Department of Economics
Johannes Kepler University of Linz

Correcting wealth survey data for the missing rich: The case of Austria

by

Paul ECKERSTORFER
Johannes HALAK
Jakob KAPELLER
Bernhard SCHÜTZ[*)]
Florian SPRINGHOLZ
Rafael WILDAUER

Working Paper No. 1401
January 2014

Johannes Kepler University of Linz
Department of Economics
Altenberger Strasse 69
A-4040 Linz - Auhof, Austria
www.econ.jku.at

bernhard.schuetz@jku.at
phone: +43 732 2468 8584

Correcting wealth survey data for the missing rich: The case of Austria[*]

Paul Eckerstorfer**, Johannes Halak†, Jakob Kapeller††, Bernhard Schütz‡, Florian Springholz†† and Rafael Wildauer‡‡

** Parliamentary Budget Office, National Council of Austria

† Government of the Federal State of Upper Austria

†† Department of Philosophy and Theory of Science, University of Linz

‡ Department of Economics, University of Linz

‡‡ Department of Economics, Kingston University London

Abstract

It is a well-known criticism that due to its exponential distribution, survey data on wealth is hardly reliable when it comes to analyzing the richest parts of society. This paper addresses this criticism using Austrian data from the *Household Finance and Consumption Survey (HFCS)*. In doing so we apply the assumption of a Pareto distribution to obtain estimates for the number of households possessing a net wealth greater than four million Euros as well as their aggregate wealth holdings. Thereby, we identify suitable parameter combinations by using a series of maximum-likelihood estimates and appropriate goodness-of-fit tests to avoid arbitrariness with respect to the fitting of the Pareto-Distribution. Our results suggest that the alleged non-observation bias is considerable, accounting for about one quarter of total net wealth. The method developed in this paper can easily be applied to other countries where survey data on wealth are available.

JEL: C46, C81, D31

Keywords: wealth distribution, non-observation bias, Pareto distribution

[*] We thank Bettina Csoka, Primin Fessler, Peter Lindner, Paul Hofmarcher, Matthias Schnetzer and Stefan Steinerberger for valuable comments. This paper was supported by the Austrian Chamber of Labor and uses data from the Eurosystem Household Finance and Consumption Survey (HFCS). The opinions expressed in this article are the author's own and do not reflect the view of the Parliamentary Budget Office of the National Council of Austria or the Government of the Federal State of Upper Austria.

1. Introduction

Social stability and the accumulation of private wealth share an ambiguous relationship. From a historical perspective, social stability is a central prerequisite for the emergence of a significant amount of private wealth (Atkinson *et al*., 2011, Borgerhoff Mulder *et al*., 2009). However, since the accumulation of wealth is also self-enforcing – those who already possess wealth on average also experience larger gains – we find evidence for a successive concentration of wealth, which in turn undermines social and economic stability (Guttmann and Plihon, 2010, Stiglitz, 2012). In their work on the relation of inequality and social outcomes of various sorts, Wilkinson and Pickett (2007, 2009) explicitly address this relationship between social cohesion and stability on the one, and distributional outcomes on the other hand. This relation is also addressed by pointing to the causal link between higher inequality and social problems such as obesity, teenager pregnancies, low life expectancy, mental diseases, suicide, xenophobia, low educational achievement and imprisonment.

In this context the improvement of computer-based evaluation methods and the increasing availability of survey data allow for a better assessment of wealth inequality. Moreover, standardized surveys like the *Household Finance and Consumption Survey* (HFCS) carried out by the European Central Bank (ECB) increase the comparability of the results between countries. For many countries of the Euro Zone, including Austria, the HFCS is the most comprehensive household-level assessment of private wealth and its disaggregated components (financial assets, tangible assets, debt etc,) that has been conducted so far, providing new and unprecedented opportunities of insight for these countries[1].

Nevertheless there are also problems concerning survey data, among those being false or no responses by the targeted households. Since the probability of false responses as well as the weight of the non-observation bias increase within the segment of the richest households, surveys are alleged to give an incomplete picture of the top of the wealth distribution. This paper aims to correct for the second bias – the case of non-observation. While the probability of non-observation is basically the same for all households in the

[1] Previous studies on wealth and income in Austria face the problem of insufficient data. Therefore, most of them necessarily cover only subcomponents of total wealth (e.g. financial wealth). For instance, Beer et al. (2006) analyse Austrian National Bank survey data on financial wealth from 2004 and find that the average amount of financial wealth held by Austrian households is 51,790 Euros (median: 21,855 Euros). Hahn and Magerl (2006) estimate the distribution of wealth and its components through the disaggregation of national accounts data leading to results largely consistent with those in Beer et al. (2006). The corresponding values to the results of Beer et al. (2006) in the HFCS are 47,991 Euros for the average and 14,071 for the median. However, a direct comparison between these two studies is difficult due to the significant time-lag as well as different definitions of financial wealth (e.g. in the HFCS the value assigned to self-employed business is not counted as financial wealth).

population and independent of the level of net wealth, only omissions at the top cause a significant bias, since in the segment of the very rich (those within the upper 0.5%-1%) only very few households decisively influence estimations for total wealth and wealth inequality due to the exponential properties of the underlying distribution (Hoeller et al., 2012, Avery *et al.* , 1986). In short: Wealth at the very top is so skewed, that the few households drawn from this segment are not considered as representative for the underlying population. While the survey design of the HFCS anticipates many known issues with survey data (like inconsistent responses or unavailability of households), the issue of underrepresentation of the wealthiest is not addressed by the Austrian variant of the HFCS.[2]

In order to deal with this problem, we employ a novel methodical strategy by using the assumption of a Pareto distribution in combination with a non-arbitrary strategy to define appropriate minimal values for the latter (Clauset et al., 2009). In turn we reapply the Pareto-Distribution to modify the original data by replacing the far-right tail of the existing HFCS data (roughly: the upper 0.5%-1%) with an estimation based on the Pareto distribution itself. While the assumption of a Pareto distribution is widely used in studying the distribution of income (see e.g. Feenberg and Poterba, 1993, Piketty, 2003, Piketty and Saez, 2003) and wealth (Alvaredo and Saez, 2009, Bach and Beznoska, 2011, Bach, Beznoska and Steiner, 2014, Durán-Cabré and Esteller-Moré, 2010, Kopczuk and Saez, 2004), we employ this assumption with respect to the specific task to correct the alleged non-observation bias inherent in survey data on wealth. Additionally, two subtle differences of our approach are related to the kind of data used and the technical aspects of estimating the Pareto distribution: With regard to the former we do not use tax returns data as most recent studies do (Alvaredo and Saez, 2009, Atkinson et al. 2011, Durán-Cabré and Esteller-Moré, 2010, Piketty and Saez 2003, Kopcuk and Saez 2004, Alvaredo et al., 2013), since such sources simply do not exist in the case of Austria (and, additionally, they vary in quality across countries due to different tax systems if existent). With regard to the latter the determination of the distribution parameters (especially the cut-off point) does not rely on theoretical considerations or a graphical identification based on a log-log graph (Cowell, 2011, Bach and Beznoska, 2011, Bach et al., 2014) but on solid statistical hypothesis testing (Clauset et al., 2009). In doing so, we find that the underrepresentation of wealthy households in the Austrian HFCS sample is significant. After taking the Pareto correction into account, aggregate net wealth increases by 27% and the Gini coefficient rises from 0.76 to 0.81, indicating the importance of correcting this bias.

---

[2] For these reasons some countries participating in the HFCS implemented an oversampling strategy in order to dampen the resulting bias. Among these are Belgium, Germany, Greece, Spain, France, Cyprus, Luxembourg, Portugal and Finland (cf. ECB, 2013, 10)

The remaining paper is structured as follows: Section 2 describes the characteristics of the HFCS data and provides some descriptive statistics. The third section explains the method used to correct the original data and tests the validity of the underlying Pareto assumption. Section 4 presents the corrected data and compares the resulting descriptive statistics to the original HFCS statistics from section 2. Section 5 assesses the robustness of our findings by comparing them to non-estimated data and presents upper and lower bounds with respect to irregularities and outliers within the dataset. The final section concludes.

2. Wealth in Austria: First evidence from the HFCS

The aim of this section is to provide an overview of the HFCS method and data. Selected descriptive statistics provide a first glance on the distribution of wealth in Austria according to HFCS data. The results of the HFCS are thereby put in context to previous studies of the distribution of wealth in Austria.

