

**On the Super-Additivity and Estimation Biases
Of Quantile Contributions**

Nassim Nicholas TALEB, Raphaël DOUADY

2014.90



On the Super-Additivity and Estimation Biases of Quantile Contributions

Nassim Nicholas Taleb*, Raphael Douady†

*School of Engineering, New York University

†Riskdata & C.N.R.S. Paris, Labex ReFi, Centre d'Economie de la Sorbonne

Abstract—Sample measures of top centile contributions to the total (concentration) are downward biased, unstable estimators, extremely sensitive to sample size and concave in accounting for large deviations. It makes them particularly unfit in domains with power law tails, especially for low values of the exponent. These estimators can vary over time and increase with the population size, as shown in this article, thus providing the illusion of structural changes in concentration. They are also inconsistent under aggregation and mixing distributions, as the weighted average of concentration measures for A and B will tend to be lower than that from $A \cup B$. In addition, it can be shown that under such fat tails, increases in the total sum need to be accompanied by increased sample size of the concentration measurement. We examine the estimation superadditivity and bias under homogeneous and mixed distributions.

Fourth version, Nov 11 2014

This work was achieved through the Laboratory of Excellence on Financial Regulation (Labex ReFi) supported by PRES heSam under the reference ANR-10-LABX-0095.

I. INTRODUCTION

Vilfredo Pareto noticed that 80% of the land in Italy belonged to 20% of the population, and vice-versa, thus both giving birth to the power law class of distributions and the popular saying 80/20. The self-similarity at the core of the property of power laws [?] and [?] allows us to recurse and reapply the 80/20 to the remaining 20%, and so forth until one obtains the result that the top percent of the population will own about 53% of the total wealth.

It looks like such a measure of concentration can be seriously biased, depending on how it is measured, so it is very likely that the true ratio of concentration of what Pareto observed, that is, the share of the top percentile, was closer to 70%, hence changes year-on-year would drift higher to converge to such a level from larger sample. In fact, as we will show in this discussion, for, say wealth, more complete samples resulting from technological progress, and also larger population and economic growth will make such a measure converge by increasing over time, for no other reason than expansion in sample space or aggregate value.

The core of the problem is that, for the class one-tailed fat-tailed random variables, that is, bounded on the left and unbounded on the right, where the random variable $X \in [x_{\min}, \infty)$, the in-sample quantile contribution is a biased estimator of the true value of the actual quantile contribution.

Let us define the *quantile contribution*

$$\kappa_q = q \frac{\mathbb{E}[X|X > h(q)]}{\mathbb{E}[X]}$$

where $h(q) = \inf\{h \in [x_{\min}, +\infty), \mathbb{P}(X > h) \leq q\}$ is the exceedance threshold for the probability q .

For a given sample $(X_k)_{1 \leq k \leq n}$, its "natural" estimator $\hat{\kappa}_q \equiv \frac{q^{th} \text{percentile}}{\text{total}}$, used in most academic studies, can be expressed, as

$$\hat{\kappa}_q \equiv \frac{\sum_{i=1}^n \mathbb{1}_{X_i > \hat{h}(q)} X_i}{\sum_{i=1}^n X_i}$$

where $\hat{h}(q)$ is the estimated exceedance threshold for the probability q :

$$\hat{h}(q) = \inf\{h : \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{x > h} \leq q\}$$

We shall see that the observed variable $\hat{\kappa}_q$ is a downward biased estimator of the true ratio κ_q , the one that would hold out of sample, and such bias is in proportion to the fatness of tails and, for very fat tailed distributions, remains significant, even for very large samples.

II. ESTIMATION FOR UNMIXED PARETO-TAILED DISTRIBUTIONS

Let X be a random variable belonging to the class of distributions with a "power law" right tail, that is:

$$\mathbb{P}(X > x) \sim L(x) x^{-\alpha} \quad (1)$$

where $L : [x_{\min}, +\infty) \rightarrow (0, +\infty)$ is a slowly varying function, defined as $\lim_{x \rightarrow +\infty} \frac{L(kx)}{L(x)} = 1$ for any $k > 0$.

There is little difference for small exceedance quantiles (<50%) between the various possible distributions such as Student's t, Lévy α -stable, Dagum,[?],[?] Singh-Maddala distribution [?], or straight Pareto.

For exponents $1 \leq \alpha \leq 2$, as observed in [?], the law of large numbers operates, though *extremely* slowly. The problem is acute for α around, but strictly above 1 and severe, as it diverges, for $\alpha = 1$.

