the review of income and wealth

Review of Income and Wealth
Series 64, Number 2, June 2018
DOI: 10.1111/roiw.12279

HOW FAT IS THE TOP TAIL OF THE WEALTH DISTRIBUTION?

BY PHILIP VERMEULEN*

Directorate General Research, Monetary Policy Research Division, European Central Bank

Differential unit non-response in household wealth surveys biases estimates of top tail wealth shares downward. Using Monte Carlo evidence, I show that adding only a few extreme observations to wealth surveys is sufficient to remove the downward bias. Combining extreme wealth observations from Forbes World's billionaires with the Survey of Consumer Finances, the Wealth and Assets Survey, and the Household Finance and Consumption Survey, I provide new improved estimates of top tail wealth in the United States, the United Kingdom, and nine euro area countries. These new estimates indicate significantly higher top wealth shares than those calculated from the wealth surveys alone.

JEL Codes: D31

Keywords: non-response, wealth distribution, Survey of Consumer Finances, Wealth and Assets Survey, Household Finance and Consumption Survey

1. INTRODUCTION

Understanding the wealth distribution is important for a number of reasons. For instance, any analysis of taxation and redistribution policies crucially depends on the shape of the wealth distribution. As wealth is usually very concentrated at the top, measures such as the share of wealth held by the top 1 or 5 percent of households carry a broader importance as measures of wealth inequality. Despite the obvious importance, accurate measurement of the wealth distribution, and especially its upper tail, has proven to be very difficult.

Recently, a new survey, the Eurosystem Household Finance and Consumption Survey (HFCS), covering in its first wave 15 countries from the euro area, expands enormously the number of countries for which wealth distribution estimates can be made. In the years to come, this survey is likely to add substantially to the knowledge of the wealth distribution in Europe. Similarly, the UK Wealth and Assets Survey (WAS) is a relatively recent addition to the set of wealth

Note: I would like to thank the members of the HFCS network for many stimulating discussions. I also thank Jirka Slacalek, Dimitrios Christelis, Markus Grabka, Anthony Shorrocks, Thomas Piketty, Daniel Waldenström, three anonymous referees, Prasada Rao (Editor), participants at the Household Wealth data and Public Policy Conference, SEEK 2014, IEAA 2014, and EEA2014 for useful comments. Many thanks to Samuel Skoda for research assistance. This paper uses data from the Eurosystem Household Finance and Consumption Survey. Any remaining errors are solely mine. The views expressed in this paper only reflect those of the author. They do not necessarily reflect the views of the European Central Bank.

*Correspondence to: Philip Vermeulen, DG-Research, ECB, Sonnemannstrasse 20, 60314 Frankfurt am Main, Germany (philip.vermeulen@ecb.europa.eu).

surveys. A better understanding of these recent European surveys, and especially their measurement of top tail wealth, is our first concern. The U.S. Survey of Consumer finances (SCF), sponsored by the Board of Governors of the Federal Reserve System, is added to the analysis as it oversamples heavily at the top of the distribution and forms an interesting comparison.

The first contribution of this paper is to provide new estimates of the share of wealth held by the top 1 and 5 percent richest households in the U.S., the U.K., the Germany, France, Italy, Spain, the Netherlands, Belgium, Austria, Finland, and Portugal. It does so based on an analysis of household survey data combined with Forbes World's billionaires list. For the euro area countries, the analysis is therefore restricted to those which have individuals on this list.

Besides providing new estimates, this paper also makes a methodological contribution. Wealth estimates from surveys will (almost) always underestimate top tail wealth. The main reason causing this downward bias is the existence of *differential unit non-response*, the fact that richer households are less likely to take part in such surveys. When non-responding households have greater wealth in some systematic but unobserved way, wealth estimates will be biased downward, and estimates of tail wealth will be particularly affected. On the methodological side, I provide new insights into the importance of differential unit non-response of the wealthy in the SCF, WAS, and HFCS for tail wealth measurement.

Finally, I propose a method to alleviate the effect of differential unit non-response on the estimates of tail wealth. The method consists in replacing the tail observations with a Pareto distribution that is estimated on a combined sample of survey tail observations and extreme wealth observations obtained from another data source. I show, using Monte Carlo simulation, that this method, under the assumption that tail wealth is Pareto distributed, is able to recover unbiased estimates of tail wealth, even if surveys suffer from differential unit non-response. I apply this method, add the Forbes World's billionaires list to the survey data, and provide new tail wealth estimates.