2.1. Household Finance and Consumption Survey (HFCS)

The HFCS is the first comprehensive survey on tangible assets, financial wealth, liabilities and expenditures of private households in 15 countries of the Euro Zone (missing countries are Ireland and Estonia). Especially the potential comparability of the data, possible due to the ex-ante harmonization of the applied method in the various countries, potentially allows for worthwhile comparison between the Euro Zone countries. In Austria the Austrian National Bank conducted the survey in cooperation with the Institute for Empirical Social Studies (IFES). In what follows we give a brief overview of the survey design (for more details see Albacete et al., 2012).

The basic reporting unit in the HFCS is the household, which is represented by the single person within the household, who felt most competent with regard to the household finances. The survey is based on personal interviews conducted between September 2010 and May 2011. The initial sample consisted of 4,436 households. Eventually, 2,380 households have been successfully interviewed, indicating a response rate of roughly 56%. The sample selection is based on a *stratified multistage cluster random sampling,* implying that a multiple-stage form of quota sampling was applied to ensure that the randomly drawn participants adequately reflect the composition of the Austrian population. Stratification was based on Austrian NUTS-3 regions and municipality size in order to assure that households from different regions enter the sample proportionally. Data collection was based on computer-assisted interviews.

It is typical for surveys, especially for those that try to evaluate sensitive information such as wealth, income or debt, that participants refuse or are unable to answer certain questions, which can bias the results. In order to reduce such a bias, missing values were inserted ex post using *multiple imputation*. In this process, missing values are replaced through estimated values. This preserves the correlation structure of the dataset since one does not have to drop all incomplete observations. Imputation was repeated five times, producing five different samples – the so called *implicates*. Systematically biased answers (for example a higher proportion of wrong answers in certain segments of the wealth distribution) and the problem of non-observation of very rich households (problem of *coverage*), however, cannot be compensated for with this method.

Finally, each observation in the dataset received a probability weight in order to reduce the sample variance and to adjust the sample to the statistical population. These weights try to correct for the different probabilities to be part of the observed sample (*unequal probability sampling bias*), the possible incompleteness of the target population (*frame bias*) and non-response of targeted households (*non-response bias*). Hence, the final weights emerge from the *design weights*, which account for the unequal probability to be part of the sample, the *post-stratification weights*, which try to correct for erroneous exclusion (e.g. wrong postal address) of households, and the *non-response weights*.

2.2. The Distribution of Wealth According to HFCS Data

In this subsection we provide several descriptive statistics based on the original HFCS data to give a first impression regarding the wealth distribution in Austria. For a more exhaustive evaluation of these data see the much more detailed report by the Austrian National Bank (Albacete et al., 2012). In section 4 we replicate some of these statistics to illustrate the difference between the original data and the corrected sample. Each illustration accounts for all five implicates as well as the respective sample weights.

Figure 1 shows the distribution of households along different ranges of net wealth. It reveals that only a small fraction of Austrians are net debtors, that the relative majority of Austrians (roughly 40%) have acquired a modest net wealth between 0 and 50,000 Euros and that a little more than the highest decile has acquired a net wealth exceeding 500,000 Euros.

<PLACE FIGURE 1 HERE>

While Figure 1 gives a first impression of the Austrian distribution of wealth, it remains relatively silent on the (relative) amount of net wealth concentrated at the top of the distribution. In this context an examination of the cumulative distribution function is more revealing: The distribution function depicted in Figure 2 is a partial replication of work done by the Austrian National Bank (the threshold values can also be found in Albacete et al., 2012, 66). As expected the distribution function indicates that the Austrian distribution of wealth takes the typical exponential form. However, we also observe that the richest households found in our sample possess a net wealth of only 15 million Euros, strongly pointing to a lack of data with regard to very rich households. This observation, of course, reinforces the validity of the critique that the non-observation of very rich households might lead to a significant bias in aggregate estimates.[3] Additionally, Figure 2 gives a first intuition how the right tail of the distribution is different compared to the rest of the distribution.

<PLACE FIGURE 2 HERE>

A closer look at the top of the distribution is even more informative with regard to the question at hand: Table 1 shows aggregate as well as average wealth for the richest 5 percentiles as well as an estimate for total wealth and the respective shares. Here, we observe that the richest percentile owns roughly 23% of total net wealth, whereas the cumulated wealth of the richest five percent amounts to 47.6% of total wealth. A full percentile list can be obtained from the authors.

<PLACE TABLE 1 HERE>

The corresponding Gini-coefficient of the distribution of wealth in Austria as obtained from the original data is equal to 0.762, which represents a rather unequal distribution of private wealth from a global perspective (Davies et al., 2007). In the next section we give a detailed description of our methodological strategy and the practical steps necessary to implement this strategy.

3. Estimating the Missing Rich: A Methodological Suggestion

As already mentioned in the previous section, the study design of the HFCS anticipates various well-known problems of survey data on household wealth. However, the problem of *non-observation* remains crucial: While the richest households in Austria possess much more than a billion Euros, the richest household in the HFCS sample lies somewhere between 11.8 million Euros and 22 million Euros within the five implicates. This section provides a statistical approach to compensate for this gap and thereby to improve data reliability on the top end of the distribution.

The approach is based on three distinct steps: first, we fit a Pareto distribution to the upper tail of the sample. Second, we apply a bootstrap procedure in order to test the validity of the distribution estimated in the first step. These two steps closely follow the method suggested in Clauset et al. (2009). In an additional third step, we eliminate any observation with

[3] According to Austrian counterpart of the (international) Forbes-list there are roughly thirty billionaires in Austria with an estimated total wealth of roughly 100 billion Euros (cf. http://www.trendtop500.at/die-reichsten-oesterreicher/).

reported net wealth exceeding 4 million Euros and replace them with data points suggested by the previously estimated Pareto distribution. Figure 3 provides a graphical illustration of these three steps. The elimination mainly affects observations from the 100th percentile and accounts for the fact that the sample does not contain observations on net wealth exceeding 20 million Euros but claims to be representative for the whole population. In contrast the corrected sample emerging from this three step procedure is more likely to represent the population since the unrealistic truncation of the sample at 20 million Euros net wealth is removed. The assumption that the data follow a Pareto distribution in the first step is not a restrictive one, since we are only interested in the upper tail. Commonly used alternative distributions such as Dagum or Singh-Maddala are designed to describe the entire interval of positive household wealth and the upper tail of those distributions converges towards a Pareto distribution. Dagum (2006) even considers this latter property as an essential characteristic for any distribution used to describe the behavior of household wealth.

<PLACE FIGURE 3 HERE>

3.1 Estimating the Distribution Parameters

Generally, a Pareto distribution has to be established for each implicate. The associated cumulative distribution function is denoted in the following way.

$$P_i(x_i) = \Pr(X_i \leq x_i) = 1 - \left(\frac{m_i}{x_i}\right)^{\alpha_i} \forall\ implicates\ i = 1 \ldots 5\ \wedge\ x_i \geq m_i, \tag{1}$$

In this context x_i represents observed net wealth of a given household, m_i the hitherto unknown cut-off point above which the sample data can be described using a Pareto distribution and, finally, α_i is a shape parameter describing the specific form of the underlying power distribution (*Pareto alpha*). Our approach closely follows Clauset et al. (2009) and can be summarized in the following way: $\hat{m}_i$ and $\hat{\alpha}_i$ are determined by estimating Pareto distributions systematically for increasing subsets of the data and choosing that subset and its corresponding parameters which exhibits the best fit to the data. We employ Cramer-Von-Mises (CvM) test-statistics to compare the relative fit of the estimated distributions.[4] The Mathematica code used to carry out the steps described below can be obtained from the authors upon request.

First, we fit Pareto distributions by maximum likelihood to increasingly large sub-samples starting from the 100th percentile. Through expanding the sub-sample by one additional percentile until we reach the 71st percentile, we get 30 different estimates for each α_i and m_i, where m_i is equal to the smallest observation within each subsample. The smallest subset

[4] A rationale for using the test statistics instead of standard p-values available in software packages is provided in section 3.2.

includes only the data points within the 100th percentile and the largest contains all observations between the 100th and 71st percentile. Maximum likelihood (ML) is our preferred estimation method since it is a well-established result that ML estimators are superior to other approaches if the data of interest follows a power law (Clauset et al., 2009, Greene, 2012). The ML estimator in our case is equivalent to the so-called Hill estimator[5].