A. Bias and Convergence

1) *Simple Pareto Distribution*: Let us first consider $\phi_\alpha(x)$ the density of a α -Pareto distribution bounded from below by $x_{\min} > 0$, in other words: $\phi_\alpha(x) = \alpha x_{\min}^\alpha x^{-\alpha-1} \mathbb{1}_{x \geq x_{\min}}$, and $\mathbb{P}(X > x) = \left(\frac{x_{\min}}{x}\right)^\alpha$. Under these assumptions, the cutpoint of exceedance is $h(q) = x_{\min} q^{-1/\alpha}$ and we have:

$$\kappa_q = \frac{\int_{h(q)}^{\infty} x \phi(x) dx}{\int_{x_{\min}}^{\infty} x \phi(x) dx} = \left(\frac{h(q)}{x_{\min}}\right)^{1-\alpha} = q^{\frac{\alpha-1}{\alpha}} \quad (2)$$

If the distribution of X is α -Pareto only beyond a cut-point x_{cut} , which we assume to be below $h(q)$, so that we have $\mathbb{P}(X > x) = \left(\frac{\lambda}{x}\right)^\alpha$ for some $\lambda > 0$, then we still have $h(q) = \lambda q^{-1/\alpha}$ and

$$\kappa_q = \frac{\alpha}{\alpha - 1} \frac{\lambda}{\mathbb{E}[X]} q^{\frac{\alpha-1}{\alpha}}$$

The estimation of κ_q hence requires that of the exponent α as well as that of the scaling parameter λ , or at least its ratio to the expectation of X .

Table I shows the bias of $\hat{\kappa}_q$ as an estimator of κ_q in the case of an α -Pareto distribution for $\alpha = 1.1$, a value chosen to be compatible with practical economic measures, such as the wealth distribution in the world or in a particular country, including developed ones.¹ In such a case, the estimator is extremely sensitive to "small" samples, "small" meaning in practice 10^8 . We ran up to a trillion simulations across varieties of sample sizes. While $\kappa_{0.01} \approx 0.657933$, even a sample size of 100 million remains severely biased as seen in the table.

Naturally the bias is rapidly (and nonlinearly) reduced for α further away from 1, and becomes weak in the neighborhood of 2 for a constant α , though not under a mixture distribution for α , as we shall see later. It is also weaker outside the top 1% centile, hence this discussion focuses on the famed "one percent" and on low values of the α exponent.

TABLE I: Biases of Estimator of $\kappa = 0.657933$ From 10^{12} Monte Carlo Realizations

$\hat{\kappa}(n)$	Mean	Median	STD across MC runs
$\hat{\kappa}(10^3)$	0.405235	0.367698	0.160244
$\hat{\kappa}(10^4)$	0.485916	0.458449	0.117917
$\hat{\kappa}(10^5)$	0.539028	0.516415	0.0931362
$\hat{\kappa}(10^6)$	0.581384	0.555997	0.0853593
$\hat{\kappa}(10^7)$	0.591506	0.575262	0.0601528
$\hat{\kappa}(10^8)$	0.606513	0.593667	0.0461397

In view of these results and of a number of tests we have performed around them, we can conjecture that the bias $\kappa_q - \hat{\kappa}_q(n)$ is "of the order of" $c(\alpha, q)n^{-b(q)(\alpha-1)}$ where constants $b(q)$ and $c(\alpha, q)$ need to be evaluated. Simulations suggest that $b(q) = 1$, whatever the value of α and q , but the rather slow convergence of the estimator and of its standard deviation to 0 makes precise estimation difficult.

2) *General Case:* In the general case, let us fix the threshold h and define:

$$\kappa_h = P(X > h) \frac{\mathbb{E}[X|X > h]}{\mathbb{E}[X]} = \frac{\mathbb{E}[X \mathbb{1}_{X > h}]}{\mathbb{E}[X]}$$

so that we have $\kappa_q = \kappa_{h(q)}$. We also define the n -sample estimator:

$$\hat{\kappa}_h \equiv \frac{\sum_{i=1}^n \mathbb{1}_{X_i > h} X_i}{\sum_{i=1}^n X_i}$$

where X_i are n independent copies of X . The intuition behind the estimation bias of κ_q by $\hat{\kappa}_q$ lies in a difference of concavity of the concentration measure with respect to

