Challenges in implementing worst-case analysis

Jon Danielsson

Systemic Risk Centre, lse ,Houghton Street, London WC2A 2AE, UK

Lerby M. Ergun*

Systemic Risk Centre, lse ,Houghton Street, London WC2A 2AE, UK

Casper G. de Vries

School of Economics, Erasmus University Rotterdam, 3000 DR Rotterdam, The Netherlands

March 1, 2017

Abstract

Worst-case analysis has increased in popularity among financial regulators in the wake of the recent financial crisis. In this paper we provide insight into this measure and provide some guidance on how to estimate it. We derive the bias for the non-parametric heavy tailed order-statistics and contrast it with the semi-parametric EVT approach. We find that if the return distribution has a heavy tail, the non-parametric worst-case analysis, i.e the minimum of the sample, is always downwards biased. Relying on semi-parametric EVT reduces the bias considerably in the case of relatively heavy tails. But for the less heavy tails this relationship is reversed. Estimates for a large sample of US stock returns indicates that this pattern in the bias is also present in financial data. With respect to risk management, this induces an overly conservative capital allocation

Keywords: Worst-case analysis, EVT, quantile estimator, risk management
JEL codes: C01, C14, C58

*Corresponding author: `m.l.ergun@lse.ac.uk`, Houghton Street, London WC2A 2AE, UK

1 Introduction

Worst-case analysis studies the worst expected outcome over a predetermined time length, with a typical question what is the worst daily market outcome in 10 years or 2,500 days. Worst-case analysis is increasingly common due to the recent financial crises. For instance, European Insurance and Occupational Pensions Authority (EIOPA) has embraced the worst-case analysis as an additional measure of risk for the insurance and pension industry. In spite of its increasing importance, little is known about worst-case analysis and what the appropriate way is to measure it. In this paper we provide insight in the bias of worst-case estimates and measure their effect for US stock returns.

There are generally three main approaches. The simplest, and the most obvious, is to directly read the object of interest from the empirical distribution, in our case the historical minima, the non-parametric approach (EA). One can also assume a model only for the tail of the distribution and not model the center of the distribution, the semi-parametric approach (SP). The third approach is based on specifying a parametric distribution for all outcomes and estimate its parameters.

Of these three alternatives, the last is the only one that cannot be recommended. The reason is that the estimates will be dominated by the center of the distribution, so that the fit is optimal for a typical observation, not the lowest, the focus of our interest. Therefore, such an approach would in most cases deliver less precise and more uncertain worst-case estimators than either the EA or the SP. We therefore focus on EA and the SP estimators and provide guidance on the appropriate use of either.

1.1 Statistical background

To derive the bias of the worst observed observation as a worst-case estimator, we start with a relatively general approach. Begin by deriving the distribution of observations in an ordered sample. Suppose we observe some i.i.d heavy tailed random variable $Y_1, ..., Y_n$ with distribution F, where

$$\lim_{t\to\infty} \frac{1-F(tx)}{1-F(t)} = x^{\alpha}, \quad \alpha > 0. \tag{1}$$

The sorted sample, i.e order-statistics, can be represented as,

$$\max(Y_1, ..., Y_n) = X^{(1,n)} \geq X^{(2,n)} \geq \cdots \geq X^{(n,n)} = \min(Y_1, ..., Y_n)$$

The distribution of the order-statistics can be studied through the number of exceedances. These follow a binomial distribution,

$$G^{(k,n)}(x) = \sum_{r=0}^{k-1} \binom{n}{r} [1 - \mathrm{F}(x)]^r [\mathrm{F}(x)]^{n-r}. \tag{2}$$

Suppose one is interested in the distribution of the maximum realization:

$$\Pr(\max(Y_1, ..., Y_n) < x) = G^{(1,n)}(x) = [\mathrm{F}(x)]^n. \tag{3}$$

Similarly to the standard central limit theorem for the asymptotic distribution of the arithmetic mean, Fisher and Tippett (1928) and Gnedenko (1943) provide a limit theorem for the asymptotic distribution of the maximum, i.e extreme value theory (EVT).

EVT gives the conditions under which there exist sequences b_n and a_n such that

$$\lim_{n\to\infty} [\mathrm{F}(a_n x + b_n)]^n \to \mathrm{G}^{(1,n)}(x),$$

where $\mathrm{G}^{(1,n)}(x)$ is the Fréchet distribution.