Next we perform a goodness of fit test for each of the 30 subsamples per implicate by computing the CvM statistic (ts_i), which increases with the difference between the observed sample and the estimated Pareto distribution. Thus, low t-statistics point to relatively good fits. Figure 4 plots these statistics for the upper 30 percentile-subsamples[6] and the corresponding $\hat{\alpha}$ across all five implicates. As is evident form Figure 4, the test-statistics vary considerably between implicates and even more do the $\hat{\alpha}$'s, especially for the first few sub-samples. The huge variation in the top of the distribution is very likely the effect of small sample sizes, since each percentile only contains roughly 24 observations.

Both, the estimations of $\hat{\alpha}_i$ as well as the computations of ts_i were performed only using the HFCS data without the corresponding weights. This operation is the only one where we ignore weights and we believe there are good reasons for doing so: First, estimation results hardly change when weights are taken into account. The difference in the average $\hat{\alpha}$ across implicates is equal to -0.00317, demonstrating not only a minor relative effect, but also leading to slightly more conservative estimates (i.e. wealth is distributed more evenly). Second, the construction of sampling weights by the Austrian central bank involves a battery of unknown regression models and assumptions about the determinants of response probabilities. By only using weights for linking the sample to the underlying population, which is required for the definition of wealth percentiles for example, we strongly limit the influence of those unknown implicit assumptions. Third, if using weighted data to carry out the CvM tests, one needs to handle the variation involved in the construction of weights by using re-sampling weights. While the question how to combine the CvM test with those re-sampling weights is far from trivial, such a procedure would also greatly increase complexity and computation time involved since total operations (estimation of α_i and computation of ts_i for each subsample) would increase from 300 to more than 150,000 if the full set of re-sampling weights provided by the Austrian central bank, is used.

---

[5] Dropping implicate indexes and starting from the pdf $p(x) = \alpha \frac{m^\alpha}{x^{\alpha+1}}$ one obtains the log-likelihood function $L(\alpha, x) = \sum_j^n \ln \alpha + \alpha \ln m - (\alpha + 1) \ln x_j$ and maximization yields $\hat{\alpha} = \frac{n}{\sum \ln {x_j}/_{m}}$. For derivations with slightly different notation see Hill (1975) or Appendix B in Clauset et al. (2009).

[6] The results do not change if one considers the upper 50 percentiles instead of the upper 30. Fitting Pareto distributions to subsamples including data below the 50th percentile is not supported theoretically nor does it yield satisfactory fits as the CvM statistics rapidly increase.

<PLACE FIGURE 4 HERE>

Our focus was on determining $\hat{m}_i$ such that on the one hand it is not sensitive to sample size problems (as seems to be the case in the highest percentiles), and on the other hand it does not rely on ad hoc assumptions alone (Bach and Beznoska, 2011, Bach et al., 2014) or a merely visual inspection of well-known log-log graphs (Cowell 2011). In this context Clauset et al. (2009) illustrate the effect of using unreliable cut-off values when estimating α. Choosing $\hat{m}_i$ below m_i leads to the inclusion of non-Pareto distributed data and thus to biased estimates. Conversely, choosing $\hat{m}_i$ above m_i ignores potentially useful information and, thus, lowers the statistical precision of the estimates. While the method provided by Clauset et al. (2009) proves a suitable guide in achieving these aims (see also Ceresetti *et al.*, 2010, Chui and Fyles, 2010, Rubinov et al., 2011), the presence of five different and autonomous implicates leads to an additional difficulty, namely how to synchronously identify a good fit across all five implicates. In the application of Wald's well known maximin model (Wald, 1945) we found a satisfying answer to the latter concern: The maximin model posits that in the face of different alternatives with uncertain consequences, one should rank those alternatives on the basis of their worst-case consequences (worst fit in our case) and choose that option were the worst-case is at least as good as all other alternatives. Thus, it introduces a certain degree of conservatism to the chosen estimation results by focusing on the relatively worst fits and nullifying the impact of single exceptionally well-fitting subsamples across all implicates. In detail, we first chose the maximal test-statistics (i.e. the worst fit) across implicates for each subsample and then identify the minimum of these test statistics (i.e. the best fit) across all samples. By applying this procedure we find that the threshold value of the 78th percentile proves to be the most suitable candidate for providing a statistically reliable estimate for m_i. Additionally the interval around the result is characterized by small test statistics as well as stable alpha parameters across implicates, which increases our confidence in the results presented in Table 2:

<PLACE TABLE 2 HERE>

So far we have established what we deem to be a non-arbitrary technique of fitting a Pareto distribution to the upper tail of the Austrian HFCS sample. However, even if this upper tail does not follow a Pareto distribution, the according parameters could still be estimated without noticing the mistake. Therefore we rigorously test the hypothesis that our data is actually drawn from a Pareto distribution prior to using these estimates for correcting the HFCS sample.

3.2 Testing the Pareto-Hypothesis

In the previous subsection we elaborated on how to find reliable distribution parameters. In turn, the obvious follow up question is whether the hypothesis that the estimated distributions truly represent the data is indeed correct. On first sight this might seem superfluous, since the p-values based on the Cramer von Mises tests already applied would provide an immediate answer to the question whether the data within a given subsample is statistically different from the estimated distribution or not. However, those standard p-values are derived under the assumption that the distribution against which the data is tested is perfectly known, which is typically not the case for survey data due to the variations inherent in the random sampling processes of any survey data collection. As a result, standard p-values are not suitable for clarifying this issue. This is why we again follow Clauset et al. (2009), who suggest comparing the goodness of fit of the original data and its estimated distribution with the goodness of fit of newly created data vectors based on the original data as well as the estimated distribution. While these new data vectors are created by means of a bootstrap – that is repeated random drawing from the estimated distribution (above $\widehat{m}_i$) and the original data (below $\widehat{m}_i$) – the general idea is to test the goodness of fit of the original estimation against the goodness of fit of an estimation based on these newly generated data vectors, where the data for top-wealth households (i.e. all households above $\widehat{m}_i$) is already known to truly follow a Pareto distribution. Thus, this procedure serves as a means for coping with the uncertainty arising from random sampling testing our hypothesis.

Following this strategy, we create $B = 10.000$ synthetic datasets (X_{ib}) for each implicate by drawing a number x_{ij} with probability t_i/n from the previously estimated distribution with the parameters $\hat{\alpha}_i$ and $\widehat{m}_i$, where $n = 2380$ is the number of total observations and t_i is the number of observations above $\widehat{m}_i$. With probability $1 - t_i/n$ we pick a random element x_{ij} from the original dataset below $\widehat{m}_i$. Repeating this process from $j = 1 \dots . n$ yields a synthetic dataset with 2,380 observations where all elements above $\widehat{m}_i$ are drawn from the originally estimated distribution. For each implicate we use these 10,000 data sets to compute an artificial p-value (p_i) for the hypothesis test that the data follows a Pareto distribution with $\hat{\alpha}_i$ and $\widehat{m}_i$ more closely than the synthetic datasets follow their estimated distributions. Since we already know that the synthetic datasets truly follow a Pareto distribution above $\widehat{m}_i$ and any differences are the result of random variations, this implies the Null-hypothesis that the data truly follows a Pareto distribution. In order to obtain this p-value with respect to this hypothesis, we repeat all the steps from section 3.1. So for each of these synthetic datasets a Pareto distribution with cutoff-point $\widehat{m}_{ib}$ and shape parameter $\hat{\alpha}_{ib}$ is estimated as described above and the corresponding CvM test statistic ts_{ib} is computed. This test-statistic is in turn compared to the original test-statistic ts_i according to the condition that

$$ts_{ib} \geq ts_i \qquad (2)$$

In any instance where the above condition (2) holds the original data is actually more similar to its estimated distribution than the new data vector. Obtaining p for our hypothesis now implies to count all instances where condition (2) holds (denoted by c_i) and divide this sum by the total number of synthetic datasets B and average over all implicates, i.e.

$$p = \frac{1}{5}\sum_{i=1}^{5} {c_i}/{B} \qquad (3)$$

The interpretation of this artificial p-value is pretty standard, namely that below the 10% level the Null-hypothesis is rejected. The main rationale behind this is, that synthetic test statistics larger than the original test statistic indicate that the difference between the original data and its estimated distribution is due to random variation, since the synthetic data truly follow a Pareto distribution above $\hat{m}_{ib}$. If enough synthetic test statistics are larger than ts_i, one cannot reject the Pareto hypothesis any more. We simply follow the common practice and use 10% as our threshold below which we reject the hypothesis that the data follows a Pareto distribution. However, the results for single implicates are partially idiosyncratic and far from consistent across all implicates.