an innovation (a new sample value), whether it falls below or above the threshold. Let $A_h(n) = \sum_{i=1}^n \mathbb{1}_{X_i > h} X_i$ and $S(n) = \sum_{i=1}^n X_i$, so that $\hat{\kappa}_h(n) = \frac{A_h(n)}{S(n)}$ and assume a frozen threshold h . If a new sample value $X_{n+1} < h$ then the new value is $\hat{\kappa}_h(n+1) = \frac{A_h(n)}{S(n) + X_{n+1}}$. The value is convex in X_{n+1} so that uncertainty on X_{n+1} increases its expectation. At variance, if the new sample value $X_{n+1} > h$, the new value $\hat{\kappa}_h(n+1) \approx \frac{A_h(n) + X_{n+1} - h}{S(n) + X_{n+1} - h} = 1 - \frac{S(n) - A_h(n)}{S(n) + X_{n+1} - h}$, which is now concave in X_{n+1} , so that uncertainty on X_{n+1} reduces its value. The competition between these two opposite effects is in favor of the latter, because of a higher concavity with respect to the variable, and also of a higher variability (whatever its measurement) of the variable conditionally to being above the threshold than to being below. The fatter the right tail of the distribution, the stronger the effect. Overall, we find that $\mathbb{E}[\hat{\kappa}_h(n)] \leq \frac{\mathbb{E}[A_h(n)]}{\mathbb{E}[S(n)]} = \kappa_h$ (note that unfreezing the threshold $\hat{h}(q)$ also tends to reduce the concentration measure estimate, adding to the effect, when introducing one extra sample because of a slight increase in the expected value of the estimator $\hat{h}(q)$, although this effect is rather negligible). We have in fact the following:

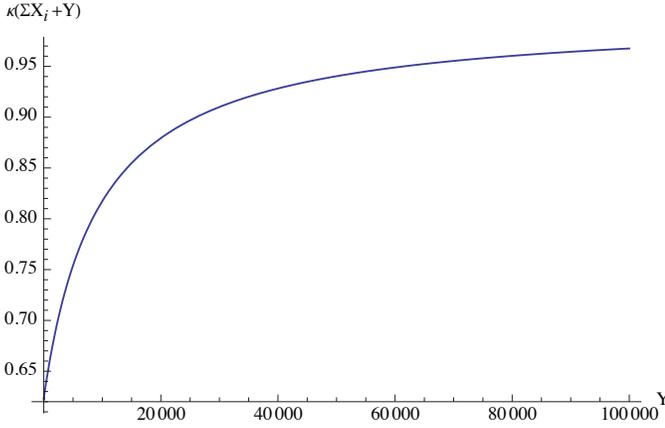
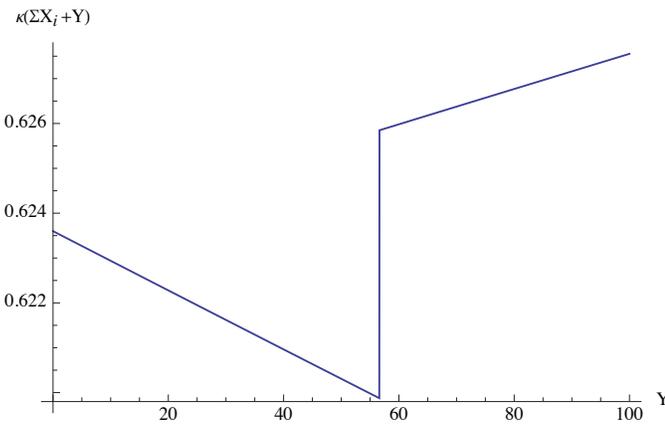
Proposition 1. Let $\mathbf{X} = (X)_{i=1}^n$ a random sample of size $n > \frac{1}{q}$, $Y = X_{n+1}$ an extra single random observation, and define: $\hat{\kappa}_h(\mathbf{X} \sqcup Y) = \frac{\sum_{i=1}^n \mathbb{1}_{X_i > h} X_i + \mathbb{1}_{Y > h} Y}{\sum_{i=1}^n X_i + Y}$. We remark that, whenever $Y > h$, one has:

$$\frac{\partial^2 \hat{\kappa}_h(\mathbf{X} \sqcup Y)}{\partial Y^2} \leq 0.$$

This inequality is still valid with $\hat{\kappa}_q$ as the value $\hat{h}(q, \mathbf{X} \sqcup Y)$ doesn't depend on the particular value of $Y > \hat{h}(q, \mathbf{X})$.