This paper belongs to a literature with a long tradition of wealth distribution estimation. Essentially, researchers have come up with widely different methods to estimate top tail wealth, mainly as a function of the data at hand.[1] Methods can broadly be divided into five groups. First, in a few countries with a wealth tax, researchers have been able to use official wealth (tax) records. This has been the case, for example, in Roine and Waldenström (2009) for Sweden, Alvaredo and Saez (2009) for Spain, and Dell *et al*. (2007) for Switzerland. Second, estate tax records, which give information on taxable inheritances, can be used through the estate multiplier method to estimate wealth holdings of the living. This is an old and large literature. Some of the more recent findings are Kopczuk and Saez (2004) for the U.S. and Piketty *et al*. (2006) for France. Third, capital income information from tax records can be used to construct wealth estimates assuming certain rates of return on wealth. See, for instance, the recent study by Saez and Zucman (2014) for the U.S. Fourth, household wealth surveys that are representative of the population can provide direct estimates of the wealth distribution. And, finally, lists of wealthy individuals provided in the media or other sources

[1]Roine and Waldenström (2014) provide a recent overview of the literature.

can be used to estimate top tail wealth. This paper combines household wealth surveys with such data.

Using household wealth surveys to estimate wealth distributions is likely to remain important in the future. First, only a few countries have a wealth tax and many have large exemptions on inheritance tax so that administrative records do not exist or are limited in scope. Second, where tax or other records exist, they might not be made available to researchers for confidentiality reasons. Gaining a deeper understanding of the limitations of survey data and proposing methods to alleviate some problems has been the main motivation for this study.

The remainder of the paper is structured as follows. Section 2 describes the data used, the SCF, the WAS, the HFCS, and Forbes World's billionaires. It also contains a discussion of the issue of oversampling and non-response. Section 3 discusses how the Pareto distribution can be estimated using survey data. The section draws on the power law literature. My contribution to this literature is a discussion of how to deal with complex survey data, where the weights of the sample points are important. It also contains a Monte Carlo study, illustrating that information from rich lists can improve Pareto estimates in the presence of differential unit non-response. Section 4 provides new estimates of the share of wealth held by the top 1 and 5 percent households. Section 5 concludes.

2. The Data

2.1. *The U.S. SCF, the U.K. WAS, and the Eurosystem HFCS*

This paper combines the 2010 wave of the U.S. SCF, the second wave of the U.K. WAS, the first wave of the HFCS, and the Forbes World's billionaires list (for the years 2009–11) to estimate wealth at the upper tail of the distribution. The SCF is a triennial survey of U.S. household wealth, sponsored by the Board of Governors of the Federal Reserve System. It provides the most comprehensive source of wealth information of U.S. households, collecting detailed data on assets and debts of around 6,000 households. The HFCS provides detailed information on household assets and debts of individual households in 15 euro area countries. In total, there are more than 62,000 households in the dataset. The collection period of the data differs across countries and ranges from 2008 to 2010. For most countries, the wealth recorded in the survey is what it is at the time of the interview. The only exceptions are Finland and the Netherlands, where wealth is provided for December 31, 2009, and Italy, for December 31, 2010. The WAS is a longitudinal sample survey of households in the U.K. Wave 2 of the survey collected household wealth data over a period from July 2008 to June 2010. Around 20,000 households responded in the second wave of the WAS survey.

I use the HFCS data for Germany (2010), France (2010), Italy (2010), Spain (2008), the Netherlands (2009), Belgium (2010), Portugal (2010), Austria (2010), and Finland (2009) (in brackets are the reference years for the wealth). I drop Greece, Cyprus, Malta, Luxembourg, Slovakia, and Slovenia from the dataset, as these countries had no Forbes billionaires at the time of the survey. The concept of wealth that is used is that of "household disposable net wealth." As discussed in Wolff (1990), that is a conventional measure of all assets that have a current market value less liabilities.[2]

The SCF, the WAS, and the HFCS survey samples are purposefully designed to be representative of the household population of the respective countries. The survey samples are obtained through probability sampling, using a complex survey design. Complex survey designs imply a combination of stratification, clustering, and weighting of the data. By design, sample inclusion probabilities vary across households. Sample weights are provided and each sample weight signifies the number of households in the population that the sample point represents. The total sum of weights for each country is equal to the total number of households in the population.

The SCF and the HFCS do multiple imputation to deal with missing observations. For each missing observation, five imputations are made. This implies that the data are provided as five replicates of the dataset, called "implicates" in the parlance of the SCF (Kennickell, 1998). For variance estimation, the survey provides bootstrap weights. In the estimation results below, these bootstrap weights are used to provide standard errors around the mean estimates. The WAS uses single imputation and does not provide bootstrap weights for variance estimation. The WAS results therefore do not allow construction of standard errors for the estimates relating to the U.K. wealth distribution.

A more detailed description of the SCF, WAS, and HFCS methodologies can be found in Kennickell (2000), ONS (2012), and HFCS (2013). For comparison purposes, the SCF data are converted into euro using the dollar/euro exchange rate of 1.3572 on February 12, 2010 (which coincides with the date of the Forbes list); the WAS data are converted into euro using the pound/euro exchange rate of 0.867183 (which is the average over the data collection period from July 2008 to June 2010).