Theorem 2.2.2 in Leadbetter et al. (1983) extends the EVT for the maximum to lower order-statistics by means of the Poisson property of the lower order-statistics. In particular, the asymptotic distribution of the k^{th} largest order-statistic is:

$$\mathrm{G}^{(k,n)}(x) \to \mathrm{G}^{(1,n)}(x) \sum_{s=0}^{k-1} \frac{\left(-\log\left[\mathrm{G}^{(1,n)}(x)\right]\right)^s}{s!}. \tag{4}$$

From (4) we determine the expectation of the order-statistics

$$\mathrm{E}\left[X^{(k,n)}\right] = \frac{a_n}{[k-1]!} \Gamma\left[k - \frac{1}{\alpha}\right]. \tag{5}$$

To determine this expectation for a specific heavy tailed distribution a_n needs to be chosen appropriately.[1] To find a good approximation, we use the first order term of the Hall expansion (Hall and Welsh, 1985):

$$\Pr(Y \leq -y) = \mathrm{F}(-y) = Ay^{-\alpha}[1 + By^{-\beta} + o(y^{-\beta})]. \tag{6}$$

[1] For the heavy tailed distributions $b_n = 0$.

For the Pareto distribution, $\mathrm{F}(-y) = Ay^{-\alpha}$, we observe that the Hall expansion perfectly fits the first order term.[2] For the Pareto distribution the scaling constant a_n is $(An)^{\frac{1}{\alpha}}$, where A is the scale parameter. Therefore, we can, through the Hall expansion, extract a good approximation of the expectation of the order-statistics of heavy tailed distributions.

1.2 The semi–parametric approach

To contrast the expectation of the maximum observation, we compare it to a semi-parametric estimator of the worst-case. By inverting the first order expansion in (6), using the empirical counter part of $A = \frac{t/n}{X^{(t,n)^{-\alpha_t}}}$ measured at some threshold t and $F(-y) = k/n$, one obtains the semi-parametric tail quantile estimator

$$\hat{x}_{SP}^{(n-k/n)}(t) = X^{(t,n)} \left(\frac{t}{k}\right)^{\frac{1}{\alpha_t}}.$$

Goldie and Smith (1987) derive the distribution of the semi-parametric quantile estimator

$$\frac{\sqrt{t}}{\log(t/np)} \left(\frac{\hat{x}_{SP}^{(n-k/n)}(t)}{x^{(p)}} - 1\right) \sim N\left(-\frac{\operatorname{sign}(B)}{\sqrt{2\beta\alpha}}, \frac{1}{\alpha^2}\right), \tag{7}$$

where B and β are the second order scale and shape parameters from (6).

1.3 Comparing the empirical- and semi-parametric quantile estimator

The two approaches, the EA and the SP, have each their own pros and cons. While the EA is much simpler to implement, the SP might be more accurate because it uses more observations in the estimation, and therefore might result in an estimate with a lower bias and uncertainty. However, the SP is dependent on correctly specifying the semi-parametric distribution and identifying the threshold $X^{(t,n)}$.

To shed more light on the use of these two estimators, we compare their bias

[2] All of the standard heavy tailed distributions satisfy the Hall expansion. It even applies to the GARCH(1,1) unconditional distribution.

at $p = 1/n$. From (5) and (7) the bias of the two approaches are:

$$\left(\frac{\hat{x}_{SP}^{(n-k/n)}(t)}{x^{(p)}} - 1\right) \sim -\frac{\text{sign}(B)}{\sqrt{2\beta\alpha}}\frac{\log(k)}{\sqrt{k}}. \qquad \text{SP} \tag{8}$$

$$\left(\frac{\hat{x}_{EA}^{(n-k/n)}}{x^{(p)}} - 1\right) \sim \frac{1}{[k-1]!}\Gamma\left[k - \frac{1}{\alpha}\right] - 1. \qquad \text{EA} \tag{9}$$

(8) and (9) indicate that neither approach is dominating in all circumstances. For the EA the asymptotic bias approaches infinity as α approaches 1. However, as α increases the Γ function decreases rapidly.[3] For the SP estimator the bias is relatively small for moderate values of β and t. This leads to a crossing point in the bias of the two estimators with respect to α.