We face a different situation from the common small sample effect resulting from high impact from the rare observation in the tails that are less likely to show up in small samples, a bias which goes away by repetition of sample runs. The concavity of the estimator constitutes an upper bound for the measurement in finite n , clipping large deviations, which leads to problems of aggregation as we will state below in Theorem 1. In practice, even in very large sample, the contribution of very large rare events to κ_q slows down the convergence of the sample estimator to the true value. For a better, unbiased estimate, one would need to use a different path: first estimating the distribution parameters $(\hat{\alpha}, \hat{\lambda})$ and only then, estimating the theoretical tail contribution $\kappa_q(\hat{\alpha}, \hat{\lambda})$. Falk [?] observes that, even with a proper estimator of α and λ , the convergence is extremely slow, namely of the order of $n^{-\delta}/\ln n$, where the exponent δ depends on α and on the tolerance of the actual distribution vs. a theoretical Pareto, measured by the Hellinger distance. In particular, $\delta \rightarrow 0$ as $\alpha \rightarrow 1$, making the convergence really slow for low values of α .

¹This value, which is lower than the estimated exponents one can find in the literature – around 2 – is, following [?], a lower estimate which cannot be excluded from the observations.

Fig. 1: Effect of additional observations on κ Fig. 2: Effect of additional observations on κ , we can see convexity on both sides of h except for values of no effect to the left of h , an area of order $1/n$

III. AN INEQUALITY ABOUT AGGREGATING INEQUALITY

For the estimation of the mean of a fat-tailed r.v. $(X)_i^j$, in m sub-samples of size n_i each for a total of $n = \sum_{i=1}^m n_i$, the allocation of the total number of observations n between i and j does not matter so long as the total n is unchanged. Here the allocation of n samples between m sub-samples does matter because of the concavity of κ .² Next we prove that global concentration as measured by $\hat{\kappa}_q$ on a broad set of data will appear higher than local concentration, so aggregating European data, for instance, would give a $\hat{\kappa}_q$ higher than the average measure of concentration across countries – an “inequality about inequality”. In other words, we claim that the estimation bias when using $\hat{\kappa}_q(n)$ is even increased when dividing the sample into sub-samples and taking the weighted average of the measured values $\hat{\kappa}_q(n_i)$.

Theorem 1. Partition the n data into m sub-samples $N = N_1 \cup \dots \cup N_m$ of respective sizes n_1, \dots, n_m , with $\sum_{i=1}^m n_i = n$, and let S_1, \dots, S_m be the sum of variables over each sub-sample, and $S = \sum_{i=1}^m S_i$ be that over the whole sample.

²The same concavity – and general bias – applies when the distribution is lognormal, and is exacerbated by high variance.

Then we have:

$$\mathbb{E}[\hat{\kappa}_q(N)] \geq \sum_{i=1}^m \mathbb{E}\left[\frac{S_i}{S}\right] \mathbb{E}[\hat{\kappa}_q(N_i)]$$

If we further assume that the distribution of variables X_j is the same in all the sub-samples. Then we have:

$$\mathbb{E}[\hat{\kappa}_q(N)] \geq \sum_{i=1}^m \frac{n_i}{n} \mathbb{E}[\hat{\kappa}_q(N_i)]$$

In other words, averaging concentration measures of sub-samples, weighted by the total sum of each subsample, produces a downward biased estimate of the concentration measure of the full sample.

Proof. An elementary induction reduces the question to the case of two sub-samples. Let $q \in (0, 1)$ and (X_1, \dots, X_m) and (X'_1, \dots, X'_n) be two samples of positive i.i.d. random variables, the X_i 's having distributions $p(dx)$ and the X'_j 's having distribution $p'(dx')$. For simplicity, we assume that

both qm and qn are integers. We set $S = \sum_{i=1}^m X_i$ and

$S' = \sum_{i=1}^n X'_i$. We define $A = \sum_{i=1}^{mq} X_{[i]}$ where $X_{[i]}$ is the i -

th largest value of (X_1, \dots, X_m) , and $A' = \sum_{i=1}^{mq} X'_{[i]}$ where $X'_{[i]}$ is the i -th largest value of (X'_1, \dots, X'_n) . We also set

$S'' = S + S'$ and $A'' = \sum_{i=1}^{(m+n)q} X''_{[i]}$ where $X''_{[i]}$ is the i -th largest value of the joint sample $(X_1, \dots, X_m, X'_1, \dots, X'_n)$.