2.2. *Oversampling the Wealthy and Differential Unit Non-response*

Wealth is heavily concentrated at the top tail of the distribution. To increase efficiency, wealth surveys usually attempt to oversample the wealthy. The word "attempt" is used purposefully here, as success is not guaranteed. In practice, extraneous information such as tax registers or other information are used to construct a sampling frame that allows oversampling of a part of the population thought, on average, to be wealthier.

Efficiency is not the only challenge (one can always increase the sample size); another big challenge in wealth estimation at the top is the existence of differential unit non-response. There is a strong presumption among survey specialists

[2]The list of assets that are included are owner-occupied housing, other real estate, vehicles, valuables, and self-employment businesses, non-self employment private businesses, checking accounts, saving accounts, mutual funds, bonds, shares, managed accounts, other assets, private lending, voluntary pension plans, or whole life insurance contracts. Liabilities include both mortgage and non-mortgage debt. Household disposable net wealth explicitly excludes future claims on public pensions or occupational pension plans, human capital, and the net present value stream of future labor income.

TABLE 1
The Oversampling Method in the SCF, the WAS, and the HFCS

Using individual information	
U.S.	List based on income tax information
Spain	List based on taxable wealth information
France	List based on taxable wealth information
U.K.	Tax returns at address level
Finland	Income information from register
Using geographic income information	
Belgium	Average regional income
Germany	Taxable income of municipalities
Using geographic information	
Austria	Vienna oversampled
Portugal	Lisbon and Porto oversampled
No oversampling	
Italy	No oversampling
Netherlands	No oversampling

Source: Author's construction based on Kennickell (2009a), HFCS (2013), and ONS (2012).

that unit non-response is positively correlated with wealth.[3] Whereas unit non-response is generally dealt with by rescaling the weights of all respondent households, differential unit non-response of wealthy households can only be dealt with effectively if weights are rescaled *selectively*.

Stratified sampling from a special sampling frame to oversample the wealthy allows for selective reweighting. This is the case for the SCF. The SCF uses a dual frame to sample households. A representative area probability sample is combined with a high-income sample which is drawn from a sampling frame constructed using federal tax returns. From the high-income sampling frame different strata are constructed, with higher strata having higher income (and greater expected wealth) and higher oversampling rates.[4] The different strata from the high-income frame allow one to address differential unit non-response by selective rescaling of weights. For the high income sample points of the SCF, a wealth index (an estimate of wealth based on income tax information) can be constructed. Kennickell and Woodburn (1997) report that sampled individuals with a wealth index between US$1 million and US$2.5 million have a response rate of 34 percent, whereas those with a wealth index between US$100 million and US$250 million have a response rate of 14 percent. This illustrates the differential unit non-response problem. Unfortunately, outside of the SCF, relatively little is known about the correlation of non-response with wealth.

Sampling frames used to oversample the wealthy differ dramatically across surveys. Table 1 provides an overview of the different methods used to oversample

[3]Household wealth survey specialists would generally agree that there is a strong presumption that non-response is positively correlated with wealth. Of course, the wealth of the non-respondent households is in principle unknown. However, for evidence that non-response is correlated with financial income in the SCF, see Kennickell and McManus (1993).

[4]Details are provided in Kennickell (2007).

the wealthy. Oversampling using information at the individual level of wealth or income is done in the U.S., the U.K., Spain, France, and Finland. Regional income information is used in Germany and Belgium. Austria and Portugal oversample the largest cities. Finally, in the Netherlands and Italy no oversampling is done.

One should expect that having wealth tax data to design different strata is better than income tax data, which in turn is clearly much better than having only auxiliary information to construct strata such as geography. The geographic criterion uses the idea that the rich tend to live in particular places. Of course, this is bound to be less precise than having direct income or wealth information to stratify samples. Alternatively stated, within a geographic stratum, the differential unit non-response problem will still exist.[5] Given these large differences in oversampling methods, it should not come as a surprise that the degree of oversampling dramatically differs across countries as does the possibility of adjusting selectively the weights for differential unit non-response. So both efficiency in top tail estimation and the magnitude of the bias will differ across countries.

Indeed, interestingly, and ultimately not surprisingly, these methods of oversampling correlate quite nicely with the fraction of the sample observations that are from the tail. Table 2 enumerates the survey sample size and the number of wealthy. Being wealthy is defined using three thresholds: having net wealth larger than €2 million, €1 million, and €500,000. In the SCF data, the fraction of observations from the tail is the largest: 15 percent of the SCF sample has wealth over €2 million. This is not just a reflection of the presence of greater wealth in the U.S. but, rather, is indicative of the very high rate of oversampling in the SCF. In Spain, the U.K., and France, three other countries using individual information to oversample the wealthy, 9, 5, and 4 percent of the sample respectively are households with wealth above €2 million. The two countries using geographic income information, Belgium and Germany, have 3 and 2 percent of the sample respectively with wealth above €2 million. The countries for which only geographic information is used, Portugal and Austria, only have a rather small 2 and 1 percent of the sample in the highest wealth category. In the cases of no oversampling, Italy and the Netherlands have 1 and 0 percent respectively. Finland is somewhat of an outlier. Although it uses individual income data from registers to oversample the wealthy, it still only has 1 percent of the sample with wealth above €2 million.

