A simple nonparametric test for the existence of finite moments

Fedotenkov, Igor

Bank of Lithuania

13 August 2015

A simple nonparametric test for the existence of finite moments

Igor Fedotenkov [1]

August 2015

Abstract

This paper proposes a simple, fast and direct nonparametric test to verify if a sample is drawn from a distribution with a finite first moment. The method can also be applied to test for the existence of finite moments of another order by taking the sample to the corresponding power. The test is based on the difference in the asymptotic behaviour of the arithmetic mean between cases when the underlying probability function either has or does not have a finite first moment. Test consistency is proved; then, test performance is illustrated with Monte-Carlo simulations and a practical application for the S&P500 index.

JEL classification: C12, C14
MSC2010 classification: 62G10, 62G20, 62G32, 65C05

Keywords: Heavy tails, tail index, finite moment, test, consistency.

[1] Bank of Lithuania.
Totoriu 4, 01121 Vilnius, Lithuania
e-mail: I.Fedotenkov: i.fedotenkov@gmail.com, ifedotenkov@lb.lt

1 Introduction

The use of heavy-tailed distributions is common in many fields of science to describe data containing unusually large observations. They are often applied in telecommunications and seismology (Neves and Alves 2008; Felgueiras 2012), insurance (McNeil 1997), economics, finance (Jansen and de Vries 1991), physics (Reed and Jorgensen 2004) and many other fields. Sometimes such data does not have a finite fourth, third, second or even first moment. In this last case even an arithmetic mean is meaningless. Many statistical methods, such as those based on the law of large numbers or central limit theorem, cannot be applied to such data, since they assume a certain number of finite moments. Therefore, there is a need for methods testing the hypothesis that a certain moment is finite. Despite the numerous methods developed for this end, a simple, fast and direct method to ascertain if the moment of interest is finite is obviously needed.

The direct way of checking if a sample is drawn from a distribution with finite moment of order k, is to plot the $k-$th empirical moment as a function of a number of observations used for its calculation. If the $k-$th finite moment exists, the empirical moment usually tends to converge to it; otherwise, the behaviour of the graph is unstable. An early example of this method was offered by Mandelbrot (1963), who applied it to cotton price analysis. However, the method is not formal and largely depends on the intuition of the researcher.

The most popular method for checking the hypothesis that a finite moment of a certain order exists, is the estimation of a tail index with the Hill estimator. This method employs the assumption of a regularity of the distribution's tail: $Pr(X > x) \sim x^{-\gamma}$, $\gamma > 0$, for large x. The tail index γ is estimated and the hypothesis about it is tested.[2] Tail index gives an inference about the number of finite moments, since the moments of an order smaller than γ are finite. The method was introduced by Hill (1975) and developed in subsequent copious literature. For example, Deheuvels, Haeusler, and Mason (1988) found strong consistency of the Hill estimator for independent data under rather weak assumptions. Resnick and Stărică (1998) found estimator consistency for dependent stationary sequences; furthermore, asymptotic normality was also proved (Haeusler and Teugels 1985). However, it is not often clear how many observations should be treated as the tail - too few observations yield a large variance, and too many observations may result in a large bias. Danielsson et al. (2001) applied bootstrap to determine the optimal number of observations for the estimation of the tail index. However, bootstrap requires the choice of a bootstrap resample size, which is rather arbitrary. Alternatively, Hill plots are often depicted with the number of observations treated as a tail on the horizontal axis, and the estimates of the tail index made with the Hill estimator on the vertical axis. The region of stability of the graph corresponds to the tail index of the distribution and makes an inference about the number of finite moments. A modification of this method with a logarithmic horizontal axis was proposed by Resnick and Stărică (1997) and discussed in detail by Drees et al. (2000). How-

[2] In fact, Hill estimator estimates $1/\gamma$.

ever, these methods also have weaknesses: the graphs may not have a region of stability; the need for researcher intervention to analyze the graphs is not suitable for a large number of tasks including Monte-Carlo simulations. Our formal and simple test for checking a hypothesis about finite moments solves a few of the problems.