The q -concentration measure for the samples $\mathbf{X} = (X_1, \dots, X_m)$, $\mathbf{X}' = (X'_1, \dots, X'_n)$ and $\mathbf{X}'' = (X_1, \dots, X_m, X'_1, \dots, X'_n)$ are:

$$\kappa = \frac{A}{S} \quad \kappa' = \frac{A'}{S'} \quad \kappa'' = \frac{A''}{S''}$$

We must prove that the following inequality holds for expected concentration measures:

$$\mathbb{E}[\kappa''] \geq \mathbb{E}\left[\frac{S}{S''}\right] \mathbb{E}[\kappa] + \mathbb{E}\left[\frac{S'}{S''}\right] \mathbb{E}[\kappa']$$

We observe that:

$$A = \max_{\substack{J \subset \{1, \dots, m\} \\ |J| = \theta m}} \sum_{i \in J} X_i$$

and, similarly $A' = \max_{J' \subset \{1, \dots, n\}, |J'| = qn} \sum_{i \in J'} X'_i$ and $A'' = \max_{J'' \subset \{1, \dots, m+n\}, |J''| = q(m+n)} \sum_{i \in J''} X_i$, where we have denoted $X_{m+i} = X'_i$ for $i = 1 \dots n$. If $J \subset \{1, \dots, m\}$, $|J| = \theta m$ and $J' \subset \{m+1, \dots, m+n\}$, $|J'| = qn$, then $J'' = J \cup J'$ has cardinal $m+n$, hence $A + A' = \sum_{i \in J''} X_i \leq A''$, whatever the particular sample. Therefore $\kappa'' \geq \frac{S}{S''} \kappa + \frac{S'}{S''} \kappa'$ and we have:

$$\mathbb{E}[\kappa''] \geq \mathbb{E}\left[\frac{S}{S''} \kappa\right] + \mathbb{E}\left[\frac{S'}{S''} \kappa'\right]$$

Let us now show that:

$$\mathbb{E}\left[\frac{S}{S''} \kappa\right] = \mathbb{E}\left[\frac{A}{S''}\right] \geq \mathbb{E}\left[\frac{S}{S''}\right] \mathbb{E}\left[\frac{A}{S}\right]$$

If this is the case, then we identically get for κ' :

$$\mathbb{E} \left[\frac{S'}{S''} \kappa' \right] = \mathbb{E} \left[\frac{A'}{S''} \right] \geq \mathbb{E} \left[\frac{S'}{S''} \right] \mathbb{E} \left[\frac{A'}{S'} \right]$$

hence we will have:

$$\mathbb{E} [\kappa''] \geq \mathbb{E} \left[\frac{S}{S''} \right] \mathbb{E} [\kappa] + \mathbb{E} \left[\frac{S'}{S''} \right] \mathbb{E} [\kappa']$$

Let $T = X_{[mq]}$ be the cut-off point (where $[mq]$ is the integer part of mq), so that $A = \sum_{i=1}^m X_i \mathbb{1}_{X_i \geq T}$ and let $B = S - A = \sum_{i=1}^m X_i \mathbb{1}_{X_i < T}$. Conditionally to T , A and B are independent: A is a sum of $m\theta$ samples constrained to being above T , while B is the sum of $m(1-\theta)$ independent samples constrained to being below T . They are also independent of S' . Let $p_A(t, da)$ and $p_B(t, db)$ be the distribution of A and B respectively, given $T = t$. We recall that $p'(ds')$ is the distribution of S' and denote $q(dt)$ that of T . We have:

$$\mathbb{E} \left[\frac{S}{S''} \kappa \right] = \iint \frac{a+b}{a+b+s'} \frac{a}{a+b} p_A(t, da) p_B(t, db) q(dt) p'(ds')$$

For given b , t and s' , $a \rightarrow \frac{a+b}{a+b+s'}$ and $a \rightarrow \frac{a}{a+b}$ are two increasing functions of the same variable a , hence conditionally to T , B and S' , we have:

$$\begin{aligned} \mathbb{E} \left[\frac{S}{S''} \kappa \middle| T, B, S' \right] &= \mathbb{E} \left[\frac{A}{A+B+S'} \middle| T, B, S' \right] \\ &\geq \mathbb{E} \left[\frac{A+B}{A+B+S'} \middle| T, B, S' \right] \mathbb{E} \left[\frac{A}{A+B} \middle| T, B, S' \right] \end{aligned}$$

This inequality being valid for any values of T , B and S' , it is valid for the unconditional expectation, and we have:

$$\mathbb{E} \left[\frac{S}{S''} \kappa \right] \geq \mathbb{E} \left[\frac{S}{S''} \right] \mathbb{E} \left[\frac{A}{S} \right]$$