In practice, successful oversampling leads to many wealthy households in the sample, all with relatively low survey weights. Unsuccessful oversampling, or no oversampling at all, leads to few wealthy households in the sample, each with relatively high weights.

To provide further evidence that the high numbers of sample observations in the tail are really the result of oversampling, Table 3 shows the number of

[5]Obviously, the extent of the problem will be a function of the granularity of the sample design. For instance, in Germany, income tax statistics were used to identify small municipalities (defined as those with fewer than 100,000 inhabitants) with a large share of wealthy households. These municipalities are oversampled. Households within those municipalities are randomly selected. So within those municipalities differential unit non-response can still occur. Details of the German sample design are given in von Kalckreuth *et al*. (2012).

TABLE 2

Summary Statistics: The Number of Wealthy Households in the Survey Samples

	Sample Size	Absolute Number			Fraction of Sample		
		> 2 million	> 1 million	> 500,000	> 2 million	> 1 million	> 500,000
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Country samples with oversampling using individual information							
U.S.	6,482	965	1,259	1,692	0.15	0.19	0.26
France	15,006	638	1,712	3,522	0.04	0.11	0.23
U.K.	20,165	949	3,467	7,609	0.05	0.17	0.38
Spain	6,197	544	1,129	2,086	0.09	0.18	0.34
Finland	10,989	59	296	1,233	0.01	0.03	0.11
Country samples with oversampling using geographic income information							
Germany	3,565	85	246	654	0.02	0.07	0.18
Belgium	2,327	71	207	599	0.03	0.09	0.26
Country samples with oversampling using geographic information							
Austria	2,380	47	113	271	0.02	0.05	0.11
Portugal	4,404	24	87	252	0.01	0.02	0.06
Country samples with no oversampling							
Italy	7,951	78	300	1,075	0.01	0.04	0.14
Netherlands	1,301	2	32	172	0.00	0.02	0.13

Source: Author's calculations based on the SCF, the WAS, and the HFCS.

TABLE 3

Summary Statistics: The Number of Wealthy Households in the Population (Estimates Derived from the Survey Samples)

		Absolute Number			Fraction of Population		
	Households	HH > 2 million	HH > 1 million	HH > 500,000	> 2 million	> 1 million	> 500,000
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
U.S.	117,609,217	3,661,191	8,407,106	15,311,762	0.031	0.071	0.130
Spain	17,017,706	139,539	621,067	2,299,825	0.008	0.036	0.135
France	27,860,408	209,668	830,661	2,891,897	0.008	0.030	0.104
U.K.	24.717,237	694,752	2,974,635	7,386,081	0.028	0.120	0.299
Finland	2,531,500	6,555	34,632	158,436	0.003	0.014	0.063
Belgium	4,692,601	85,386	264,728	890,283	0.018	0.056	0.190
Germany	39,673,000	368,693	1,051,250	3,261,600	0.009	0.026	0.082
Austria	3,773,956	70,939	174,550	427,248	0.019	0.046	0.113
Portugal	3,932,010	14,141	64,443	185,746	0.004	0.016	0.047
Italy	23,817,962	265,782	901,176	3,100,288	0.011	0.038	0.130
Netherlands	7,386,144	2,895	83,813	508,482	0.000	0.011	0.069

Source: Author's calculations based on the SCF, the WAS, and the HFCS.

TABLE 4

The Forbes Billionaires List: Number of People and Nominal Wealth

	Date	Number of Individuals	Total Wealth of Forbes Billionaires	As Percentage of Country Wealth
U.S.	February 12, 2010	396	978.6	2.3
Germany	February 14, 2011	52	183.3	2.4
U.K.	February 13, 2009	37	84.8	0.7
Italy	February 14, 2011	14	46.6	0.7
Spain	February 13, 2009	12	28.3	0.6
France	February 12, 2010	11	60.1	0.9
Austria	February 14, 2011	5	13.0	1.2
Netherlands	February 12, 2010	3	4.8	0.4
Portugal	February 12, 2010	2	4.1	0.7
Finland	February 12, 2010	1	1.0	0.2
Belgium	February 12, 2010	1	1.9	0.1

Source: Author's calculations based on Forbes World's Billionaires, the SCF, the WAS, and the HFCS.

Note: Total wealth in billion euro.

households that those observations in the tail represent (i.e. their weight). For instance, for the category above €2 million, Spain has 544 sample observations (Table 2) representing 139,539 households (Table 3); whereas Germany has a sample of 85, representing almost three times as many households. The Netherlands, with no oversampling, only has two households in the sample above €2 million. One immediately observes how efficiency of tail estimation will be dramatically affected by the different rates of oversampling.

2.3. *Forbes Data*

Media lists of wealthy individuals provide another source of information on the wealth of the very top of the distribution. The SCF, the WAS, and the HFCS do not capture the absolute top. The SCF explicitly excludes individuals of the Forbes 400 wealthiest people in the U.S., presumably to preserve confidentiality (Kennickell, 2009a).