Apart from the Hill estimator, numerous tail index and moment estimators have been developed. The most prominent are the Pickands and DEdH estimators (Pickands 1975; Dekkers, Einmahl, and de Haan 1989). There are also a number of other more recent estimators, for instance, the tail index estimator proposed by Alves et al. (2009) who present their estimator as a tool to check for superheavy tails (where the tail index is equal to zero), because the estimator does not impose a restriction that the tail index must be only positive. Interesting tail index estimators were also proposed by Meerschaert and Scheffler (1998) and McElroy and Politis (2007). They use the fact that if the second moment does not exist, the second empirical moment tends to diverge to infinity. Our proposed method for testing the hypothesis that the finite moments exist is similar to the estimators of Meerschaert and Scheffler, and McElroy and Politis in that it employs asymptotic behavior of the first empirical moment (arithmetic mean). However, it skips the step of the tail index estimation and checks the hypothesis that a finite moment of a specific order exists directly.

A direct bootstrap-based method for testing if the first finite moment exists was developed by Fedotenkov (2013), using the fact that under some general conditions, if the first moment does not exist, the arithmetic mean of the sample diverges to infinity faster than the arithmetic means of the subsamples of a smaller size. Following the method of Fedotenkov (2013), the test proposed in this paper is based on the strong law of large numbers, when the first moment does not exist as developed by Derman and Robbins (1955). The advantage of our proposal is that it is incomparably faster - the time gained being proportional to the number of bootstrap subsamples.

In the next section, the test is introduced and the main concepts are defined. In section 3 consistency of the test is shown. Section 4 presents the results of Monte-Carlo experiments and a practical application for the S&P500 index. Section 5 concludes.

2 The test

Let X, $X > 0$ be a random variable with a distribution $\mathcal{F}$. Suppose, $(X_1, X_2, ..., X_n)$ is a random i.i.d. sample drawn from a distribution $\mathcal{F}$, where n is a number of observations. The test is designed to check if the first moment of X is finite. If a researcher wishes to check if the moment of order p is finite, the test may be applied for the sample raised to the corresponding power, i.e. $(X_1^p, X_2^p, ..., X_n^p)$.

Hypothesis:

- H_0: $\mathcal{F}$ has a finite first moment;

- H_1: The first moment is infinite.

Denote $\hat{\mathcal{F}}_n(x) = \sum_{i=1}^n I(X_i < x)/n$, where I is a unity indicator function, which is equal to one when the condition $X_i < x$ is fulfilled, and zero otherwise. Denote the mean of the distribution $\mathcal{F}$ as μ. As we supposed that the observations can only take positive values, $\mu > 0$, which can also be infinite. Because of the properties of the distribution function $Pr(X > \mu) = 1 - \mathcal{F}(\mu) > 0$ if $\mu < \infty$ and $Pr(X > \mu) = 0$ otherwise. The theoretical distribution function and mean can be substituted with their empirical equivalents: $\hat{\mathcal{F}}_n$ and $\bar{X}$. Furthermore, to ensure correct asymptotic behavior, the argument of the empirical distribution function needs to be adjusted: $1 - \hat{\mathcal{F}}(\bar{X}/g(n))$, where $g(n)$ is a function defined below. Comparing this value with a tolerance level α the finite mean hypothesis can be tested.

Define $g(n)$ as a function exhibiting such properties:

- $\lim_{n\to\infty} g(n) = \infty$;

- $\lim_{n\to\infty} g(n)n^{-\rho} = 0$ for each $\rho > 0$.

Examples of such a function could be a logarithmic function, double and triple logarithms. In section 4 we perform simulations with $g(n) = \sqrt{\log\log n}$, and show that it gives reasonable results. As $g(n)$ is used to reduce the argument of the empirical distribution function, for practical reasons $g(n) \geq 1$ shall be used. However, this assumption is not important for the asymptotic properties of the test.

