., 2013. A bootstrap method to test for the existence of finite moments. *Journal of Nonparametric Statistics*, 25, 315–322.
- Fedotenkov I., 2015a, A simple nonparametric test for the existence of finite moments, *Mimeo*.

Fedotenkov I., 2015b. Bootstrap for testing the existence of finite moments: a caution for possible misapplication, Mimeo.

Grcar, J.F., 2003. A matrix lower bound. Lawrence Berkeley National Laboratory report LBNL-50635.

de Haan, L., Ferreira, A., 2006. Extreme value theory: an introduction. New York: Springer.

Hill, B.M., 1975. A simple general approach to inference about the tail of a distribution. The Annals of Statistics, 3, 1163–1174.

Hill, J.B., 2010. On tail index estimation for dependent, heterogeneous data. Econometric Theory, 26, 1398–1436.

Hurst, S. R., Platen, E. 1997. The marginal distributions of returns and volatility. In Y. Dodge (Ed.), L_1 -Statistical Procedures and Related Topics, Volume 31 of IMS Lecture Notes - Monograph Series, pp. 301-314. Institute of Mathematical Statistics Hayward, California.

Iosifescu, M., Theodorescu, R., 1969. Random Processes and Learning. Berlin: Springer.

Linton, O.B., Xiao, Z., 2013. Estimation of and inference about the expected shortfall for time series with infinite variance. Econometric Theory, 29, 771–807.

Markowitz, H., Usmen, N. 1996a. The likelihood of various stock market return distributions, Part 1: Principles of inference. Journal of Risk & Uncertainty, 13, 207-219.

Markowitz, H., Usmen, N. 1996b. The likelihood of various stock market return distributions, Part 2: Empirical results. Journal of Risk & Uncertainty, 13, 221-247.

McElroy T., Politis D.N., 2007. Moment-based tail index estimation. Journal of Statistical Planning and Inference, 137, 1389–1406.

Meerschaert M.M., Scheffer H-P., 1998. A simple robust estimation method for the thickness of heavy tails. Journal of Statistical Planning and Inference, 71, 19–34.

Pearson, E.S., 1950. On questions raised by the combination of tests based on discontinuous distributions. Biometrika, 37, 383–398.

Phillips, P.C.B., Loretan, M., 1991. The Durbin-Watson ratio under infinite variance

errors. *Journal of Econometrics*, 47, 85–114.

Phillips, P.C.B., Loretan, M., 1994. Testing the covariance stationarity of heavy-tailed time series. *Journal of Empirical Finance*, 1, 211–248.

Phillips, P.C.B., Loretan, M., 1995. On the theory of testing covariance stationarity under moment condition failure. In G. S. Maddala, P.C.B. Phillips and T.N. Srinivasan (eds.), *Advances in Econometrics and Quantitative Economics: Essays in Honor of Professor C.R. Rao*. Basil Blackwell.

Resnick, S.I., 1997. Heavy tail modeling and teletraffic data. *Annals of Statistics*, 25, 1805–1869.

Rio, E., 1995. A maximal inequality and dependent Marcinkiewicz-Zygmund strong laws. *Annals of Probability*, 23, 918–937.

Trapani, L., 2014. Chover-type Laws of the k -Iterated Logarithm for strongly mixing sequences. *Journal of Mathematical Analysis and Applications*, 420, 908–916.

Xuejun, W., Shuhe, H., Yan, S., Wenzhi, Y., 2009. Moment inequality for α -mixing sequences and its applications. *Journal of Inequalities and Applications*, 2009:379743.

Wu, Q., Jiang, Y., 2010. Chover-type laws of the k -iterated logarithm for α -mixing sequences of random variables. *Journal of Mathematical Analysis and Applications*, 366, 435–443.

n	(;)	2	3	4	(0; 0)	5	6
	<i>Test speci cs</i>						
100	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 1; 1g$	0:031	0:026	0:027	0:077	0:104	
	$r = n^{\underline{4=5}}; U \mathcal{Z} f \ 2; 2g$	0:088	0:124	0:254	0:444	0:493	
	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 2; 2g$	0:031	0:048	0:069	0:307	0:412	
250	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 1; 1g$	0:039	0:039	0:044	0:232	0:351	
	$r = n^{\underline{4=5}}; U \mathcal{Z} f \ 2; 2g$	0:059	0:094	0:123	0:587	0:716	
	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 2; 2g$	0:039	0:046	0:082	0:426	0:464	
500	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 1; 1g$	0:060	0:049	0:062	0:367	0:432	
	$r = n^{\underline{4=5}}; U \mathcal{Z} f \ 2; 2g$	0:062	0:082	0:097	0:696	0:841	
	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 2; 2g$	0:061	0:054	0:086	0:473	0:597	
1000	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 1; 1g$	0:059	0:061	0:070	0:395	0:526	
	$r = n^{\underline{4=5}}; U \mathcal{Z} f \ 2; 2g$	0:041	0:076	0:091	0:783	0:895	
	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 2; 2g$	0:059	0:061	0:072	0:565	0:668	
10000	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 1; 1g$	0:032	0:051	0:055	0:512	0:651	
	$r = n^{\underline{4=5}}; U \mathcal{Z} f \ 2; 2g$	0:042	0:057	0:079	0:898	1:000	
	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 2; 2g$	0:032	0:051	0:063	0:584	0:680	
100000	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 1; 1g$	0:055	0:048	0:054	0:584	0:874	
	$r = n^{\underline{4=5}}; U \mathcal{Z} f \ 2; 2g$	0:055	0:055	0:061	1:000	1:000	
	$r = n^{\underline{1=2}}; U \mathcal{Z} f \ 2; 2g$	0:055	0:048	0:045	0:687	0:995	

Table 1b. Empirical rejection frequencies at a nominal 5% value; data are generated as Student t with degrees of freedom for the case (;) = (0; 0). For each sample size, the entries refer to: (a) the case where $r = n^{\underline{1=2}}$ and $U = f \ 1; 1g$; (b) the case $r = n^{\underline{4=5}}$ and $U = f \ 2; 2g$; (c) the case $r = n^{\underline{1=2}}$ and $U = f \ 2; 2g$.