Given values of t and β, we define switching point α^*. For $\alpha < \alpha^*$ the absolute bias of the SP estimator is smaller than the EA estimator. When $\alpha > \alpha^*$, the relationship is reversed. For a fixed t, this relationship is depicted in Figure 1. This figure portraits at which combination of α and β the bias of the non-parametric worst-case estimator becomes smaller than the semi-parametric approach.

For example, in the case of the family of Student-t distributions $\alpha \in [1, \infty)$ and $\beta = 2$. From Figure 1, we read that for $t = e^2$ the switching of the biases occurs around $\alpha^* \approx 5$. For higher and lower values of t, the α^* increases. For the family of symmetric stable distributions, the bias is always smaller for the SP, as $\beta = \alpha$ and $\alpha < 2$.

2 Application

In this section we investigate the bias in the EA and SP estimators for US stocks, using the CRSP security database. For this empirical application we use the Hill estimator to estimate the tail exponent α. This estimator depends on a selection of a high order-statistic as a threshold, i.e $X^{(t,n)}$. This nuisance statistic is obtained by the KS-distance metric.[4] Given the estimate of the tail index, the quantile can be estimated semi-parametrically. We compare the difference between the previously introduced worst-case estimators

[3] The bias of EA is invariant to dependence in the time-series as measured by the extremal index. See Leadbetter et al. (1983) page 416.

[4] The KS-distance metric chooses the threshold which minimizes the maximum quantile distance between the empirical and Pareto distribution.This approach is further explained in Danielsson et al. (2015).

Figure 1: Bias comparison

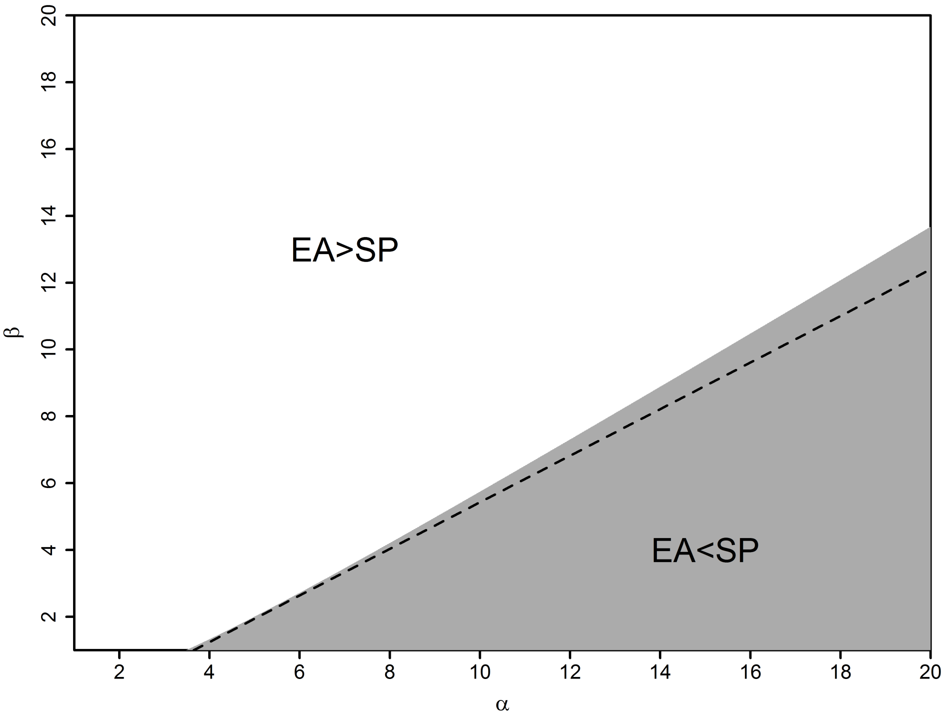

This Figure depicts the area where the absolute bias of the semi-parametric estimator becomes larger than the bias of the order statistics (colored region). The biases of the estimators are at $p = 1/n$ as in (8) and (9). For this figure we fix t at e^2. To the right of the lines the combination of α and β produces a larger bias for the semi-parametric approach. The dotted line shows where the boundary shifts to when the threshold t is doubled to e^2.

for each stock at the $1/n$ quantile. These differences are collected in different buckets sorted by the $\hat{\alpha}$ of each stock. This way we are able to determine a switching point in the size of the biases between the two estimators.