<INSERT TABLE 3 HERE>

One can see from Table 3 that the Pareto distribution is a plausible model for implicates 2, 4 and 5 and that it is strongly rejected for implicate 1 and weakly for implicate 3. On the average however the hypothesis still holds. We focus on this latter result since the variability expressed by the single implicates is due to imputing missing data based on a series of different statistical models. Thus, only an average across those implicates seems to be a justifiable criterion, since the different implicates have to be interpreted conjointly to appropriately consider the variability between implicates.

3.3 Correcting the Data

After identifying $\hat{\alpha}_i$ and $\hat{m}_i$ we use this information to correct the missing observations of rich people in the upper tail of the HFCS sample. In doing so, we first remove all observations exceeding 4 million Euros in net wealth from the original data set. We choose that threshold because the frequency of observations starts to markedly decline around this level of net wealth. Thereby patterns however slightly differ across implicates (see Figure 5). All these observations are part of the 100th percentile (except for the second implicate, where the 100th percentile starts at 4.6 million Euros) and represent between 11,374 and 44,081 households depicted by 8 to 30 observations depending on the specific implicate one considers. This treatment implies that we assume that the alleged non-observation bias affects this group of households and instead suggest to rely on the estimated Pareto distribution for observations above 4 million Euros in net wealth.

<INSERT FIGURE 5 HERE>

To determine how many households should be added to the sample based on the estimated Pareto distribution we look at the number of households (HH_i) with net wealth holdings above $\hat{m}_i$ and below $\mu \equiv 4\, million\, Euro$ according to the HFCS data set. HH_i varies between 785,924 for $i = 2$ and 817,418 for $i = 1$. By using the properties of the underlying probability distribution function we compute the number of households above 4 million Euros (H_i) by:

$$H_i = HH_i \frac{1-P_i(\mu)}{P_i(\mu)} \qquad (4)$$

H_i varies between 22,982 for $i = 5$ and 40,251 for $i = 2$. This approach ensures that the correction for rich households only depends on high quality observations from the HFCS data. Given H_i, one can derive the wealth x_i for each household above μ by exploiting the fact that

$$1 - P_i(x_i) = \Pr(X_i > x_i) = \left(\frac{m_i}{x_i}\right)^{\alpha_i} \equiv \frac{H_{x_i}}{HH_i + H_i}.$$

Rearranging terms gives

$$x_i = m_i \left(\frac{HH_i + H_i}{H_{x_i}}\right)^{1/\alpha_i}, \qquad (5)$$

where H_{x_i} is the number of households reporting a net wealth of at least x_i or, put differently, the rank of a given household. As a result one obtains new observations for net wealth above μ for each implicate. It is important to note, however, that we limited net assets by 1 billion Euros. Specifically any observation above that value was set equal to 1 billion Euros. The motivation for this truncation of the newly generated sample is our preference for conservative estimates. In the next section we investigate the implications of this data correction and the change of major descriptive statistics.

Before that it is important to adapt the sample weights since the number of households added and the number of households removed from the sample differ. The net-change, which is the difference between the number of households above 4 million Euros according to the original HFCS sample and HH_i varies between +16,280 ($i = 1$) and – 3,828 ($i = 2$). In the latter case, the original HFCS data set reported a higher number of households above μ compared to the estimated Pareto distribution. Thus in this case the weights of the remaining observations below the cut-off point need not to be reduced but increased. In either situation the alteration of sampling weights is done proportionally to the total number of households less the observations above $\hat{m}_i$. For example the net change in the first implicate (+16,281) in relation to the number of households below $\hat{m}_i$ equals 0.55%. As a result the weights for observations below the cut-off are reduced by 0.55%. On average (across the implicates) the weights are reduced by 0.21%. The next section contrasts the results of the data correction to the original sample.

4. The Impact of the Missing Rich: Distributional Statistics

This section summarizes the distributional statistics of the corrected data and compares the presented results to those obtained from the original HFCS-data. Let us first take a look at the impact on the overall structure of the sample: Figure 6 shows, analogously to Figure 1, how the population is distributed among different segments of wealth, where the relative changes between original and corrected data have been highlighted. In this representation we only observe minor changes: The share of the population possessing net wealth greater than 500,000 Euros slightly increases due to the increase in very rich households. Correspondingly the remaining shares decrease.

<INSERT FIGURE 6 HERE>

While it is immediately obvious from Figure 6 that the number of households with net wealth above 500,000 Euros slightly changes, it remains unclear how the estimated wealth changed. In this context Table 4 – basically a replication of Table 1 – is more revealing: it indicates that the total wealth of the richest percentile grows by more than 100%, namely from 237 billion Euros to 497 billion. From this it follows that total wealth also increases significantly from roughly 1000 billion Euros to about 1270 billion Euros. Thus, this analysis indicates an increase of 27% in the estimate for total wealth due to the correction for the alleged non-observation bias, which is concentrated in the top wealth classes. In this context the increase in the share of percentiles 96-99 is due to the correction of the number of households richer than 4 million Euros, which in turn implies that all households drop in ranks. A full list of all percentiles based on the corrected data can be obtained from the authors.

<INSERT TABLE 4 HERE>

Correspondingly, the share of wealth held by selected population groups changes significantly, with the most remarkable change in the share of the richest percentile, which increases from 22.9% to 38.2%. The share of the poorest 50% of the wealth distribution, on the other hand, decreases from 2.8% to 2.2%.

These changes in estimates for total wealth and its distribution are mirrored by the change of central indicators like the Gini-coefficient, which rises from 0.762 to 0.811 if its calculation is based on the corrected data. To further illustrate the changes in the estimated wealth distribution, Figure 7 compares a Lorenz curve for the original data to a Lorenz curve for the corrected data. As can already be inferred from Table 2, we observe a significant shift in the Lorenz curve due to the correction of the data as exemplified in the foregoing section.

<INSERT FIGURE 7 HERE>

In sum our estimations suggest that the size of wealth omitted due to the non-observation bias inherent in survey designs related to private wealth holdings is indeed significant. The estimate for total wealth changed by roughly a quarter although the correction of the data affected less than 1% of the underlying population. Consequently, one has to conclude that the implications of the non-observation bias are far from trivial and can hardly be ignored when dealing with top-wealth data.

However, one possible objection to this conclusion is to question the robustness of the estimated distribution used for the correction of the initial data, since the reliability of the results presented here strongly depend on the validity of our distributional assumptions. In what follows we offer a series of such robustness checks, thereby scrutinizing the adequacy and plausibility of our methodological setup as well as the resulting estimations.

5. Robustness Checks

Given the sharp increase in net wealth holdings in the 100th percentile due to the data correction we suggest in section 3, the reader may be interested in the uncertainty of our estimates or be sceptical about our results anyway since they depend entirely on the statistical Pareto model. In this section we try to validate the robustness of our results in two ways: first we compare some specific implications of our estimated Pareto distribution with real data where it is available. Second, we provide upper and lower bounds of estimated net wealth with respect to data variability across subsamples by means of a bootstrap. In pursuing the first route we use the number of billionaires as provided by the Austrian equivalent of the (international) Forbes-list on top wealth households as an independent data source accessing the plausibility of our estimation with regard to the tail of the distribution. Additionally, we try to estimate the total wealth of all households between $\hat{m}_i$ and μ and compare the results to the estimates provided by the original HFCS-data.

