

WORKING PAPER 2001

A comment on fitting Pareto tails to complex survey data

Rafael Wildauer and Jakob Kapeller

January 2020



A Comment on Fitting Pareto Tails to Complex Survey Data

Rafael Wildauer¹ and Jakob Kapeller²

¹University of Greenwich

²University Duisburg-Essen and Johannes Kepler University Linz

January 2020

Abstract

Taking survey data on household wealth as our major example, this short paper discusses some of the issues applied researchers are facing when fitting (type I) Pareto distributions to complex survey data. The major contribution of this paper is twofold: First, we propose a new and intuitive way of deriving Gabaix and Ibragimov's (2011) bias correction for Pareto tail estimations from which the generalization to complex survey data follows naturally. Second, we summarise how Kolmogorov-Smirnov and Cramer-von-Mises goodness of fit tests can be generalized to complex survey data. Taken together we think the paper provides a concise and useful presentation of the fundamentals of Pareto tail fitting with complex survey data.

Keywords: Pareto distribution, complex survey data, wealth distribution

JEL Classification: D31, C46, C83

The usual disclaimer applies. We would like to thank Stephan Steinerberger for useful comments and discussions.

1 Introduction

Taking survey data on household wealth as our major example, this short paper discusses some of the issues applied researchers are facing when fitting (type I) Pareto distributions to complex survey data. First we show that Kratz and Resnick’s (1996) standard QQ regression approach is often implemented based on a formulation of the complementary cumulative distribution function (CCDF) that is tailored to avoid the occurrence of $\text{Log}(0)$. While giving plausible results this approach lacks generality and, hence, is not appealing on statistical grounds. In what follows, we provide a different approach towards the problem of adequately deriving the CCDF from complex survey data, which is more general than past approaches and, at the same time, incorporates existing innovations in the literature.

While different formulations of the CCDF typically arise from different orderings of the data vector, we introduce an alternative formulation of the CCDF by averaging the CCDF obtained from data vectors in ascending and descending order, which allows not only to avoid the $\text{Log}(0)$ problem, but also corresponds to the bias correction of the QQ regression proposed by Gabaix and Ibragimov (2011). Thus we are able to provide an alternative, more general and probably also more intuitive explanation of why and how Gabaix and Ibragimov’s (2011) bias correction works. At the same time our alternative formulation is easier to implement in practical work. Finally, with our proposed implementation of Gabaix and Ibragimov (2011) the generalization towards complex survey weights arises naturally. In general the formulation of the averaged CCDF allows for an easy implementation of Vermeulen’s (2018) rich list or Wildauer and Kapeller’s (2019) rank correction approach.

In addition, we also summarise the extension of goodness of fit tests in the form of the Kolmogorov-Smirnov and the Cramer-von-Mises test statistics to complex survey data. These extensions are important for researchers who are interested in determining the scale parameter of the Pareto distribution not based on ad hoc assumptions or a purely graphical analysis but on goodness of fit tests as demonstrated by Clauset, Shalizi, and Newman (2009).

Hence, the aim of this paper is twofold: First, we provide an alternative perspective on key aspects of Pareto tail fitting such as Gabaix and Ibragimov’s (2011) bias correction and a new and easy way of implementing the latter. Second, we summarise key results on goodness of fit tests with complex survey data. Taken together we think the paper provides a concise and useful presentation of the fundamentals of Pareto tail fitting with complex survey data.

2 Revisiting Fundamentals

Much of the literature about fitting Pareto tails to wealth survey data (Jayadev, 2008; Bach, Thiemann, & Zucco, 2018; Dalitz, 2018; Vermeulen, 2018) relies on simple OLS regressions, which fit the complementary cumulative distribution function (CCDF) of the Pareto distribution to the empirical CCDF derived from an available sample (Kratz & Resnick, 1996). The theoretical CCDF for a random variable X following a Pareto distribution is defined as follows:

$$CCDF_T(x_i) = Pr(X > x_i) = \left(\frac{x_m}{x_i}\right)^\alpha \quad (1)$$

where x_m is the scale and α the shape parameter. In addition assume we have a random sample of households with net wealth $x = (x_1, \dots, x_n)$ and corresponding weights $w = (w_1, \dots, w_n)$, where the number of households represented by the available sample is defined as $N = \sum_{i=1}^n w_i$ and the data vector is organised in descending order (i.e., from the most to the least affluent observation). This setup yields a data vector denoted as $x_d = (x_{(1)}, \dots, x_{(n)})$ with the corresponding vector of weights $w_d = (w_{(1)}, \dots, w_{(n)})$. Then the empirical CCDF is written as:

$$CCDF(x_{(i)})_d = \frac{\sum_{1 \leq j \leq i} w_{(j)}}{N} \quad (2)$$

Combining the theoretical and empirical CCDFs provides the basis for a regression equation:

$$\frac{\sum_{1 \leq j \leq i} w(j)}{N} = \left(\frac{x_m}{x(i)} \right)^\alpha \quad (3)$$

which can be translated into the following regression equation:

$$\ln \left(\sum_{1 \leq j \leq i} w(j) \right) = C_1 - \alpha \ln(x(i)) + \epsilon_i \quad (4)$$

where $C_1 = \ln(N) + \alpha \ln(x_m)$. Equation (4) is estimated by OLS in order to obtain an estimate of the shape parameter (α). Vermeulen (2018) extended this standard approach along two lines. First, he proposed to add observations from rich lists to the data vector. Secondly, he incorporated Gabaix and Ibragimov's (2011) (G&I from here onwards) bias correction and generalized it to situations of complex survey weights. For now we will focus on this extension. G&I argue that it has long been known that OLS estimation of equation (4) yields a biased estimate of the shape parameter α (e.g. Aigner & Goldberger, 1970). They show that subtracting the value 1/2 from $\sum_{j=1}^i w(j)$ will eliminate this bias. However, G&I also assume that $w_i = w_j = 1$ so that the expression $\sum_{j=1}^i w(j)$ is equated with the rank of observation i as represented by the index number $(i) = (1), \dots, (n)$. Vermeulen (2018) extends G&I's bias correction approach to the case of complex survey weights which means that weights are not equal to 1 and differ across observations. Starting from a data vector arranged in descending order (x_d) with a corresponding weights vector (w_d) he defines $\bar{N} = \frac{\sum_{j=1}^i w(j)}{n} = \frac{N}{n}$ as the average weight and $\bar{N}_i = \frac{\sum_{j=1}^i w(j)}{i}$ as the average weight up to observation i . Then the empirical CCDF is defined as:

$$CCDF(x(i))_d = \frac{\sum_{1 \leq j \leq i} w(j)}{N} = i \frac{\bar{N}_i \bar{N}}{N} \quad (5)$$

Which allows for a straightforward inclusion of the Gabaix & Ibragimov bias correction:

$$CCDF(x(i))_V = (i - 0.5) \frac{\bar{N}_i \bar{N}}{N} \quad (6)$$

And based on that Vermeulen (2018) derives the equation to be estimated by OLS as:

$$\ln \left((i - 0.5) \frac{\bar{N}_i \bar{N}}{N} \right) = C_2 - \alpha \ln(x(i)) + \epsilon_i \quad (7)$$

where $C_2 = \ln(N) + \alpha \ln(x_m) - \ln(\bar{N}) = C_1 - \ln(\bar{N})$. This is the standard setting in much of the applied literature to fit a Pareto tail to wealth survey data. Augmenting the regression (7) by observations from rich lists as suggested by Vermeulen (2018) is easy to do.