One widely known list is the annual Forbes World's billionaires list. An individual is on this list if his or her wealth is estimated to be above US$1 billion. For the purpose of this paper, the wealth of individuals on the list is recalculated in euro.[6]

Table 4 shows the number of individuals on the Forbes World's billionaires list, the total wealth they have, and their wealth as a percentage of total household wealth of the country (as estimated directly from the survey). Note that the SCF, WAS, and HFCS surveys differ slightly with respect to the reference years, the range depending on the country from 2008 to 2010. For most countries, the wealth recorded in the survey is what is recorded at the time of

[6]The Forbes list calculates wealth at the end of February for each year. I use the dollar/euro exchange rates of 1.2823 for 2009, 1.3572 for 2010, and 1.344 for 2011. Therefore, an individual is on the Forbes list if he or she has a wealth of approximately €740 million.

TABLE 5

The GAP: Maximum Nominal Wealth versus Minimum at Forbes (million euro)

	Maximum Wealth, SCF/WAS/HFCS	Minimum Wealth, Forbes
U.S.	806	737
France	153	810
U.K.	92	780
Spain	409	780
Finland	15	958
Germany	76	818
Belgium	8	1,920
Austria	22	1,560
Portugal	27	1,110
Italy	26	893
Netherlands	5	958

Source: Author's calculations based on Forbes World's Billionaires, the SCF, the WAS, and the HFCS.

Note: The maximum is over all five replicates of the dataset (for the HFCS and the SCF).

the interview. The period over which the set of households is interviewed lasts for multiple months. Therefore, I match the survey of each country with the date of the Forbes list that comes closest to the interview period. For the Netherlands, Finland, and Italy, where wealth is measured on December 31, I match the survey with the Forbes list of the following February. As the largest country, the U.S. has the most individuals on the list, with Germany and the U.K. second and third. Note that the individuals on the Forbes list can add significant information on the tail. For instance, the HFCS survey sample in Germany only has 85 individuals with wealth above €2 million, whereas there are 52 individuals on the Forbes billionaires list. For Italy, these numbers are 78 versus 14. For the Netherlands, there are more individuals on the Forbes billionaires list, namely three, than there are households in the HFCS sample above €2 million, namely only two.

Table 5 compares the maximum wealth found in the SCF, the WAS, and the HFCS with the minimum wealth of a person on the Forbes World's billionaires list. In principle, the SCF, the HFCS, and the WAS cover all resident households, and thus also potentially billionaires. In practice, only the SCF survey contains billionaires. In the SCF, there are sample observations that have greater wealth than the *poorest* Forbes billionaire. The very high oversampling rate of the wealthy in the SCF is clearly very effective. Contrary to the SCF, there is a serious gap between the richest household in the HFCS and the WAS and the poorest person on the Forbes list. Such a gap can be found in all countries. So the first observation is that none of the households in the HFCS or the WAS comes even close to the wealth levels of individuals on the Forbes billionaires list. The gap between the poorest person on the Forbes list and the wealthiest household in the surveys is very large. So, with the only notable example of the SCF, households that fall in between the richest household surveyed and the poorest Forbes billionaire are not in the sample. Note that among the HFCS surveys, the Spanish one shows the highest maximum wealth (€409 million). This is likely not a coincidence, as this survey arguable does a very good job in oversampling the rich (using wealth tax records).

The method of oversampling of the rich is correlated with this gap. The highest maximum wealth in the HFCS is found in Spain and France (€409 million and €153 million, respectively), two countries where oversampling is done based on individual wealth tax records. Also, the WAS for the U.K. has a relatively large maximum wealth of €92 million. The Netherlands, with no oversampling, has a rather low value of the maximum of wealth, namely €5 million. The other country with no oversampling, Italy, also has a low maximum value of wealth (€26 million). Also, the use of only geographic information, which is the case for Portugal and Austria, or geographic income information, in the cases of Belgium and Germany, does not guarantee observation of a high maximum of wealth.

To conclude, very rich households are not in the HFCS sample due to a combination of non-response and lack of effective oversampling, with the effectiveness varying greatly across countries. The few wealthy households in the tail that were sampled (in case of low oversampling) likely refused to answer the wealth surveys. Effectively, they are replaced by other households that have lower wealth. Only when a dramatic effort is being made to oversample, such as in the SCF, the WAS, and for France and Spain in the HFCS, can one observe a larger maximum of wealth.