QTOLandscapepenlandscape	n	(:)	(0; 0)						(0.5; 0)					
			2	3	4	5	6	2	3	4	5	6	2	
	100	Test speci cs												
		PRE_{Hill}^W	0.062 0.031	0.065 0.051	0.154 0.064	0.394 0.072	0.416 0.067	0.058 0.032	0.063 0.043	0.146 0.066	0.413 0.055	0.423 0.073	0.06 0.03	
	250	PRE_{Hill}^W	0.063 0.028	0.061 0.039	0.126 0.059	0.497 0.083	0.561 0.081	0.061 0.022	0.060 0.048	0.132 0.057	0.516 0.071	0.577 0.095	0.06 0.02	
		500	PRE_{Hill}^W	0.050 0.032	0.049 0.060	0.098 0.083	0.527 0.116	0.613 0.124	0.057 0.033	0.062 0.056	0.105 0.082	0.541 0.122	0.625 0.115	0.05 0.03
	1000		PRE_{Hill}^W	0.044 0.028	0.058 0.060	0.068 0.090	0.626 0.142	0.783 0.154	0.039 0.028	0.059 0.059	0.071 0.105	0.634 0.136	0.794 0.150	0.04 0.03
		10000	PRE_{Hill}^W	0.043 0.014	0.041 0.071	0.054 0.159	0.687 0.244	0.836 0.303	0.043 0.015	0.041 0.062	0.056 0.147	0.708 0.229	0.852 0.324	0.04 0.01
	100000		PRE_{Hill}^W	0.043 0.004	0.040 0.074	0.052 0.213	0.760 0.377	0.903 0.534	0.042 0.006	0.041 0.066	0.058 0.206	0.798 0.375	0.911 0.560	0.04 0.00

Table 1c. Empirical rejection frequencies at a nominal 5% value; data are generated from a power law with tail index .

Descriptive Statistics	Company	Barclays	HSBC	Lloyds	RBS	Std. Chart.
Mean		0:0004	0:0001	0:0008	0:0016	0:0000
Std. Dev.		0:0418	0:0216	0:0438	0:0516	0:0297
Skewness		1:2506	0:3093	0:8238	6:7949	0:3129
Kurtosis		28:675	14:691	24:156	145:216	13:395
$\hat{\alpha}_{Hill}$		3:008	3:486	3:546	2:256	6:226
Tests						
1		2:61 10 ² [0:000]	1:40 10 ² [0:000]	2:70 10 ² [0:000]	2:87 10 ² [0:000]	1:90 10 ² [0:000]
2		17:46 10 ⁴ [0:000]	4:68 10 ⁴ [0:000]	19:22 10 ⁴ [0:000]	26:66 10 ⁴ [0:001]	8:81 10 ⁴ [0:000]
3		277:27 10 ⁶ [0:316]	30:24 10 ⁶ [0:107]	318:74 10 ⁶ [0:225]	1152:87 10 ⁶ [0:554]	77:78 10 ⁶ [0:139]
4		8736:86 10 ⁸ [0:485]	321:95 10 ⁸ [0:233]	8945:16 10 ⁸ [0:555]	103630:23 10 ⁸ [0:420]	1040:76 10 ⁸ [0:597]

Table 2a. Results for companies in the banking sector - the stocks considered are, respectively, Barclays plc, HSBC Holdings plc, Lloyds Banking Group plc, The Royal Bank of Scotland Group plc, Standard Chartered plc. For each stock, we report the mean, the standard deviation (computed after adjusting for the degree of freedom), the skewness and the kurtosis (both unadjusted). Hill estimator of the tail index has been computed based on (21), using the same specifications as in the Monte Carlo section. In the Tests section, we report the values taken by the k -th absolute sample moment; number in square brackets are p-values.

	Company	Aberdeen	Ashmore	ICAP	Schroders
Descriptive Statistics					
Mean		0:0001	0:0002	0:0004	0:0005
Std. Dev.		0:0238	0:0264	0:0304	0:0268
Skewness		0:3224	0:1938	0:0941	0:0649
Kurtosis		8:946	9:179	13:467	22:357
$\hat{\alpha}_{Hill}$		3:393	5:434	3:043	2:472
Tests					
1		1:63 10^{-2} [0:000]	1:75 10^{-2} [0:000]	1:99 10^{-2} [0:000]	1:75 10^{-2} [0:000]
2		5:65 10^{-4} [0:000]	6:98 10^{-4} [0:000]	9:22 10^{-4} [0:000]	7:15 10^{-4} [0:000]
3		32:93 10^{-6} [0:034]	46:81 10^{-6} [0:041]	81:21 10^{-6} [0:282]	64:52 10^{-6} [0:284]
4		286:11 10^{-8} [0:098]	446:85 10^{-8} [0:111]	1146:78 10^{-8} [0:424]	1143:70 10^{-8} [0:707]

Table 2b. Results for companies in the financial services sector - the stocks considered are, respectively, Aberdeen Asset Management plc, Ashmore Group plc, ICAP plc and Schroders plc. For each stock, we report the mean, the standard deviation (computed after adjusting for the degree of freedom), the skewness and the kurtosis (both unadjusted). Hill estimator of the tail index has been computed based on (21), using the same specifications as in the Monte Carlo section. In the Tests section, we report the k -th absolute sample moment; number in square brackets are p-values.