Another relevant tool is Clauset et al.'s (2009) method for determining the scale parameter x_m . In practice this is often done based on ad hoc assumptions or based on a purely graphical inspection of the data. Clauset et al. (2009) on the other hand provide a procedure which determines x_m based on statistical goodness of fit tests. The idea is to choose a generous lower bound in the available data, for example the 90th wealth percentile, and then fit Pareto distributions to increasingly smaller subsamples starting beyond the chosen lower bound. For example one could fit Pareto distributions for increasingly smaller subsets by removing the smallest observation above the chosen lower bound in each step. Then one computes goodness of fit statistics based on Kolmogorov-Smirnov or Cramer-von-Mises test statistics to compare the different fitted distributions. The distribution which exhibits the smallest goodness of fit statistic (and thus the best fit to the data) is chosen as the preferred specification. In this way x_m is defined as the smallest observation in the subset which exhibits the best fit to the data. It is crucial to note here that Clauset et al.'s (2009) original paper does not discuss the situation of complex survey weights, but rather proceeds by suggesting a method for more concisely testing the adequacy of assuming a Pareto tail in the first place.

Against this backdrop the paper presents comments on three aspects of these standard procedures. First, we clarify how the ordering of the data impacts on the exact shape of the empirical CCDF. Second, we provide a simpler and more intuitive formulation of G&I’s bias correction from which the generalization to complex survey data emerges naturally. This yields a reformulation of the basic QQ regression incorporating the G&I bias correction and being compatible with Vermeulen’s (2018) rich list approach as well as Wildauer and Kapeller’s (2019) rank correction approach. Third we discuss the generalization of Kolmogorov-Smirnov or Cramer-von-Mises test statistics to complex survey data. We think all three aspects are of high relevance for practitioners in the field.

3 Rethinking Fundamentals

3.1 Data ordering and alternative CCDF definitions

Most exercises in fitting Pareto distributions to wealth survey data start by formatting the data vector in descending order and then define the empirical CCDF based on equation (2). However, this definition does not exactly conform to the theoretical idea expressed in equation (1), which demands a strict inequality and, hence, denotes the probability of observing someone with wealth greater than x_i . In contrast, the formulation used in (2), does not demand a strict inequality and, hence, exhibits a systematic deviation from its theoretical counterpart. In the case of the richest household ($CCDF(x_{(1)})_d = \frac{w_{(1)}}{N}$), this implies the assignment of a non-zero probability to observing a household more affluent than $x_{(1)}$, although the theoretical formulation would imply a probability of zero. In the case of the poorest household ($CCDF(x_{(n)})_d = \frac{N}{N} = 1$), the same inconsistency arises as a probability of 1, instead of a probability slightly less than 1, is assigned to this observation.

The simple reason why the Pareto literature deviates from the theoretical definition of the CCDF and defines the latter based on equation (2) is that it ensures that the empirical CCDF is never equal to 0 and thus can be log-linearised without having to deal with $\log(0) = -\infty$. This is the reason why other applications which do not require log-linearization organise the data vector in ascending order $x_a = (x_{(n)}, \dots, x_{(1)})$ with a corresponding vector of weights $w_a = (w_{(n)}, \dots, w_{(1)}) = (1, \dots, 1)$ and then define the empirical CCDF as:

$$CCDF(x_{(i)})_a = 1 - \frac{\sum_{i \leq j \leq n} w_{(j)}}{N} \quad (8)$$

This approach to the empirical CCDF is perfectly in line with its theoretical counterpart: For the case of the richest household ($CCDF(x_{(1)})_a = 1 - \frac{N}{N} = 0$), a zero probability is assigned to observing a more affluent household and for the poorest household ($CCDF(x_{(n)})_a = 1 - \frac{w_{(n)}}{N} = 1 - \frac{1}{N}$) a below unit probability is assigned to observe a more affluent household. These examples demonstrate why a definition based on equation (8) more accurately represent the available sample information.

That said, we can also derive formulas to compute $CCDF(x_{(i)})_a$ from data organised in descending order and vice versa. For example given the data vector $x_d = (x_{(1)}, \dots, x_{(n)})$ with corresponding weights $w_d = (w_{(1)}, \dots, w_{(n)})$ we can define $w_{(0)} = 0$. Thus, we amend the weights vector with a zero entry and then compute $CCDF(x_{(i)})_a$ as:


$$CCDF(x_{(i)})_a = \frac{\sum_{1 \leq j \leq i} w_{(j-1)}}{N} \quad (9)$$

Note that the main difference between equation (8) and equation (9) are the two different subsets of observations over which the sum operator is defined. $i \leq j \leq n$ in the first case and $1 \leq j \leq i$ in the latter. If both equations are applied on the same weights vector, they will give identical results.