If the two samples have the same distribution, then we have:

$$\mathbb{E} [\kappa''] \geq \frac{m}{m+n} \mathbb{E} [\kappa] + \frac{n}{m+n} \mathbb{E} [\kappa']$$

Indeed, in this case, we observe that $\mathbb{E} \left[\frac{S}{S''} \right] = \frac{m}{m+n}$. Indeed $S = \sum_{i=1}^m X_i$ and the X_i are identically distributed, hence $\mathbb{E} \left[\frac{S}{S''} \right] = m \mathbb{E} \left[\frac{X}{S''} \right]$. But we also have $\mathbb{E} \left[\frac{S''}{S''} \right] = 1 = (m+n) \mathbb{E} \left[\frac{X}{S''} \right]$ therefore $\mathbb{E} \left[\frac{X}{S''} \right] = \frac{1}{m+n}$. Similarly, $\mathbb{E} \left[\frac{S'}{S''} \right] = \frac{n}{m+n}$, yielding the result.

This ends the proof of the theorem. \square

Let X be a positive random variable and $h \in (0, 1)$. We remind the theoretical h -concentration measure, defined as:

$$\kappa_h = \frac{P(X > h) \mathbb{E}[X | X > h]}{\mathbb{E}[X]}$$

whereas the n -sample θ -concentration measure is $\widehat{\kappa}_h(n) = \frac{A(n)}{S(n)}$, where $A(n)$ and $S(n)$ are defined as above for an n -sample $\mathbf{X} = (X_1, \dots, X_n)$ of i.i.d. variables with the same distribution as X .

Theorem 2. For any $n \in \mathbb{N}$, we have:

$$\mathbb{E} [\widehat{\kappa}_h(n)] < \kappa_h$$

and

$$\lim_{n \rightarrow +\infty} \widehat{\kappa}_h(n) = \kappa_h \quad \text{a.s. and in probability}$$

Proof. The above corollary shows that the sequence $n \mathbb{E} [\widehat{\kappa}_h(n)]$ is super-additive, hence $\mathbb{E} [\widehat{\kappa}_h(n)]$ is an increasing sequence. Moreover, thanks to the law of large numbers, $\frac{1}{n} S(n)$ converges almost surely and in probability to $\mathbb{E}[X]$ and $\frac{1}{n} A(n)$ converges almost surely and in probability to $\mathbb{E}[X \mathbb{1}_{X > h}] = P(X > h) \mathbb{E}[X | X > h]$, hence their ratio also converges almost surely to κ_h . On the other hand, this ratio is bounded by 1. Lebesgue dominated convergence theorem concludes the argument about the convergence in probability. \square

IV. MIXED DISTRIBUTIONS FOR THE TAIL EXPONENT

Consider now a random variable X , the distribution of which $p(dx)$ is a mixture of parametric distributions with different values of the parameter: $p(dx) = \sum_{i=1}^m \omega_i p_{\alpha_i}(dx)$. A typical n -sample of X can be made of $n_i = \omega_i n$ samples of X_{α_i} with distribution p_{α_i} . The above theorem shows that, in this case, we have:

$$\mathbb{E} [\widehat{\kappa}_q(n, X)] \geq \sum_{i=1}^m \mathbb{E} \left[\frac{S(\omega_i n, X_{\alpha_i})}{S(n, X)} \right] \mathbb{E} [\widehat{\kappa}_q(\omega_i n, X_{\alpha_i})]$$

When $n \rightarrow +\infty$, each ratio $\frac{S(\omega_i n, X_{\alpha_i})}{S(n, X)}$ converges almost surely to ω_i respectively, therefore we have the following convexity inequality:

$$\kappa_q(X) \geq \sum_{i=1}^m \omega_i \kappa_q(X_{\alpha_i})$$

The case of Pareto distribution is particularly interesting. Here, the parameter α represents the tail exponent of the distribution. If we normalize expectations to 1, the cdf of X_α is $F_\alpha(x) = 1 - \left(\frac{x}{x_{\min}} \right)^{-\alpha}$ and we have:

$$\kappa_q(X_\alpha) = q^{\frac{\alpha-1}{\alpha}}$$

and

$$\frac{d^2}{d\alpha^2} \kappa_q(X_\alpha) = q^{\frac{\alpha-1}{\alpha}} \frac{(\log q)^2}{\alpha^3} > 0$$

Hence $\kappa_q(X_\alpha)$ is a convex function of α and we can write:

$$\kappa_q(X) \geq \sum_{i=1}^m \omega_i \kappa_q(X_{\alpha_i}) \geq \kappa_q(X_{\bar{\alpha}})$$

where $\bar{\alpha} = \sum_{i=1}^m \omega_i \alpha$.