5.1 Validating the Number of Billionaires

When correcting the original HFCS-data for the non-observation bias we assumed maximal wealth to be equal to one billion dollar. Practically, this implied that net wealth of all households, exceeding 1 billion Euros according to the estimated Pareto distributions where set to one billion. Although this restriction might lead to a significant underestimation of total wealth (assets of Austrian billionaires amount to roughly 70 billion Euros according to

Austria's counterpart to the Forbes-list), we nonetheless imposed it for reasons of statistical conservatism. However, even though we do not want our estimates to rely on the upper end of the distribution's tail, we can still use it in validating our results. In doing so we compared the estimated number of households with net wealth greater than one billion as implied by our estimated distributions with the available media information. The latter varies considerably between years (and it is unclear whether these variations are due to actual changes in wealth or just to changes in journalists' informational status), from 19 in 2011 (the year of the HFCS survey) to 30 in 2013. Our own calculations point to 31 billionaires and thus are very well in line with the figures reported by journalists.

5.2 Validating the Amount of Wealth Between m_i and μ

In addition to the robustness check performed in the previous subsection we applied the same logic to the less risky question of how well our estimated results match the original data in those segments of the distribution which were not affected by our correction procedure. Specifically, we compared the estimate for the total wealth of all households between $\widehat{m}_i$ and μ as implied by the Pareto distribution with the estimates obtained from the original HFCS-data. Here our results point in a similar direction as in Section 5.1: We only find a minor difference of -33.4 billion Euros between the estimate obtained from the HFCS-data and that implied by the corrected data (565.8 billion Euros). The suggested estimation strategy, thus, seems to be a relatively robust one.

5.3 Assessing the Sampling Variation

Due to the complex survey design of the HFCS, which involves stratified cluster sampling as well as multiple imputation to correct item-non-response, one is confronted with serious complications when trying to compute confidence intervals reflecting the uncertainty of the estimation process. While current literature offers procedures to compute confidence intervals with either multiply imputed data (Rubin, 1987) or data from complex surveys (Rao and Wu, 1998, Rao et al., 1992, Kolenikov, 2010), there is, according to our knowledge, no contribution which shows how to construct appropriate confidence intervals when multiple imputation as well as a complex survey design are used in the data collection process. Therefore, we implemented an approach to validate the robustness of our estimates with regard to sampling variation and suggest focusing only on the uncertainty arising from the variability of the original data. In doing so, we apply a bootstrap in order to test the robustness of our results with respect to random resampling. Even though we cannot express the uncertainty of the estimation process itself this way, we are still able to demonstrate the robustness of our results due to potential outliers and irregularities within certain subsets of the original HFCS sample. Therefore, it is important to bear in mind that the bounds reported below do not serve as direct substitutes for traditional confidence intervals.

The bootstrap procedure for computing an upper and lower bound of $\hat{\alpha}_i$ involves the construction of $U = 1000$ random samples consisting of $n_U = {}^2\!/_3 * n \approx 1587$ observations randomly picked from the original HFCS data set for each implicate. Then we re-estimate $\hat{\alpha}_i$ for each random sample. After ordering them in ascending order, the 26th estimate of α_i is identified as the lower and the 975th as the upper bound of $\hat{\alpha}_i$. By repeating the data correction procedure described in section 3.3 for the upper and lower bound $\hat{\alpha}_i$'s, we obtain new estimates for the distribution of net wealth. The results are reported in Table 5 (already averaged across the implicates). As one can see, the upper bound of net wealth within the 100th percentile deviates approximately by 140 billion Euros from the point estimate while the lower bound deviates by 111 billion, indicating that the 237 billion Euros reported in the original HFCS data are very likely to be downward biased.

<INSERT TABLE 5 HERE>

6. Conclusion

In this paper we tried to correct for the underrepresentation of the wealthiest households in the Austrian HFCS data by fitting a Pareto distribution to the upper fifth of the wealth distribution. For the exact determination of the lower bound of the distribution we used the Cramer-von-Mises test instead of graphical evidence or ad hoc assumptions. Finally, all households within the sample with reported net wealth above 4 million Euros were eliminated from the dataset and replaced with the households from the estimated distribution.

It can be seen that this adjustment exhibits a significant influence on the results. The estimated aggregate wealth increase from about 1000 billion Euros to 1277 billion, where the increase is mainly due to the increase of wealth within the highest percentile (wealth within this percentile increases by 110%). Amongst other things it follows that the richest 10% of Austrian households possess 69.3% of total net wealth instead of the 61% that follow from the original HFCS data. The change in the share of the richest percentile is even more remarkable: it increases from 22.9% (HFCS) to 38.2%.

Finally we compare our results to a list of Austrian billionaires published by a national news magazine and additionally assess the fit of the estimated distribution to the original data for those parts of the distribution which are not corrected. These exercises reveal that our non-arbitrary approach of fitting a Pareto distribution is very well in line with non-sample evidence and also closely fits the data. Furthermore the computation of upper and lower bound wealth estimates indicates that the underrepresentation of wealthy individuals that we detect especially within the 100th percentile cannot be explained by potential irregularities and outliers in the sample and thus seems to be a robust finding.

References

Albacete, N., Lindner, P., Wagner, K. and S. Zottel, “Household Finance and Consumption Survey des Eurosystems 2010: Methodische Grundlagen für Österreich,“ *Geldpolitik & Wirtschaft*, 3/2012-Addendum, 2012.

Alvaredo, F. and E. Saez, “Income and Wealth Concentration in Spain in a Historical and Fiscal Perspective,” *Journal of the European Economic Association*, 7, 1140–67, 2009.

Alvaredo, F., Atkinson, A. B., Piketty, T. and E. Saez, "The Top 1 Percent in International and Historical Perspective." *Journal of Economic Perspectives*, 27(3), 3-20, 2013.

Atkinson, A., Piketty, T. and E. Saez, “Top incomes in the long run of history,” *Journal of Economic Literature*, 49(1), 3–71, 2011.

Avery, R. B., Elliehausen, G. E. and A. B. Kennickell, “Measuring wealth with survey data: an evaluation oft the 1983 survey of consumer finances,” *Review of Income and Wealth*, 34(4), 339-369, 1986.

Bach, S. and M. Beznoska, “Aufkommens- und Verteilungswirkung einer Wiederbelebung der Vermögenssteuer,“ DIW: Politikberatung Kompakt 68, Berlin, 2011.

Bach, S., Beznoska, M. and V. Steiner, “A Wealth Tax on the Rich to Bring Down Public Debt? Revenue and Distributional Effects of a Capital Levy,“ *Fiscal Studies, forthcoming,* 2014.

Beer, C., Mooslechner, P., Schürz, M. and K. Wagner, “Das Geldvermögen privater Haushalte in Österreich: eine Analyse auf Basis von Mikrodaten,“ *Geldpolitik und Wirtschaft*, 2/2006, 101-119, 2006.

Borgerhoff Mulder, M., Bowles, S., Hertz, T., Bell, A., Beise, J., et al., “Intergenerational wealth transmission and the dynamics of inequality in small-scale societies,” *Science*, 326(5953), 682-688, 2009.

Ceresetti, D., Molinié, G. and J.-D. Creutin, “Scaling properties of heavy rainfall at short duration: A regional analysis,” *Water Resources Research*, 46, W09531, doi:10.1029/2009WR008603, 2010.

Clauset, C., Shalizi, C. R. and M. E. J. Newman, “Power Law Distributions in Empirical Data,” SIAM Review, 51(4), 661-703, 2009.

Chui, J. K. W. and T. M. Fyles, “Apperent fractal distribution of open durations in cyclodextrin-based ion channels,” *Chemical Communications*, 46(23), 4169-4171, 2010.

Cowell, F. A., “Inequality Among the Wealthy,” Centre for Analysis of Social Exclusion at the London School of Economics, 2011.

Dagum, C., “Wealth distribution models: analysis and applications,” *Statistica,* 66(3), 235-268, 2006.

Durán-Cabré, J. M. and A. Esteller-Moré, “Tax Data for Wealth Concentration Analysis: An Application to Spanish Wealth Tax,“ *Review of Income and Wealth*, 56(3), 620-631, 2010.

ECB (2013): The Eurosystem Household Finance and Consumption Survey – Methodological Report for the first wave. *Statistics Paper Series,* No. 1/April, 2013.

Feenberg, D. R. and J. M. Poterba. "Income Inequality and the Incomes of Very High-Income Taxpayers: Evidence from Tax Returns," *Tax Policy and the Economy*, 7, 145-177, 1993.

Greene, W., *Econometric Analysis.* Pearson, Essex, Seventh edition, 2012.