So the issue of defining the empirical CCDF is not about how the data is organised but from a conceptual point of view whether one wants to define the empirical CCDF as equal to the probability $Pr(X > x_{(i)})$ as in the case of $CCDF(x_{(i)})_a$ or equal to the probability $Pr(X \geq x_{(i)})$

as in the case of $CCDF(x_{(i)})_d$. The reason why the literature on fitting Pareto tails chooses the definition $CCDF(x_{(i)})_d = Pr(X \geq x_{(i)})$ by default, simply lies in the need to log-linearise the CCDF and to avoid the logarithm of 0. Before moving on to the issue of G&I's bias correction in the next section, Figure 1 plots $CCDF_a$ and $CCDF_d$ for a random sample $N=10$ from a univariate distribution over the interval $[0,15]$. One can clearly see that $CCDF(14)_a = 0$ whereas $CCDF(14)_d > 0$ and $CCDF(0)_a < 1$ whereas $CCDF(0)_d = 1$. However we can also see that $CCDF(x_i)_a = CCDF(x_{i+1})_d$. That means these two CCDFs are shifted versions of each other. A feature which will become clearer in the next section.

Figure 1: Different empirical CCDFs



Empirical CCDFs for a sample ($N=10$) from a univariate distribution over the interval $[0,15]$.

3.2 Bias Correction by Averaging

G&I argue that in order to reduce the bias in OLS based estimates of the shape parameter of the Pareto distribution, researchers should subtract the value $1/2$ from the rank of each observation prior to performing the rank-wealth regression. The rank of observation i in their terminology is equivalent to the cumulative weight of that observation i given by $\sum_{j=1}^i w_{(j)}$ based on a descending data vector. So the bias correction proposed by Gabaix and Ibragimov (2011) corresponds to computing the empirical CCDF in the following way:

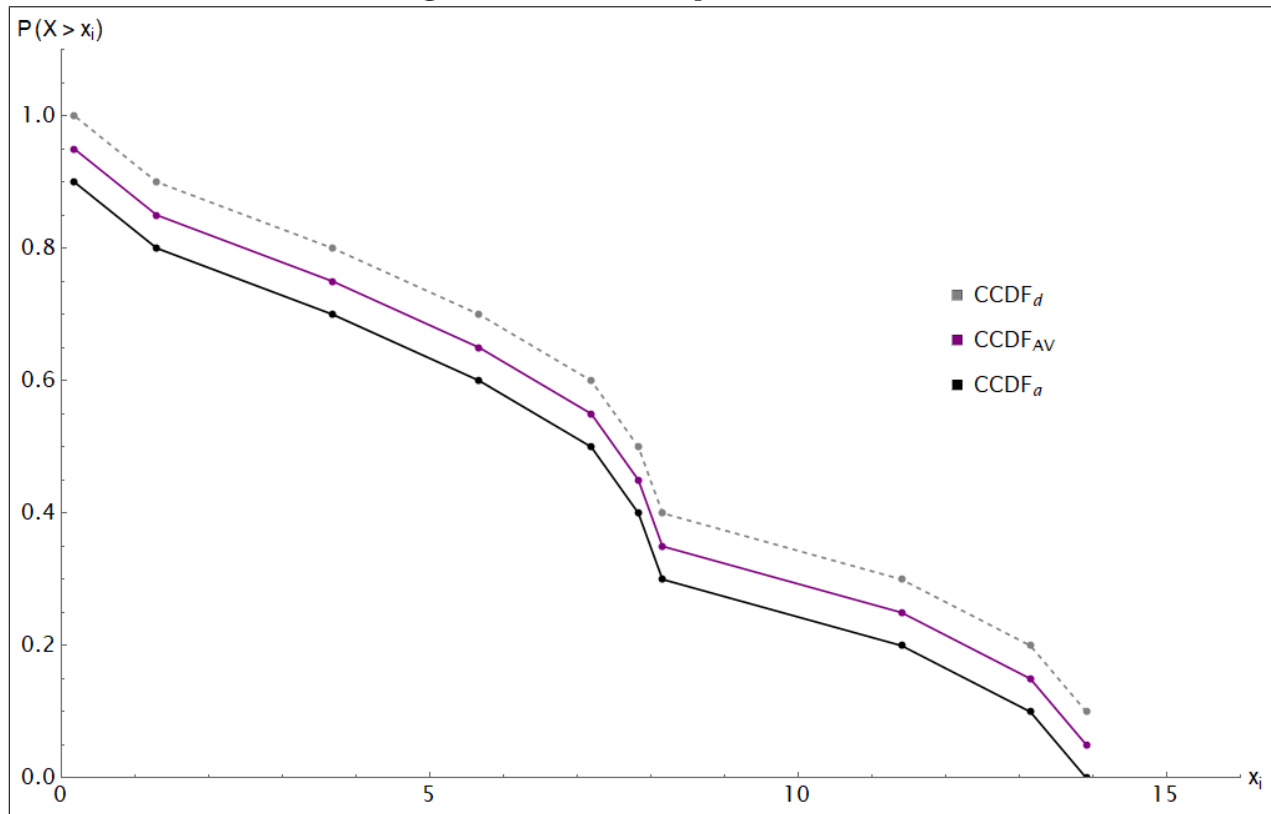
$$CCDF(x_{(i)})_{G\&I} = \frac{\left(\sum_{1 \leq j \leq i} w_{(j)} \right) - 0.5w_{(i)}}{N} \quad (10)$$