3. A Pareto Law for the Tail of the Wealth Distribution

3.1. *The Pareto Distribution*

Davies and Shorrocks (1999) list two "enduring features of the shape of the distribution of wealth: 1) it is positively skewed 2) the top tail is well approximated by a Pareto distribution." The Pareto distribution has been used to approximate the tail of the wealth distribution in a number of distinct settings. First, extreme wealth observations have been modeled as a Pareto distribution. For instance, Ogwang (2011) estimates Pareto distributions for the 100 wealthiest Canadians for the years 1999–2008, Levy and Solomon (1997) estimate a Pareto distribution for the Forbes 400 wealthiest people in the U.S. for the year 1996, and Klaas *et al*. (2006) estimate Pareto laws using the Forbes 400 in the U.S. for the period 1988–2003.[7]

Another use of Pareto distribution has been to extrapolate existing tail observations "backward." For example, Kopczuk and Saez (2004) use long historical estate tax data to estimate the evolution of wealth of the top 1 percent of the U.S. wealth distribution. As before 1945 less than 1 percent of the population needed to file, they use a Pareto extrapolation to estimate the wealth share of the top 1 percent. A third use of the Pareto distribution has been to extrapolate truncated survey data "forward." For instance, Avery *et al*. (1988) extrapolated the first SCF data of 1983 (and the 1963 Survey of Financial Characteristics of Consumers) beyond US$60

[7]I follow the mainstream literature and approximate the top using a Pareto distribution. With available sample sizes, other distributions with long fat tails are often hard or impossible to distinguish from the Pareto distribution. For a study which compares the Pareto distribution with the log-normal and the stretched exponential, see Brzezinski (2014).

million by estimating first a Pareto distribution on the sample above US$10 million. I will show below that extrapolation using *only* the survey data leads to too low tail wealth estimates in the presence of differential unit non-response.

The Pareto distribution has the following complementary cumulative distribution function (CCDF):

$$P(W > w) = \left(\frac{w_{\min}}{w}\right)^{\alpha}, \tag{1}$$

defined on the interval $[w_{\min}, \infty)$ and $\alpha > 0$. The parameter $w_{\min}$ determines the lower bound on the distribution. The parameter α, also called the tail index[8], determines the "fatness" of the tail. The lower the value of α, the fatter is the tail, and the more concentrated is wealth.

Note that it is useful to keep the distinction between the theoretical Pareto distribution and the notion of a power law in a finite population. Finite populations that follow a power law can be seen as a (potentially very large) sample drawn from a Pareto distribution.

Imagine a finite population of N households, each having wealth at or above $w_{\min}$.[9] Let w_i be the wealth of household i, and denote by $N(w_i)$ the number of households that have wealth at or above w_i. We say that wealth in this population follows an (approximate)[10] power law if the empirical CCDF of the population follows approximately the CCDF of a Pareto distribution:

$$\frac{N(w_i)}{N} \cong \left(\frac{w_{\min}}{w_i}\right)^{\alpha}, \quad \forall w_i. \tag{2}$$

3.2. *Estimation of the Power Law on Samples from Complex Survey Designs*

Estimation on Simple Random Samples versus Samples from Complex Survey Designs

There exists a large literature on the estimation of power laws in simple random samples. For detail on different methods, the interested reader is referred to Gabaix (2009) and Clauset *et al*. (2009). However, estimation methods on simple random samples cannot simply be applied to samples from complex survey designs where observations have weights and are not i.i.d. In this section, I show how to adapt estimation methods of power laws for simple random samples to methods suitable for samples from complex survey designs, taking into account the weights of the sample points. As far as I can tell, this exposition is new to the literature.

The density of the Pareto distribution is given by:

[8]This term appears in Gabaix and Ibragimov (2011). Alternative terms appearing in the literature are "Pareto exponent" and "tail exponent."

[9]Note that these N households could be part of a larger population. Generally, $w_{\min}$ could thus be a large number. We only consider here the tail, that is, the N richest households.

[10]In reality, power laws will always be approximate in the data. However, for simplicity, "approximate" is dropped from the further discussion.

$$f(w) = \frac{\alpha w_{\min}^{\alpha}}{w^{\alpha+1}}, \tag{3}$$

so that it is straightforward to show that the maximum likelihood estimator of α from a simple random sample of n observations $\{w_i, i=1, \ldots n\}$ drawn from a Pareto distribution with known w_{min}, is given by

$$\tilde{\alpha_{ml}} = \left[\sum_{i=1}^{n} \frac{1}{n} \ln \left(\frac{w_i}{w_{\min}} \right) \right]^{-1}. \tag{4}$$

Now, $\frac{n-1}{n} \tilde{\alpha_{ml}}$ gives an unbiased estimate of α (Rytgaard, 1990).

Without some adjustment, the maximum likelihood estimator should not be used on complex survey data. The sampling observations of the SCF, the WAS, and the HFCS, due to the complex survey design, are not i.i.d., a requirement for maximum likelihood. Because the exact detail of the sampling method is unknown (the SCF, the WAS, and the HFCS only provide weights, but not the exact sampling detail), a true likelihood cannot be constructed. Due to stratification, clustering, and possible oversampling, some observations will have a much higher likelihood of occurring in the sample than others. The use of a maximum likelihood estimator on such samples would clearly lead to erroneous results.