Guttmann, R. and D. Philon, “Consumer debt and financial fragility,” *International Review of Applied Economics*, 24(3), 269-283, 2010.

Hahn, F. R. and C. Magerl, “Vermögen in Österreich,“ WIFO Monatsbericht 1/2006, 2006.

Hill, B., A Simple General Approach to Inference About The Tail of A Distribution. *The Annals of Statistics*, 3(5), 1163-1174, 1975.

Hoeller, P., Joumard, I., Bloch, D. and M. Pisu, “Less Income Inequality and More Growth - Are They Compatible?: Part1: Mapping Income Inequality Across the OECD,” OECD Working Paper No. 924, 2012.

Kolenikov, S., “Resampling variance estimation for complex survey data,” *The Stata Journal*, 10(2), 165-199, 2010.

Kopczuk, W. and E. Saez, “Top Wealth Shares in the United States, 1916-2000: Evidence from Estate Tax Returns,” NBER Working Paper 10399, 2004.

Piketty, T., “Income Inequality in France, 1901-1998,” *Journal of Political Economy*, 111(5), 1004-1042, 2003.

Piketty, T. and E. Saez, “Income inequality in the United States, 1913-1998,” *Quarterly Journal of Economics*, 118(1), 1-41, 2003.

Rao, J. N. K. and Wu, C. F. J., “Resampling Inference With Complex Survey Data,” *Journal of the American Statistical Association,* 83(401), 231-241, 1988.

Rao, J. N. K., Wu, C. F. J. and Yue, K., “Some Recent Work on Resampling Methods for Complex Surveys,” *Survey Methodology,* 18(2), 209-217, 1992.

Rubin, D. B., *Multiple Imputation for Nonresponse in Surveys,* J. Wiley & Sons, New York, 1987.

Rubinov, M., Sporns, O. Thivierge, J. P. and M. Breakspear, “Neurobiologically realistic determinants of self-organized criticality in networks of spiking neurons,” *PLoS computational biology*, 7(6), doi:10.1371/journal.pcbi.1002038, 2011.

Stiglitz, J. E., *The price of inequality: How today’s divided society endangers our future*, WW Norton & Company, New York and London, 2012.

Wald, A., “Statistical Decision Functions Which Minimize the Maximum Risk”, *Annals of Mathematics,* 46(2), 265-280, 1945.

Wilkinson, R. G. and K. E. Picket K. E., “The problems of relative deprivation: why some societies do better than others,” *Social Science & Medicine*, 65(9), 1965-1978, 2007.

Wilkinson, R. G. and K. E. Pickett, *The Spirit Level: Why More Equal Societies Almost Always Do Better,* London, Allen Lane, 2009.

Tables

Table 1: Austrian's richest 5% according to the HFCS-data

Percentile	Total net wealth (in billion Euros)	Average net wealth per household (in million Euros)	Share of total net wealth[1]	Culminated share of total net wealth
96	38.8	1.0	3.9%	47.6%
97	48.7	1.3	4.9%	
98	65.5	1.7	6.6%	
99	94.1	2.5	9.3%	
100	237.0	6.4	22.9%	
Total Sample	1,000	0.265	100%	100%

[1] Note that total wealth and the relative shares of wealth are not fully consistent, since each final estimate requires the application of Rubin's Rule (i.e. each estimate is based on averaging five single estimates obtained from the five different implicates, see: Rubin, 1987).

Table 2: Estimating the Pareto distribution for the Austrian wealth distribution: Results

Implicate	$\hat{\alpha}$	$\hat{m}_i$
#1	1.28808	281,242 Euros
#2	1.14815	287,809 Euros
#3	1.3332	289,811 Euros
#4	1.24881	293,161 Euros
#5	1.36649	288,422 Euros
Average	1.276946	288,089 Euros

Table 3: p-values across all implicates

Implicate	#1	#2	#3	#4	#5	Mean
p-value	0.019	0.2776	0.0983	0.5421	0.1781	**0.223**

Table 4: Austrian's richest 5% according to the corrected data.
Values using original HFCS data in parenthesis.

Percentile	Total net wealth (in billion Euros)	Average net wealth per household (in million Euros)	Share of total net wealth	Culminated share of total net wealth
96	40.8 (38.8)	1.1 (1.0)	3.2% (3.9%)	58.8% (47.6%)
97	50.3 (48.7)	1.3 (1.3)	4.0% (4.9%)	
98	66.7 (65.5)	1.8 (1.7)	5.2% (6.6%)	
99	101.2 (94.1)	2.7 (2.5)	7.9% (9.3%)	
100	497.3 (237)	13.4 (6.4)	38.2% (22.9%)	
Total Sample	1,277.7 (1,000)	0.339 (0.265)	100%	100%

Table 5: Upper and lower estimation bounds.

Lower Bound	Point Estimate	Upper Bound
Pareto's Alpha		
α = 1.35	α = 1.28	α = 1.21
wealth attributed to the richest percentile		
386 billion Euros	497 billion Euros	635 billion Euros

Figures

Figure 1: Distribution of households along different segments of ne wealth (see also Albacete *et al.* 2012, 41).

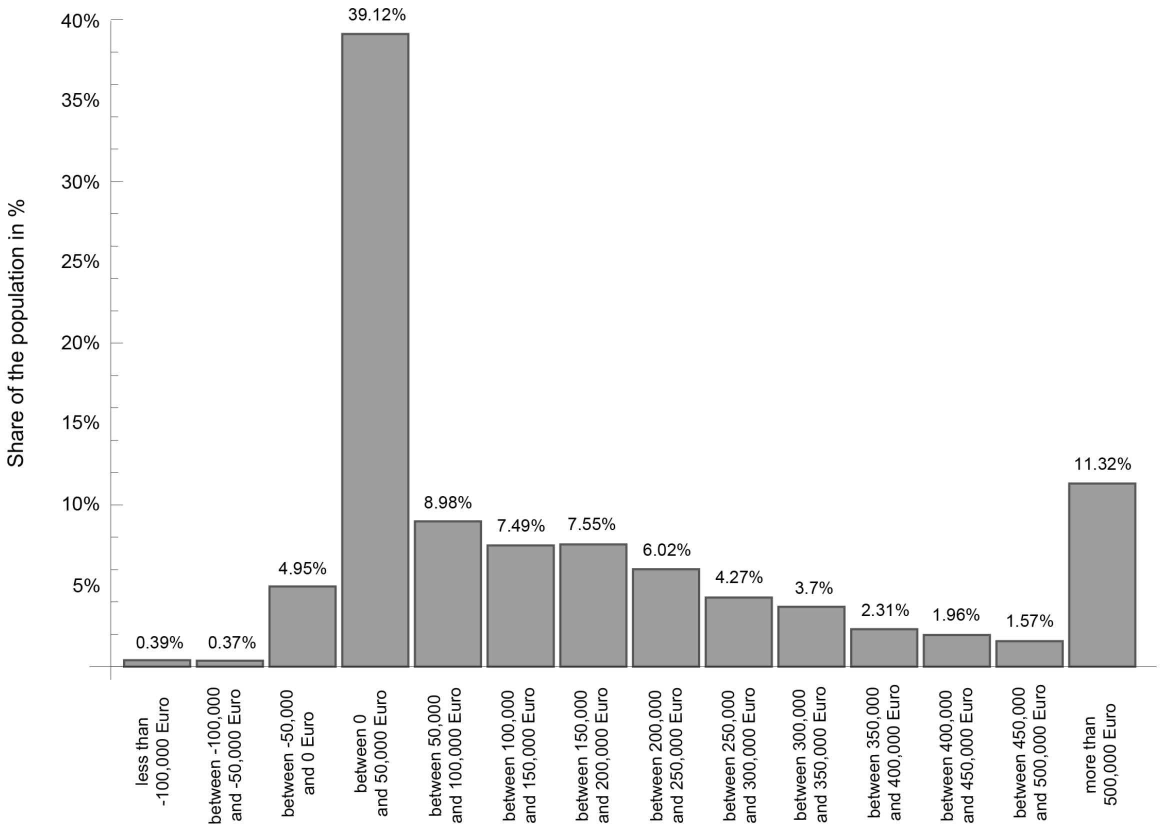

Figure 2: The cumulative distribution function of net wealth in Austria.