Formally, the testing procedure is:

1. Choose a tolerance level α, $0 < \alpha < 1$;
2. Calculate the test
$$P_n = \hat{\mathcal{F}}_n\left(\frac{\bar{X}_n}{g(n)}\right);$$
3. Accept H_0 if $P_n < 1 - \alpha$ and reject it otherwise.

The term "tolerance level" used here is similar to the notion of a significance level, however, the usage of the term "significance level" in this case would not be correct since for each α, $0 < \alpha < 1$, sample size n, and function $g(n)$ such highly-skewed probability functions satisfying H_0 can be constructed that the probability to reject H_0 is larger than α. Choosing the tolerance level α a researcher can control for the set of the distributions (s)he sacrifices, which likewise works for the choice of significance level. Furthermore, as it is shown in section 4, tolerance level sometimes can also be interpreted as an upper bound for first order mistakes, similar to the significance level.

Let us look what happens if $\mathcal{F}$ is a distribution function with a finite mean:

$$P_n \xrightarrow[n\to\infty]{a.s.} \mathcal{F}(0) = 0.$$

When $\mathcal{F}$ has a finite mean $\bar{X}_n \xrightarrow{a.s.} \mu$ as $n \to \infty$, according to Kolmogorov's law of large numbers, $\lim_{n\to\infty} g(n) = \infty$, $\hat{\mathcal{F}}_n \xrightarrow{a.s.} \mathcal{F}$ according to the Glivenko-Cantelli lemma, and $\mathcal{F}(0) = 0$ because X can take positive values only.

3 Consistency

Theorem 1. *Suppose the existence of constants $C > 0$ and $0 < \beta < 1$, such that $F(X) < 1 - C/X^{\beta}$ for sufficiently large X. Then $P_n \xrightarrow{a.s.} 1$ as $n \to \infty$.*

Proof. Similarly to Derman and Robbins (1955):

$$Pr\Big(\sum_{i=1}^{n} X_i \leq ng^2(n)\Big) \leq Pr\Big(max(X_1, ..., X_n) \leq ng^2(n)\Big) =$$

$$F\big(ng^2(n)\big)^n \leq \left(1 - \frac{C}{(ng^2(n))^{\beta}}\right)^n \leq \exp\left(-\frac{n^{1-\beta}C}{g^{2\beta}(n)}\right).$$

Because of the properties of function $g(n)$, it is easy to verify that

$$\sum_{n=N}^{\infty} \exp\left(-\frac{n^{1-\beta}C}{g^{2\beta}(n)}\right) < \infty.$$

Thus, the Borel-Cantelli lemma implies that

$$Pr\Big(\sum_{i=1}^{n} X_i \leq ng^2(n)\Big) = 0$$

for sufficiently large n. In other words, for sufficiently large n

$$Pr\left(\frac{\bar{X}_n}{g(n)} > g(n)\right) = 1.$$

As $n \to \infty$, $g(n) \to \infty$, thus $\bar{X}_i/g(n) \xrightarrow{a.s.} \infty$. Therefore, according to the Glivenko-Cantelli lemma and the definition of a distribution function

$$\hat{\mathcal{F}}_n\left(\frac{\bar{X}_n}{g(n)}\right) \xrightarrow[n\to\infty]{a.s.} \mathcal{F}(\infty) = 1.$$

This proves the theorem. □

The condition of the theorem does not cover all possible distributions without a finite first moment; for example, distributions with a tail index equal to unity do not satisfy the conditions of the theorem. In fact, the test can be inconsistent in this case, but the rate of H_0 rejection may still be greater than the tolerance level.

The weakest point of this testing procedure is that a selection of the tolerance level is rather arbitrary, and it does not have such an intuitive interpretation as a significance level. I hope to address this question in more detail in the following works. One possible solution for this weakness can be a randomization of the test statistics, first introduced by Corradi and Swanson (2002), and widely used nowadays. Despite such a modification would make the choice of the significance level more plausible I expect that it shall reduce the power of the test, since apart of the randomness existing in the sample itself, another source of randomness is introduced. In the next section, it is shown that the choice of the tolerance level employed in this paper provides reasonable results.