Remember that survey weights represent the number of households that the sample point represents. One can therefore construct a pseudo-maximum likelihood estimator that incorporates the weights of the observations as follows. Denote by N_i the survey weight of a household sample observation. Sort the sample observations from highest to lowest wealth $w_1, w_2, w_3, \ldots$. Thereafter, consider the first n sample observations (i.e. those with the highest wealth). Denote by N the sum of the survey weights of the first n observations, $\sum_{i=1}^{n} N_i = N$. This represents an estimate of the number of households that have wealth at least as high as w_n.

The pseudo-maximum likelihood estimate of the tail index is defined by

$$\tilde{\alpha_{\text{pml}}} = \left[\sum_{i=1}^{n} \frac{N_i}{N} \ln \left(\frac{w_i}{w_n} \right) \right]^{-1}. \tag{5}$$

The pseudo-maximum likelihood estimator has the same form as the maximum likelihood estimator but takes into account the weights of the sample observations. Sample observations that represent more households have a larger weight and are therefore weighted more in the estimation.

The power law relationship (2) also leads heuristically to an alternative estimation method in simple random samples. Start from a population of N households that follows a power law as in equation (2). Assume that a simple random sample $\{w_i, i=1, \ldots n\}$ is drawn from the population; observations are sorted from largest to smallest: $w_1 \geq w_2 \geq w_3$, and so on. Then i denotes the number of sample observations that have wealth at or above w_i, also called the rank of the observation. So the rank of the richest household in the sample is one, the rank for the second richest is two, and so on. Now, the tail distribution (or one minus the cumulative relative frequency) in the sample provides an estimate of the tail distribution in the population. In other words,

$$\frac{i}{n} \cong \frac{N(w_i)}{N}, \quad \forall w_i. \tag{6}$$

As the sample gets larger, the estimate will obviously become closer to the true population figure. Combining this with the power law relationship in the population (2), we obtain

$$\frac{i}{n} \cong \left(\frac{w_{\min}}{w_i}\right)^{\alpha}, \quad \forall w_i. \tag{7}$$

Taking logs on both sides and rearranging, we obtain a "log-rank–log-size" relationship, that is, the log of the rank of the observation is a downward-sloping function of the log of wealth:

$$\ln(i) = C - \alpha \ln(w_i), \tag{8}$$

with $C = \ln(n) + \alpha \ln(w_{\min})$.

It is well known that for a simple random sample drawn from a Pareto distribution, a linear regression of the log-rank–log-size relationship leads to a biased estimate of α (see, e.g., Aigner and Goldberger, 1970). Gabaix and Ibragimov (2011) show that the bias can be removed (up to first order) by subtracting 1/2 from the rank of the observation. They propose the following regression:

$$\ln(i - 1/2) = C - \alpha \ln(w_i). \tag{9}$$

In a complex survey sample, the survey weights have to be taken into account. For such a survey, rank the sample households according to wealth. That is, the wealthiest household has wealth w_1 and a survey weight of N_1, the second-wealthiest household has wealth w_2 and survey weight of N_2, and so on. Define $\bar{N}$, the average weight of a sample point (i.e. $\bar{N} = \frac{\sum_{j=1}^{n} N_j}{n}$), and $\overline{N_{fi}}$, the average weight of the first i sample points (i.e. $\overline{N_{fi}} = \frac{\sum_{j=1}^{i} N_j}{i}$). Then one can show that taking the weights into account leads to the following regression (the derivation of which can be found in Part III of the Appendix (in the Online Supporting Information):

$$\ln\left((i - 1/2)\frac{\overline{N_{fi}}}{\bar{N}}\right) = C - \alpha \ln(w_i). \tag{10}$$

Combining Survey with Forbes Data

As discussed above, the SCF, the WAS, and the HFCS do not contain the very top of the wealth distribution. The Forbes data can easily be combined with the survey data in the regression method of estimation. First pool the Forbes with the survey data and rank the households from highest to lowest wealth. The richest Forbes individual will have wealth w_1 and a weight of 1, the second-wealthiest Forbes individual has wealth w_2 and a weight of 1, and so on. In other words, the Forbes observations are treated *as if* they were sample points with a weight of 1. The richest household in the survey will have wealth w_{K+1} (if there are K Forbes individuals richer than this household) and survey weight N_{K+1}, and so on. Equation (10) can then be estimated on the pooled dataset. Note that combining survey with Forbes data raises the issue of measurement error in both datasets. Capehart (2014) discusses measurement error problems in rich lists. In addition, a combination of both datasets in the regression method is only warranted under the assumption that both the sample and the rich list are consistent with the same Pareto distribution.

Equation (2) implies that if the data follows a power law, there is a linear relationship between the empirical CCDF and wealth (scaled by $w_{\min}$) on a graph with a log–log scale. Figures of this relationship are provided in Part II of the Appendix (in the Online Supporting Information).