We can reformulate $CCDF(x_{(i)})_{G\&I}$ in the following way after defining $w_{(0)} = 0$:

$$CCDF(x_{(i)})_{AV} = \frac{\left(\sum_{1 \leq j \leq i} w_{(j)} \right) - 0.5w_{(i)}}{N} = \frac{2 \left(\sum_{1 \leq j \leq i} w_{(j)} \right) - w_{(i)}}{2N} = \frac{\sum_{1 \leq j \leq i} w_{(j-1)} + \sum_{1 \leq j \leq i} w_{(j)}}{2N} = \left[CCDF(x_{(i)})_a + CCDF(x_{(i)})_d \right] / 2 \quad (11)$$

Equation (11) can now be interpreted in two ways: First, we can interpret it as an average between the two previous definitions of the empirical CCDF. Second, we can interpret it as the CCDF obtained after averaging neighbouring pairs of weights after amending the descending weights vector by $w_{(0)} = 0$.

Figure 2: Different empirical CCDFs



Empirical CCDFs for a sample ($N=10$) from a univariate distribution over the interval $[0,15]$.

Another advantage of equation (11) is that it is straightforward to use it in conjunction with complex survey weights as the formula takes weights explicitly into account and thus it does not matter whether all weights are implicitly assumed to be equal to 1 $w = (w_1, \dots, w_n) = (1, \dots, 1)$ or they are allowed to be different for each observation: $w = (w_1, \dots, w_n)$ and $w_j \neq w_i$.

Figure 2 provides a plot of all three discussed variants of the empirical CCDF based on a sample of 10 observations drawn from a univariate distribution over the interval $[0,15]$. Weights were treated as being equal to 1 for each observation. Unsurprisingly, the figure illustrates the close relationship between these three concepts. First, it demonstrates that $CCDF(x_{(i)})_{AV}$ (purple line) is simply the average between $CCDF(x_{(i)})_d$ (grey line) and $CCDF(x_{(i)})_a$ (black line). Second, $CCDF(x_{(i)})_a$ and $CCDF(x_{(i)})_d$ are based on exactly the same probabilities but they assign these probabilities to different values of x_i . Putting it differently, they are shifted versions of each other.