4 Monte-Carlo simulations and an example for the stock exchange

Monte-Carlo experiments were performed to examine the power of the test. As previously mentioned, the function $g(n)$ used in simulations was chosen as $\sqrt{\log\log n}$. Figure 1 presents the results of Monte-Carlo simulations for samples i.i.d. drawn from a log-logistic probability function. The scale parameter was set to unity. The horizontal axis corresponds to the shape parameter of the log-logistic distribution, which is responsible for the heaviness of the tails, and conforms to the tail index. Shape parameters were taken in the interval [0.02, 2] with the grid equal to 0.02. For each shape parameter 10^5 Monte-Carlo simulations were performed. The vertical axis shows the share of H_0 rejected, i.e. the power of the test when the tail index is smaller or equal to unity, and the rate of the first order mistakes if the tail index is greater than unity. Three curves are depicted for the sample sizes $n = 100$, $n = 1000$ and $n = 10.000$. As seen in the figure, the power of the test increases as n grows.

Figure 1: Power of the test and the rate of the first order mistakes

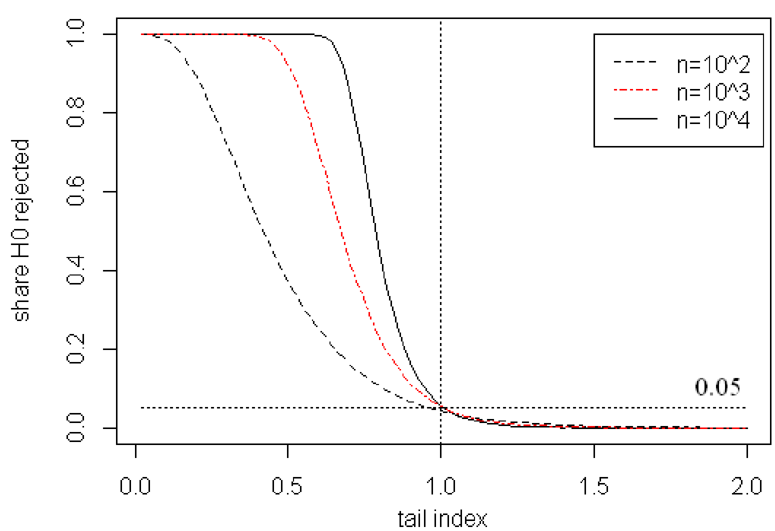

It can also be noted from Figure 1 that the rate of first order mistakes is close to the tolerance level (equal 0.05) only when the shape parameter of the log-logistic distribution is close to 1. Undoubtedly, this property depends on the choice of the scaling function $g(n)$, and may not be valid with other scaling functions. For the larger values of the tail index the rate of the first order mistakes is much lower than the tolerance level, and converges to zero as n grows.

4.1 Example for the S&P500 index

As a practical example the S&P500 index was considered for the period from 3 January 2007 till 11 December 2012. The absolute values of daily returns were calculated as $|100(Z_t - Z_{t-1})/Z_{t-1}|$, where Z_t is the S&P500 index value on day t at the closure of the stock exchange. Next, the test was applied. The hypotheses that 1-3 moments are finite were accepted at all the reasonable tolerance levels. The hypothesis that the fourth finite moment exists is accepted at 5% but rejected at 10% tolerance level, and the hypothesis that the fifth moment is finite is rejected at 5% tolerance level. This is in line with the findings of Cont (2001) who analyzed stylized empirical facts of price variations in various types of financial markets and argued that often the behaviour of the fourth moment is rather erratic. A similar result was also reported by Pagan (1996) who found evidence for the existence of the second finite moment in the US stock exchange daily data, but argues that the fourth moment may not exist.

