ARTICLE IN PRESS
JID: ECOSTA [m3Gsc;March 27, 2017;15:28]

Econometrics and Statistics 000 (2017) 1–20

Contents lists available at ScienceDirect

Econometrics and Statistics

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

Improved estimators of extreme Wang distortion risk measures for very heavy-tailed distributions

Jonathan El Methni[a], Gilles Stupfler[b,*]

[a] Université Paris Descartes, Sorbonne Paris Cité, Laboratoire MAP5, UMR CNRS 8145, 75006 Paris, France
[b] School of Mathematical Sciences, The University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom

ARTICLE INFO

Article history:
Received 30 November 2016
Revised 6 March 2017
Accepted 6 March 2017
Available online xxx

Keywords:
Asymptotic normality
Extreme value statistics
Heavy-tailed distribution
Trimming
Wang distortion risk measure
Winsorising

ABSTRACT

A general way to study the extremes of a random variable is to consider the family of its Wang distortion risk measures. This class of risk measures encompasses several indicators such as the classical quantile/Value-at-Risk, the Tail-Value-at-Risk and Conditional Tail Moments. Trimmed and winsorised versions of the empirical counterparts of extreme analogues of Wang distortion risk measures are considered. Their asymptotic properties are analysed, and it is shown that it is possible to construct corrected versions of trimmed or winsorised estimators of extreme Wang distortion risk measures who appear to perform overall better than their standard empirical counterparts in difficult finite-sample situations when the underlying distribution has a very heavy right tail. This technique is showcased on a set of real fire insurance data.

1. Introduction

Early developments of extreme value analysis focused on estimating a quantile at a level so high that the straightforward empirical quantile estimator could not be expected to be consistent. Motivating problems include estimating extreme rainfall at a given location (Koutsoyiannis, 2004) or extreme daily wind speeds (Beirlant et al., 1996), modeling large forest fires (Alvarado et al., 1998), analysing extreme log-returns of financial time series (Drees, 2003) and studying extreme risks related to large losses for an insurance company (Rootzén and Tajvidi, 1997). A large part of practical applications of extreme value theory can actually be modelled using heavy-tailed distributions, which shall be the focus of this paper. A distribution is said to be heavy-tailed if its survival function $1-F$, where F is the related cumulative distribution function, roughly behaves like a power function with exponent $-1/\gamma$ at infinity where the positive parameter γ is the so-called tail index of the distribution. In such a model, the function $1-F$ essentially satisfies a homogeneity property and it therefore becomes possible to use an extrapolation method (Weissman, 1978) to estimate quantiles at arbitrarily extreme levels, provided an estimate of γ is computed. Under appropriate stationarity assumptions this analysis can be used to draw predictive conclusions: extreme value analysis has been applied to determine how high the dykes surrounding the areas below sea level in the Netherlands should be so as to protect these zones from flood risk in case of extreme storms affecting Northern Europe (de Haan and Ferreira, 2006). It is also used nowadays by insurance companies operating in Europe so as to determine their own solvency capital necessary to meet the European Union Solvency II directive requirement that an insurance company should be able to survive the upcoming calendar year with a probability not less than 0.995.

* Corresponding author.
E-mail addresses: jonathan.el-methni@parisdescartes.fr (J. El Methni), gilles.stupfler@nottingham.ac.uk, gilles.stupfler@gmail.com (G. Stupfler).

http://dx.doi.org/10.1016/j.ecosta.2017.03.002
2452-3062/

Please cite this article as: J. El Methni, G. Stupfler, Improved estimators of extreme Wang distortion risk measures for very heavy-tailed distributions, Econometrics and Statistics (2017), http://dx.doi.org/10.1016/j.ecosta.2017.03.002

Of course, the knowledge of a single high quantile is clearly not enough to characterise the behaviour of a random variable in its right tail, since two distributions may well share a quantile at some common level although their respective tail behaviours are different. This is why other quantities such as the Tail-Value-at-Risk, Conditional Value-at-Risk or Conditional Tail Moment (see El Methni et al., 2014) were developed and studied; a common feature of these indicators is that their computation takes into account the whole right tail of the random variable of interest. This, of course, also entails increased sensitivity to a change in tail behaviour compared to what is observed in quantiles, at the population level and at the finite-sample level alike. These measures are of great value in practice, especially in actuarial science: for instance, as mentioned in Dowd and Blake (2006), the Tail-Value-at-Risk would be used if one is interested in the average loss after a catastrophic event or to estimate the cover needed for an excess-of-loss reinsurance treaty. As shown in El Methni and Stupfler (2017), the aforementioned quantities can actually be written as simple combinations of Wang distortion risk measures of a power of the variable of interest (abbreviated by Wang DRMs hereafter; see Wang, 1996). Wang DRMs are weighted averages of the quantile function, the weighting scheme being specified by the so-called distortion function; on the practical side, Wang DRMs can, among others, be useful to price insurance premiums, bonds, and tackle capital allocation problems, see e.g. Wang et al. (1997), Wang (2004) and Belles-Sampera et al. (2014). It is therefore not surprising that the estimation of Wang DRMs above a *fixed* level of risk has been the subject of a number of papers: in particular, we refer the reader to Jones and Zitikis (2003), Necir and Meraghni (2009), Necir et al. (2010) and Deme et al. (2013, 2015).