An intuitive rationalization of the appropriateness of the formulation given in equation (11), starts by pointing out that there is a range of possible probabilities we can assign to any observation x_i when computing empirical CCDFs. In Figure 2, for instance, we have one observation with a value of 0 (out of ten observations overall), with the range of meaningful probabilities we can assign to that observation being $0.9 \leq P \leq 1$. Now, the differences between the three CCDFs plotted in Figure 2 resides in the exact assignment of these probabilities, where $CCDF_a$ and $CCDF_d$ assign the observation x_i to one of the extremes of this range, whereas $CCDF_{AV}$ positions the observation in the centre of the respective range.

Most importantly, $CCDF(x_{(i)})_{AV}$ can be readily used as the basis for fitting a Pareto distribution in every practically relevant case: 1) when using a given sample of (complex) survey data as it is, 2) as a starting point for applying Vermeulen's (2018) rich list approach or 2) to implement Wildauer

and Kapeller's (2019) rank correction approach. Effectively it allows to implement the G&I bias correction into all three approaches.

3.3 Goodness of fit statistics and complex survey weights

When fitting Pareto tails to available data, the obvious follow up question is to ask whether the estimated distribution represents a good fit to the data. In addition determining the scale parameter of the Pareto distribution not simply by graphical means but relying on Clauset et al.'s (2009) method requires a goodness of fit test. In both instances goodness of fit tests need to be adapted to deal with complex survey data. These results have been established in the literature already. Restating them here serves the purpose to provide a concise summary and guide for practitioners who might not be aware of the specialised statistical literature (Monahan, 2011; D'Agostino & Stephens, 1986).

We will begin with the Kolmogorov-Smirnov (KS from here on) goodness of fit test. The test statistic (T_{KS}) for unweighted data or trivial weights $w = (1, \dots, 1)$ is defined as follows (Monahan, 2011, p. 351):

$$T_{KS} = \sqrt{n}D \quad (12)$$

where D is the KS statistic and is defined as:

$$D = \sup_y |CDF_E(x) - CDF_T(x)| \quad (13)$$

which is the maximum distance between the empirical distribution function CDF_E and the theoretical distribution function CDF_T . With a data vector in descending order $x_d = (x_{(1)}, \dots, x_{(n)})$ with a corresponding vector of weights $w_a = (1, \dots, 1)$ and $w_0 = 0$ the empirical CDF is defined as:

$$CDF(x_{(i)})_a = 1 - \frac{\sum_{1 \leq j \leq i} w_{(j-1)}}{N} \quad (14)$$

For trivial weights the KS statistic can be computed as (Monahan, 2011, p. 351):

$$D = \max_i \left(\frac{i}{n} - CDF_T(x_{(i)}), CDF_T(x_{(i)}) - \frac{i-1}{n} \right) \quad (15)$$

where n represents the number of observations in the sample. For nontrivial weights $w_a = (w_{(1)}, \dots, w_{(n)})$ the KS statistic is defined as (Monahan, 2011, p. 358):

$$D = \max_i \left(CDF(x_{(i)})_a - CDF_T(x_{(i)}), CDF_T(x_{(i)}) - CDF(x_{(i-1)})_a \right) \quad (16)$$

The test statistic for nontrivial weights can then be obtained as (Monahan, 2011, p. 358):

$$T_{KS} = D \sqrt{\frac{N^2}{\sum_{i=1}^n w_{(i)}^2}} \quad (17)$$

For the Cramer-von-Mises (CvM from here on) goodness of fit test, the test statistic (T_{CvM}) for trivial weights is defined as (D'Agostino & Stephens, 1986, p. 101):

$$T_{CvM} = nW^2 = \frac{1}{12n} + \sum_{i=1}^n \left(CDF_T(x_{(i)}) - \frac{i-0.5}{n} \right)^2 \quad (18)$$

Deriving this expression relies on the the probability integral transformation. The CvM criterion (W^2) is defined as:

$$W^2 = \int_{-\infty}^{\infty} [CDF_E - CDF_T]^2 dCDF_T \quad (19)$$

We define $CDF_T(x_{(i)}) = U_i$ and the probability integral transformation allows us to rewrite the definition of the CvM test statistic as:

$$W^2 = \int_{-\infty}^{\infty} [CDF_E - U]^2 dU \quad (20)$$

which simplifies to equation 18 after multiplying by n .