5 Conclusion

This paper proposed a simple nonparametric test for checking if a sample is drawn from a distribution with the finite mean, and showed its consistency. The test makes use of the fact that when the first moment is finite, an arithmetic mean converges to it, and if there is no finite first moment the behaviour of the arithmetic mean becomes unstable, and under some weak assumptions it diverges to infinity.

Acknowledgements

I would like to thank Marco Minozzo and Remigijus Leipus for useful comments and Steve Guttenberg for help with the English.

References

Alves, I. F., L. de Haan, and C. Neves (2009). A test procedure for detecting super-heavy tails. *Journal of Statistical Planning and Inference 139*(2), 213–227.

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance 1*(2), 223–236.

Corradi, V. and N. R. Swanson (2002). A consistent test for nonlinear out of sample predictive accuracy. *Journal of Econometrics 110*, 353–381.

Danielsson, J., L. de Haan, L. Peng, and C. G. de Vries (2001). Using a bootstrap method to choose the sample fraction in tail index estimation. *Journal of Multivariate Analysis 76*(2), 226–248.

Deheuvels, P., E. Haeusler, and M. D. Mason (1988). Almost sure convergence of the Hill estimator. *Mathematical Proceedings of the Cambridge Philosophical Society 104*(2), 371–381.

Dekkers, A. L. M., J. H. J. Einmahl, and L. de Haan (1989). A moment estimator for the index of an extreme-value distribution. *The Annals of Statistics 17*(4), 209–222.

Derman, C. and H. Robbins (1955). The strong law of large numbers when the first moment does not exist. *Proceedings of the National Academy of Sciences 41*(8), 586–587.

Drees, H., S. Resnick, and L. de Haan (2000). How to make a Hill plot. *The Annals of Statistics 28*(1), 254–274.

Fedotenkov, I. (2013). A bootstrap method to test for the existence of finite moments. *Journal of Nonparametic Statistics 25*(2), 315–322.

Felgueiras, M. M. (2012). Explaining the seismic moment of large earthquakes by heavy and extremely heavy tailed models. *GEM - International Journal on Geomathematics 3*(2), 209–222.

Haeusler, E. and J. L. Teugels (1985). On asymptotic normality of Hill's estimator for the exponent of regular variation. *The Annals of Statistics 13*(2), 743–756.

Hill, B. (1975). A simple general approach to inference about the tail of a distribution. *The Annals of Statistics 3*(5), 1163–1174.

Jansen, D. and C. G. de Vries (1991). On the frequency of large stock returns: putting booms and busts into perspective. *The Review of Economics and Statistics 73*(1), 18–24.

Mandelbrot, B. (1963). The variation of certain speculative prices. *The Journal of Business 36*(4), 394–419.

McElroy, T. and D. N. Politis (2007). Moment-based tail index estimation. *Journal of Statistical Planning and Inference 137*(4), 1389–1406.

McNeil, A. (1997). Estimating the tails of loss severity distributions using extreme value theory. *Astin Bulletin 27*(1), 125–139.

Meerschaert, M. M. and H. P. Scheffler (1998). A simple robust estimation method for the thickness of heavy tails. *Journal of Statistical Planning and Inference 71*(1-2), 19 – 34.

Neves, C. and I. F. Alves (2008). Ratio of maximum to the sum for testing super heavy tails. Advances in Mathematical and Statistical Modeling, Statistics for Industry and Technology, 181-194.

Pagan, A. (1996). The econometrics of financial markets. *Journal of Empirical Finance 3*(1), 15–102.

Pickands, J. (1975). Statistical inference using extreme order statistics. *The Annals of Statistics 3*(1), 119–131.

Reed, W. J. and M. Jorgensen (2004). The double Pareto-lognormal distribution - A new parametric model for size distributions. *Communications in Statistics - Theory and Methods 33*(8), 181–194.

Resnick, S. and C. Stărică (1997). Smoothing the Hill estimator. *Advances in Applied Probability 29*(1), 271–293.

Resnick, S. and C. Stărică (1998). Tail index estimation for dependent data. *Annals of Applied Probability 8*(4), 1156–1183.