To the best of our knowledge though, the only study providing estimators of *extreme* distortion risk measures is the recent work of El Methni and Stupfler (2017). More precisely, they show that a simple and efficient solution to estimate extreme Wang DRMs when the right tail of the underlying distribution is moderately heavy is to consider a so-called functional plug-in estimator. Two weaknesses of this study can be highlighted however. The first problem, a practical one, is that it is a consequence of the results in the simulation study of El Methni and Stupfler (2017) that the finite-sample performance of the suggested class of estimators decreases sharply in terms of mean squared error as the tail of the underlying distribution gets heavier. This is due to the propensity of heavier-tailed distributions to generate highly variable top order statistics and, therefore, to increase dramatically the variability of the estimates. No solution is put forward in El Methni and Stupfler (2017) in order to tackle this issue. The second problem, which is theoretical, is that their asymptotic results about this class of estimators are restricted to asymptotic normality and are thus somewhat frustrating in the sense that they are stated under an integrability condition on the quantile function which is substantially stronger than the simple existence of the Wang DRM to be estimated. In particular, a consistency result under the latter condition, in the spirit of the one Jones and Zitikis (2003) obtained for the estimation of fixed-order Wang DRMs, is not provided in El Methni and Stupfler (2017).

Herein it is shown that robustifying the functional plug-in estimator of El Methni and Stupfler (2017) by deleting certain top order statistics and/or replacing them by lower order statistics, namely trimming or winsorising the estimator, enables one to obtain estimators with reduced variability, as well as to show a consistency result under weaker hypotheses and to retain the asymptotic normality result under the same technical conditions. Trimming and winsorising have both been (and arguably still are) the easiest and most intuitive ways to give a statistical technique some degree of robustness to high-value outliers. A historical account is given in Stigler (1973). The motivation here is rather that the integrability condition of El Methni and Stupfler (2017) depends solely on the behaviour of the quantile function around 1 and becomes more and more stringent as the rate of divergence of this function to infinity increases. At the sample level, this means that this integrability condition has to be fulfilled in order to control the highest order statistics. Deleting the most extreme part of the sample or replacing it by lower (but still high) order statistics can thus be thought of informally as a way to reduce the difficulty of the problem, both from the theoretical and practical point of view.

To be more specific, we shall essentially consider a Wang DRM of a random variable given that it lies between two high-level quantiles, instead of assuming that it simply lies above a high threshold like El Methni and Stupfler (2017) did. This is then estimated by its empirical counterpart, which leads to a trimmed estimator of a Wang DRM. The winsorised estimator, meanwhile, is obtained by considering the empirical counterpart of a Wang DRM given that the random variable lies above a high threshold and is clipped above yet another higher level. By construction, these two estimators do not depend on some of the highest observations, and therefore can be expected to suffer from less finite-sample variability than the original estimator of El Methni and Stupfler (2017) does. To ensure consistency, the highest level (that is, the trimming/winsorising level) is then made to increase to 1 faster than the lowest one does as the sample size increases. Both of these estimators can actually be embedded into a common class of estimators, whose consistency and asymptotic normality are studied. A somewhat surprising feature of this technique is that one can also obtain the consistency of the estimator using the full data above a high level by approximating it by such robustified estimators whose fraction of deleted data becomes smaller as the sample size increases; this argument is actually similar in spirit to a proof by Etemadi (1981) of the law of large numbers for independent copies of an integrable random variable, starting with the case when the variance is finite and concluding by a truncation argument.

These new estimators, for all their improved properties as far as variability is concerned, should be expected to suffer from finite-sample bias issues, since they are in fact sample counterparts of a different quantity than the originally targeted Wang DRM. The second step is then to devise a correction method which allows the estimator to have a bias intuitively similar to that of the basic functional plug-in estimator and therefore to be (almost) unbiased in practice, while retaining its low variability. The gist of the correction step is to note that the newly proposed estimators are in reality approximately equal to the Wang DRM to be estimated multiplied by a quantity converging to 1 and depending on the extremes of the sample only. This makes it possible to estimate the error made when using the purely trimmed or winsorised estimators