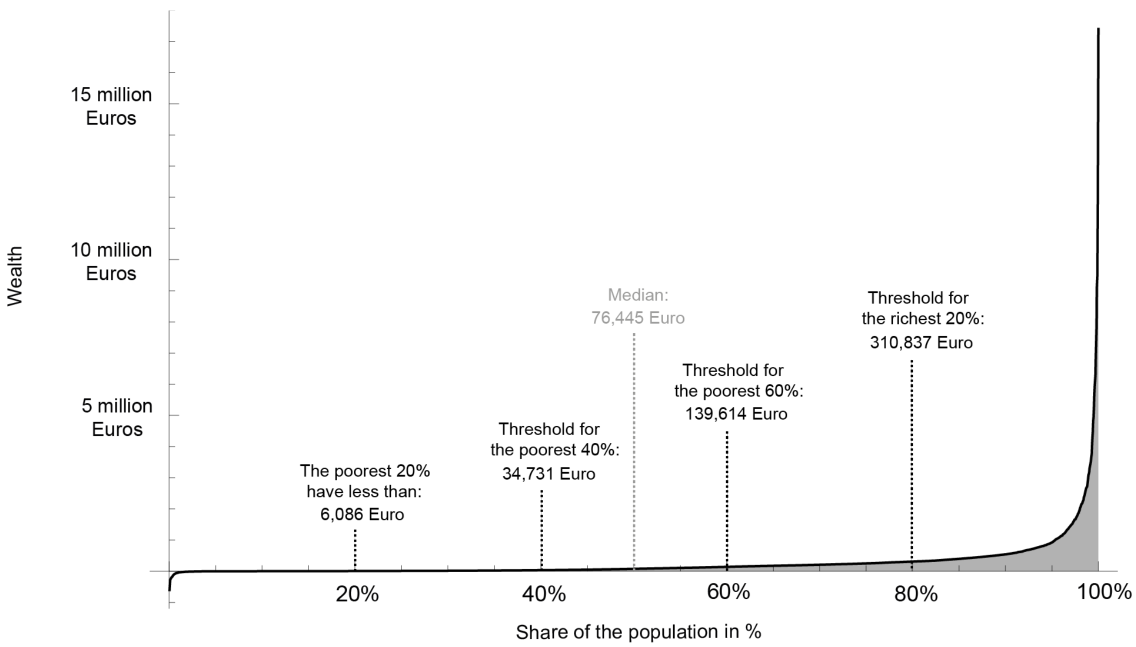

Figure 3: A sketch of the methodological strategy.

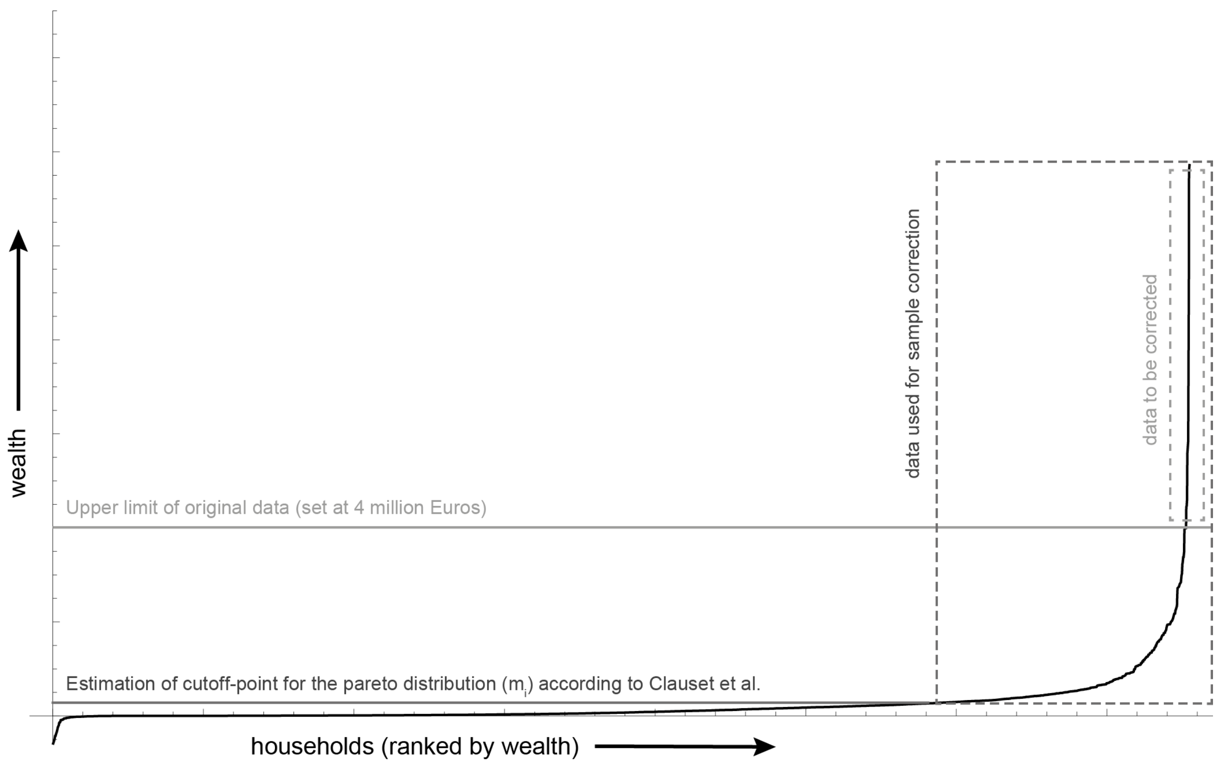

Figure 4: Estimates for Pareto's α and corresponding Cramer von Mises test statistics across all implicates.

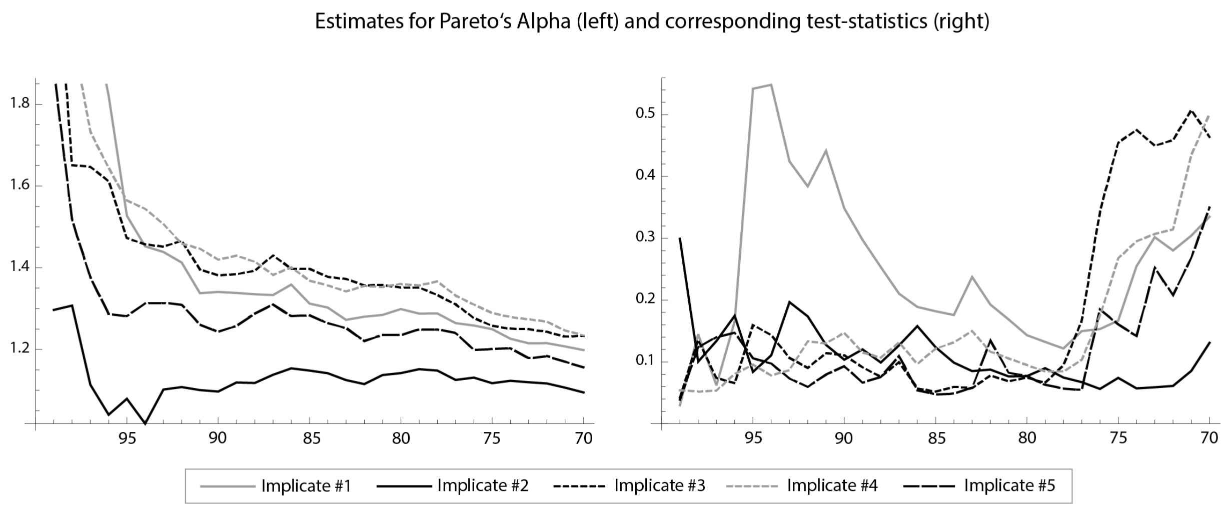

Figure 5: Plots of the richest 50 households in each implicate and the 4 million Euros cut-off.