Suppose now that X is a positive random variable with unknown distribution, except that its tail decays like a power law with unknown exponent. An unbiased estimation of the exponent, with necessarily some amount of uncertainty (i.e., a distribution of possible true values around some average), would lead to a downward biased estimate of κ_q .

Because the concentration measure only depends on the tail of the distribution, this inequality also applies in the case of a mixture of distributions with a power decay, as in Equation 1:

$$\mathbb{P}(X > x) \sim \sum_{j=1}^N \omega_j L_j(x) x^{-\alpha_j} \quad (3)$$

The slightest uncertainty about the exponent increases the concentration index. One can get an actual estimate of this bias by considering an average $\bar{\alpha} > 1$ and two surrounding values $\alpha^+ = \alpha + \delta$ and $\alpha^- = \alpha - \delta$. The convexity inequality writes as follows:

$$\kappa_q(\bar{\alpha}) = q^{1-\frac{1}{\bar{\alpha}}} < \frac{1}{2} \left(q^{1-\frac{1}{\alpha+\delta}} + q^{1-\frac{1}{\alpha-\delta}} \right)$$

So in practice, an estimated $\bar{\alpha}$ of around $3/2$, sometimes called the "half-cubic" exponent, would produce similar results as value of α much closer to 1, as we used in the previous section. Simply $\kappa_q(\alpha)$ is convex, and dominated by the second order effect $\frac{\ln(q)q^{1-\frac{1}{\alpha+\delta}}(\ln(q)-2(\alpha+\delta))}{(\alpha+\delta)^4}$, an effect that is exacerbated at lower values of α .

To show how unreliable the measures of inequality concentration from quantiles, consider that a standard error of 0.3 in the measurement of α causes $\kappa_q(\alpha)$ to rise by 0.25.

V. A LARGER TOTAL SUM IS ACCOMPANIED BY INCREASES IN $\hat{\kappa}_q$

There is a large dependence between the estimator $\hat{\kappa}_q$ and the sum $S = \sum_{j=1}^n X_j$: conditional on an increase in $\hat{\kappa}_q$ the expected sum is larger. Indeed, as shown in theorem 1, $\hat{\kappa}_q$ and S are positively correlated.

For the case in which the random variables under concern are wealth, we observe as in Figure 3 such conditional increase; in other words, since the distribution is of the class of fat tails under consideration, the maximum is of the same order as the sum, additional wealth means more measured inequality. Under such dynamics, is quite absurd to assume that additional wealth will arise from the bottom or even the middle. (The same argument can be applied to wars, epidemics, size or companies, etc.)

VI. CONCLUSION AND PROPER ESTIMATION OF CONCENTRATION

Concentration can be high at the level of the generator, but in small units or subsections we will observe a lower κ_q . So examining times series, we can easily get a historical illusion of rise in, say, wealth concentration when it has been there all along at the level of the process; and an expansion in the size of the unit measured can be part of the explanation.³

Even the estimation of α can be biased in some domains where one does not see the entire picture: in the presence of uncertainty about the "true" α , it can be shown that, unlike other parameters, the one to use is not the probability-weighted

³Accumulated wealth is typically thicker tailed than income, see [?].

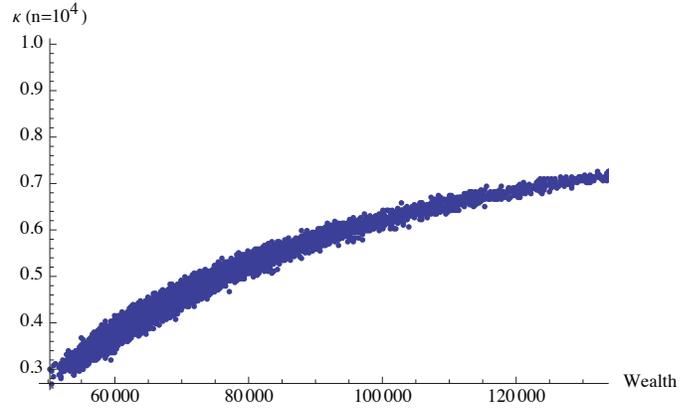


Fig. 3: Effect of additional wealth on $\hat{\kappa}$

exponents (the standard average) but rather the minimum across a section of exponents [?].