The theoretical results stipulate that the relative size of the bias changes as a function of $\hat{\alpha}$. Table 1 portraits this pattern for the securities in the CRSP database. For these stocks the switch point is around $\hat{\alpha} = 3$. It is difficult to determine the exact switch point for real data. This is because β, in the bias of the semi-parametric quantile estimator, is not estimated. In addition, the Hill estimator is estimated with a bias. This makes it difficult to determine the exact switch point. It is encouraging that we see a monotonic decline in the average difference as $\hat{\alpha}$ increases. This is supportive of the result that

Table 1: CRSP data

	All $r = 2,230$	$\hat{\alpha} < 2$ $r = 13$	$2 < \hat{\alpha} < 3$ $r = 712$	$3 < \hat{\alpha} < 4$ $r = 1,001$	$4 < \hat{\alpha} < 5$ $r = 442$	$\hat{\alpha} > 5$ $r = 62$
Mean	-0.003	0.074	0.010	-0.005	-0.017	-0.026
Median	-0.011	0.074	0.005	-0.010	-0.017	-0.026
Ste. dev.	0.031	0.072	0.039	0.022	0.019	0.018
1% quantile	-0.059	-0.038	-0.063	-0.051	-0.063	-0.085
99% quantile	0.109	0.194	0.138	0.064	0.045	0.031
Rank sum test	0.070	0.305	0.397	0.162	0.003	0.047

This table reports the difference between the largest order statistic and semi-parametric quantile estimator for US stocks. These are stocks selected from the CRSP database. The securities need to be traded on NYSE, AMEX, NASDAQ, and NYSE Arca exchanges over the period from 01-01-1995 till 01-01-2011. The table reports various statistics on the distribution of the difference between the two estimators. Here r is the number of different stocks in the different buckets. To determine the number of order statistics for the Hill estimator we use the KS distance metric described in Danielsson et al. (2016).

the bias of the EVT based quantile estimator overtakes the bias of the non-parametric quantile estimator. The results for the median convey the same story.

The 1% and 99% quantiles of the buckets show that although the mean and median showcase a switch between the severity of the bias of the quantile estimators, this might be statistically insignificant. Therefore, we employ the rank-sum test to test for the difference in size of the observations of the empirical distribution from the SP and EA estimator. We find that for the lighter heavy tailed stock returns the estimates from the estimators are significantly different from one another. The empirical distribution of the semi-parametric quantile estimator tends to have larger values than the distribution of the non-parametric quantile estimator. For low values of $\hat{\alpha}$ the difference is in the expected direction, but insignificant.

3 Conclusion

With worst-case analysis becoming increasingly common in both policymaking and practice, it is of interest to evaluate the qualities of common methods for such applications. The simplest, and perhaps the most common way is to estimate the worst-case by taking the most negative outcome in the historical sample. Alternatively, one could estimate the lower tail of the distribution by semi-parametric methods and use that to calculate the worst-case.

In our main conclusion, either method is best, depending on how heavy the tails are and their specific shape. Generally, for the heaviest, the semi-parametric approach is best, and as it thins, the historical minima eventually becomes better.

4 Acknowledgements

This work was supported by the Economic and Social Research Council (ESRC) in funding the Systemic Risk Centre [grant number ES/K002309/1] and the Netherlands Organisation for Scientific Research Mozaiek grant [grant number: 017.005.108]

References

Danielsson, J., L. M. Ergun, L. De Haan, and C. De Vries (2016). Tail index estimation: Quantile driven threshold selection. *SRC Working paper*.

Fisher, R. A. and L. Tippett (1928). Limiting forms of the frequency distribution of the largest of smallest member of a sample. In *Proceedings of the Cambridge Philosophical Society*, Volume 24, pp. 180–190.

Gnedenko, B. (1943). Sur la distribution limite du terme maximum d'une serie aleatoire. *Annals of Mathematics 44*(1), 423–453.

Goldie, C. M. and R. L. Smith (1987). Slow variation with remainder: Theory and applications. *The Quarterly Journal of Mathematics 38*(1), 45–71.

Hall, P. and A. Welsh (1985). Adaptive estimates of parameters of regular variation. *The Annals of Statistics 13*(1), 331–341.

Leadbetter, M. R., G. Lindgren, and H. Rootzén (1983). Extremes and related properties of random sequences and processes. *Springer Series in Statistics*.