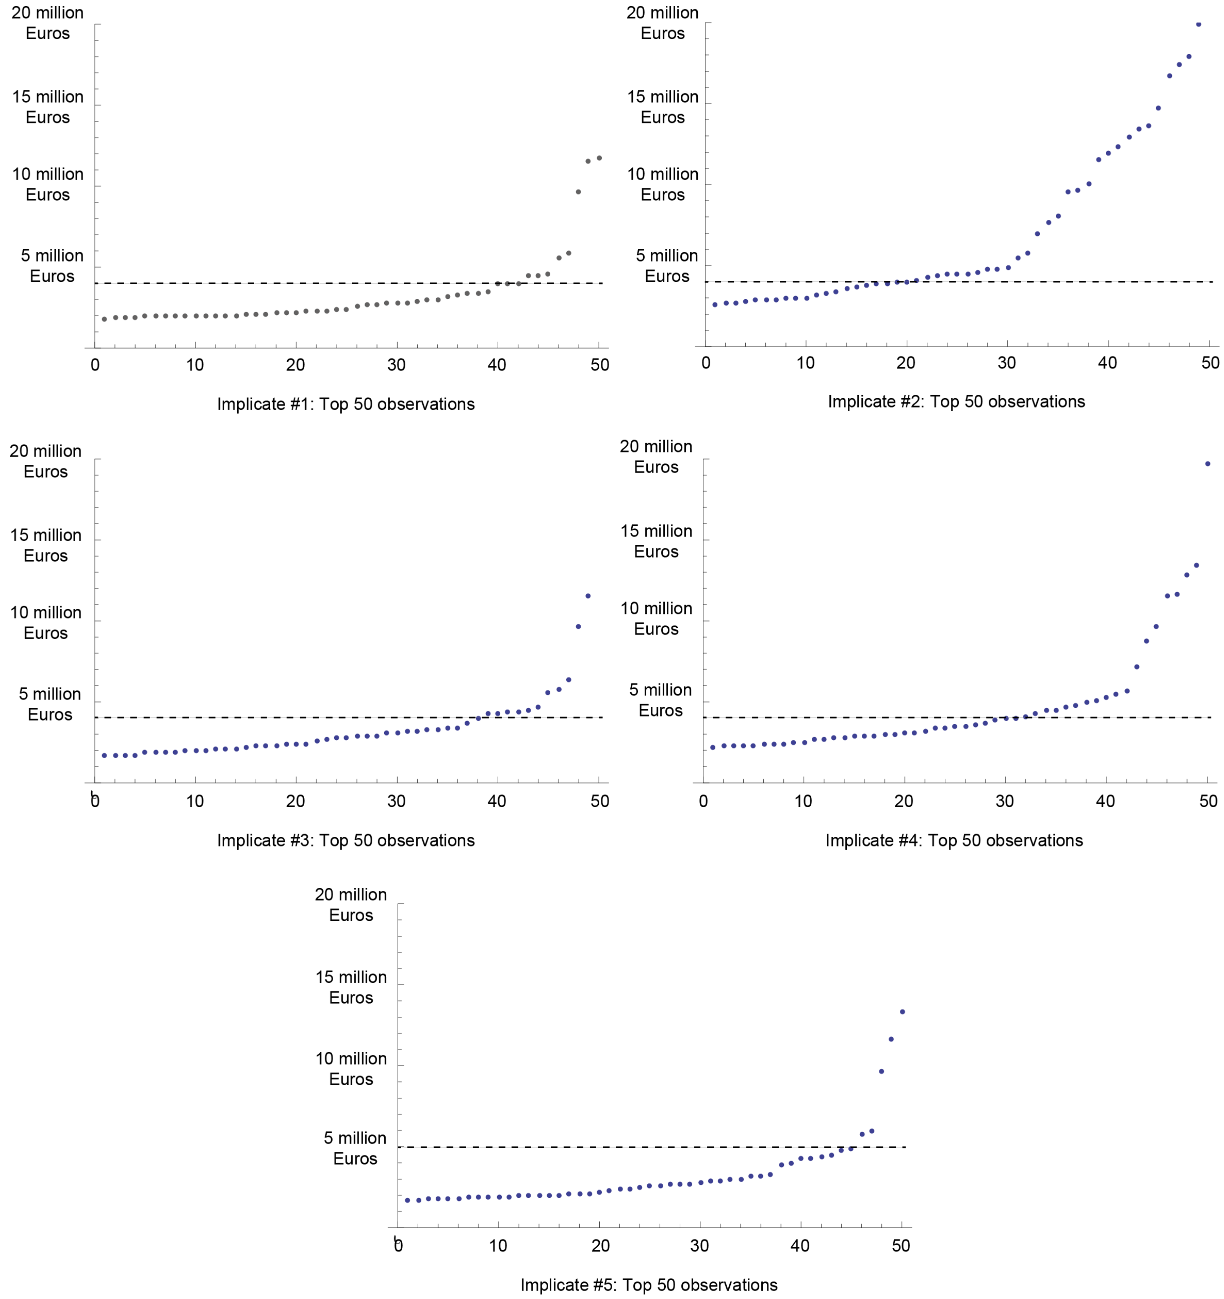

Figure 6: Distribution of households along different segments of net wealth based on the corrected data.

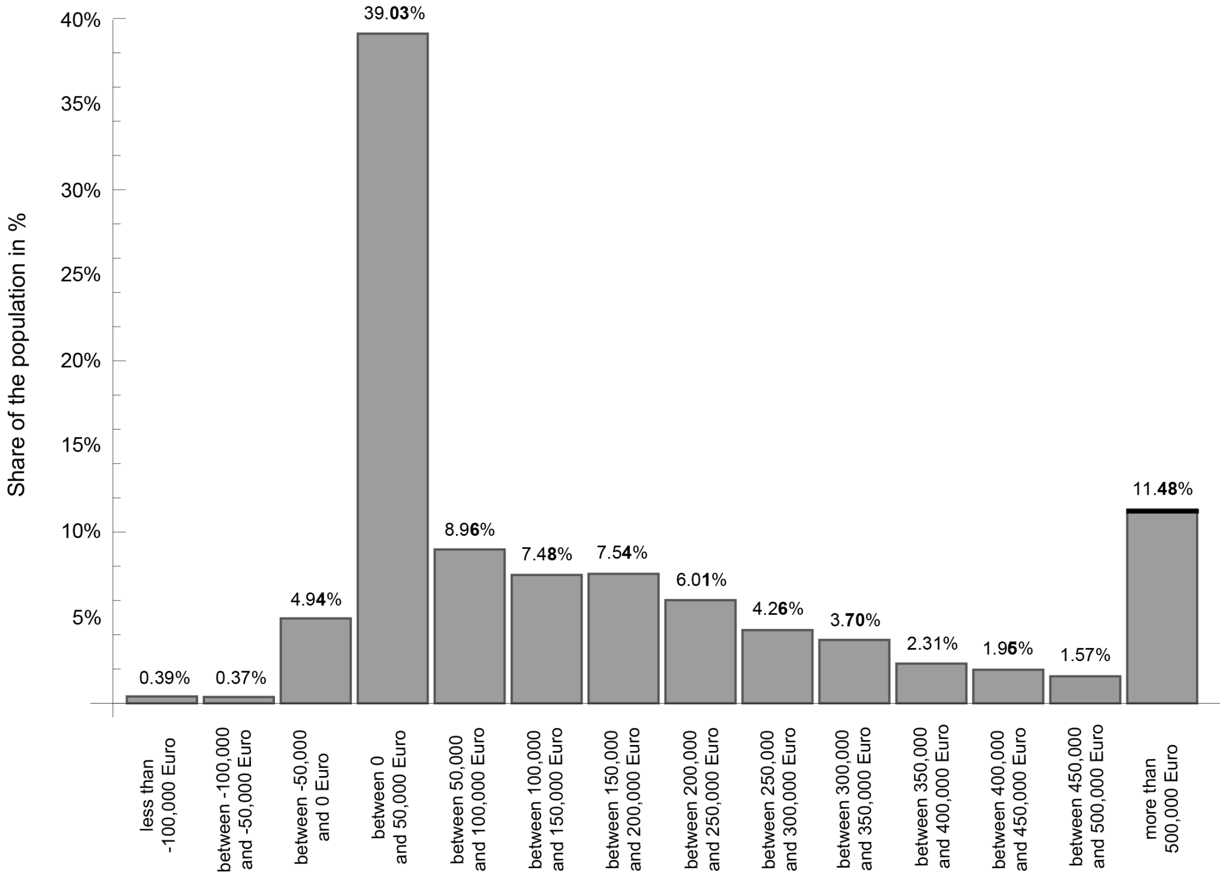

Figure 7: Lorenz curves for the original data (shaded area) as well as the corrected data.