One must not perform analyses of year-on-year changes in $\hat{\kappa}_q$ without adjustment. It did not escape our attention that some theories are built based on claims of such "increase" in inequality, as in [?], without taking into account the true nature of κ_q , and promulgating theories about the "variation" of inequality without reference to the stochasticity of the estimation – and the lack of consistency of κ_q across time and sub-units. What is worse, rejection of such theories also ignored the size effect, by countering with data of a different sample size, effectively making the dialogue on inequality uninformatively statistically.⁴

The mistake appears to be commonly made in common inference about fat-tailed data in the literature. The very methodology of using concentration and changes in concentration is highly questionable. For instance, in the thesis by Steven Pinker [?] that the world is becoming less violent, we note a fallacious inference about the concentration of damage from wars from a $\hat{\kappa}_q$ with minutely small population in relation to the fat-tailedness.⁵ Owing to the fat-tailedness of war casualties and consequences of violent conflicts, an adjustment would rapidly invalidate such claims that violence from war has statistically experienced a decline.

A. Robust methods and use of exhaustive data

We often face argument of the type "the method of measuring concentration from quantile contributions $\hat{\kappa}$ is robust and based on a complete set of data". Robust methods, alas, tend to fail with fat-tailed data, see [?]. But, in addition, the problem here is worse: even if such "robust" methods were deemed unbiased, a method of direct centile estimation is still linked to a static and specific population and does not

⁴Financial Times, May 23, 2014 "Piketty findings undercut by errors" by Chris Giles.

⁵Using Richardson's data, [?]: "(Wars) followed an 80:2 rule: almost eighty percent of the deaths were caused by *two percent* (his emph.) of the wars". So it appears that both Pinker and the literature cited for the quantitative properties of violent conflicts are using a flawed methodology, one that produces a severe bias, as the centile estimation has extremely large biases with fat-tailed wars. Furthermore claims about the mean become spurious at low exponents.

aggregate. Accordingly, such techniques do not allow us to make statistical claims or scientific statements about the true properties which should necessarily carry out of sample.

Take an insurance (or, better, reinsurance) company. The "accounting" profits in a year in which there were few claims do not reflect on the "economic" status of the company and it is futile to make statements on the concentration of losses per insured event based on a single year sample. The "accounting" profits are not used to predict variations year-on-year, rather the exposure to tail (and other) events, analyses that take into account the stochastic nature of the performance. This difference between "accounting" (deterministic) and "economic" (stochastic) values matters for policy making, particularly under fat tails. The same with wars: we do not estimate the severity of a (future) risk based on past in-sample historical data.

B. How Should We Measure Concentration?

Practitioners of risk managers now tend to compute CVaR and other metrics, methods that are extrapolative and nonconcave, such as the information from the α exponent, taking the one closer to the lower bound of the range of exponents, as we saw in our extension to Theorem 2 and rederiving the corresponding κ , or, more rigorously, integrating the functions of α across the various possible states. Such methods of adjustment are less biased and do not get mixed up with problems of aggregation –they are similar to the "stochastic volatility" methods in mathematical finance that consist in adjustments to option prices by adding a "smile" to the standard deviation, in proportion to the variability of the parameter representing volatility and the errors in its measurement. Here it would be "stochastic alpha" or "stochastic tail exponent"⁶ By extrapolative, we mean the built-in extension of the tail in the measurement by taking into account realizations outside the sample path that are in excess of the extrema observed.^{7 8}

ACKNOWLEDGMENT

The late Benoît Mandelbrot, Branko Milanovic, Dominique Guégan, Felix Salmon, Bruno Dupire, the late Marc Yor, Albert Shiryayev, an anonymous referee, the staff at Luciano Restaurant in Brooklyn and Naya in Manhattan.

⁶Also note that, in addition to the centile estimation problem, some authors such as [?] when dealing with censored data, use Pareto interpolation for insufficient information about the tails (based on tail parameter), filling-in the bracket with conditional average bracket contribution, which is not the same thing as using full power-law extension; such a method retains a significant bias.

⁷Even using a lognormal distribution, by fitting the scale parameter, works to some extent as a rise of the standard deviation extrapolates probability mass into the right tail.

⁸We also note that the theorems would also apply to Poisson jumps, but we focus on the powerlaw case in the application, as the methods for fitting Poisson jumps are interpolative and have proved to be easier to fit in-sample than out of sample, see [?].