Deriving the CvM test statistic for the case of complex survey weights is a bit more complicated. Starting from equation (20) we have to split up the integral:

$$W^2 = \int_{-\infty}^{x^{(1)}} [0 - U]^2 dU + \sum_{i=1}^{n-1} \int_{x^{(i)}}^{x^{(i+1)}} [CDF(x^{(i)})_E - U]^2 dU + \int_{x^{(n)}}^{\infty} [1 - U]^2 dU \quad (21)$$

After evaluating the integrals we obtain:

$$W^2 = \frac{1}{3}U_1^3 + \frac{1}{3} \sum_{i=1}^{n-1} \left[\left(CDF(x^{(i)})_a - U_i \right)^3 - \left(CDF(x^{(i)})_a - U_{i+1} \right)^3 \right] + \frac{1}{3} [1 - U_n]^3 \quad (22)$$

The test statistic is then obtained as

$$T_{CvM} = nW^2 \quad (23)$$

4 Conclusion

We show that Gabaix and Ibragimov's (2011) bias correction can be understood as using an average of the empirical CCDFs computed from data vectors in ascending and descending order. This does not only provide a simple and easy to understand rationale for the bias correction but more importantly provides a framework which can be directly generalized to the case of complex survey data (which comes with observation specific survey weights). In addition the proposed definition of the empirical CCDF can also easily be used in order to fit type I Pareto tails to the data based on Vermeulen's (2018) rich list approach or Wildauer and Kapeller's (2019) rank correction approach. In addition we summarise the key results for generalizing the Kolmogorov-Smirnov and Cramer-von-Mises test statistics for the complex survey case, which are needed to replace ad-hoc assumptions about the scale parameter of type I Pareto distributions by sound estimates (Clauset et al., 2009).

References

- Aigner, D. J., & Goldberger, A. S. (1970). Estimation of pareto's law from grouped observations. *Journal of the American Statistical Association*, 65(330), 712–723.
- Bach, S., Thiemann, A., & Zucco, A. (2018). Looking for the missing rich: Tracing the top tail of the wealth distribution. *DIW Berlin Discussion Paper*.
- Clauset, A., Shalizi, C. R., & Newman, M. E. J. (2009). Power-law distributions in empirical data. *SIAM Review*, 51(4), 661–703. doi: 10.1137/070710111
- D'Agostino, R. B., & Stephens, M. A. (1986). Goodness-of-fit-techniques. In D. B. Owen (Ed.), *Statistics: Textbooks and monographs* (Vol. 68). Marcel Dekker, Inc.
- Dalitz, C. (2018). Estimating wealth distribution: Top tail and inequality. *arXiv preprint arXiv:1807.03592*.
- Gabaix, X., & Ibragimov, R. (2011). Rank - 1/2: A simple way to improve the ols estimation of tail exponents. *Journal of Business & Economic Statistics*, 29(1), 24–39. doi: 10.1198/jbes.2009.06157
- Jayadev, A. (2008). A power law tail in india's wealth distribution: Evidence from survey data. *Physica A: Statistical Mechanics and its Applications*, 387(1), 270–276. doi: 10.1016/j.physa.2007.08.049
- Kratz, M., & Resnick, S. I. (1996, jan). The qq-estimator and heavy tails. *Communications in Statistics. Stochastic Models*, 12(4), 699–724. doi: 10.1080/15326349608807407
- Monahan, J. F. (2011). *Numerical methods of statistics* (2nd ed.). Cambridge University Press. doi: 10.1017/cbo9780511977176
- Vermeulen, P. (2018). How fat is the top tail of the wealth distribution? *Review of Income and Wealth*, 64(2), 357–387.
- Wildauer, R., & Kapeller, J. (2019). Rank correction: A new approach to differential nonresponse in wealth survey data.