3.3. *Monte Carlo Results: The Power Law When Survey Data Has Differential Unit Non-response*

The presence of unobserved and uncorrected differential unit non-response correlated with wealth will have serious consequences for tail wealth estimation. Such non-response causes the empirical sample distribution of the tail to systematically differ from the actual tail distribution in the population. As wealthy households respond less frequently when being sampled than less wealthy ones, the tail in the survey sample will be truncated. This causes the tail index to be biased upward, that is, showing a lower degree of wealth concentration. Total wealth in the tail will be biased downward.

How biased are estimates of α in the presence of differential unit non-response? How much bias reduction is possible when oversampling the wealthy and selectively correcting for non-response (as in the SCF)? Can extra observations of wealthy individuals from rich lists in the regression method reduce bias and increase precision, especially when oversampling is lacking or limited? How much improvement of estimates can we expect? These questions are important. First, they determine our degree of confidence in estimates of concentration of wealth in the tail of the population. Second, combining rich lists with survey data potentially provides a method of improving on estimates of the level of tail wealth.

To get a handle on those questions, a Monte Carlo study is performed. The central idea is to model a wealth survey in the presence of differential unit non-response under two possible sampling schemes: no oversampling versus oversampling of the wealthy. The no-oversampling case corresponds to a simple random sample of the population, whereas the oversampling of the wealthy corresponds to a stratified sampling where the population is divided in wealth strata. Obviously, both sampling schemes are approximations to the complicated (unfortunately unknown) complex survey designs. In reality, to allow for oversampling of rich households, stratification is based on some characteristic (e.g. income) correlated with wealth.

The no-oversampling case is more relevant to understanding the results of surveys such as those in the Netherlands and Italy, whereas the oversampling case is more relevant for surveys such as the SCF and the Spanish and French HFCS. The other surveys are somewhat in the middle, as they attempt to oversample the rich but are not as successful as in the case of the SCF.

First, I explain the main features of the Monte Carlo experiment. Consider a large country with a tail population of 1 million households, each with wealth above €1 million. Each individual household's wealth is drawn from a Pareto distribution with given tail index α, and threshold $w_{\min} = 1$ million. For instance, such a country could be imagined to be of roughly the size of Germany or France.[11] Imagine further that all households with wealth above €740 million are also on a media rich list, say, a dollar billionaires list. It is assumed that the rich list is exhaustive.

A sample is drawn from this tail population, with a sample size of 750 households. Some households respond to the survey while others do not, according to the non-response mechanism in place. From the sample of survey respondents, the tail index is estimated using the two estimation methods (equations (5) and (10)). For the regression method there are two estimates, one using only the survey observations and another one combining the survey observations with the rich list. To construct mean estimates and standard errors of the tail index, the experiment is performed 20,000 times; that is, a new population of 1 million households is drawn from the same Pareto distribution, a new sample of 750 households is drawn from that population, and the tail index is estimated from the respondents (with or without the rich list).

The experiment is performed for ten different values of α (i.e. $\alpha = 1.1, 1.2, \ldots 2.0$). According to Gabaix (2009), the tail exponent of wealth found in earlier studies is around 1.5, so that the interval of values of α considered here should suffice. Each experiment for a given α is performed for the two sampling strategies.

The differential unit non-response mechanism attempts to model a reasonable relation of wealth with non-response in the population, that is, a differential unit non-response that mimics reality. There is relatively little existing earlier research on this issue that would guide one in choosing a reasonable function that links wealth with non-response. However, Kennickell and Woodburn (1997) provide response rates for different strata of the wealth index from the list sample for the 1989, 1992, and 1995 SCF. The response rates across different strata are relatively stable across different SCF waves, indicating that the positive correlation of wealth with non-response is a relatively robust feature of the SCF, and one can assume that it is likely for surveys in other countries. In the 1992 SCF, individuals with a wealth index between US$1 million and US$2.5 million have a response rate of 34.4 percent. This rate gradually declines to 14.3 percent for individuals with a wealth index between US$100 million and US$250 million. Households' non-response probability as a function of wealth is then calibrated to mimic the non-response rate as a function of

[11]According to the HFCS survey results in Germany, about 1 million households have wealth above €1 million; in France, this is about 800,000 households. As we are only interested in the tail, the Monte Carlo only models the tail of the distribution.

the wealth index in the 1992 SCF. This is done in the following way. The non-response rates of the six strata in the 1992 SCF are regressed on the log of wealth, taking the midpoint of the stratum and translating back into 2010 euros. This regression results in the following relationship between the probability of non-response and the log of wealth: $P(\text{non-response}) = 0.097167 + 0.036594 * \ln(\text{Wealth})$. This relationship is our differential unit non-response mechanism.[12] The combination of a random sample of 750 households with the non-response function defined above leads to roughly 280 households responding and 470 non-responding. According to the HFCS in Germany, there are 246 households in the sample with wealth above €1 million.

In the no-oversampling case, the survey sample is a simple random sample where only the aggregate non-response rate is observed. For the case of oversampling of the wealthy, one first needs to define the oversampling mechanism. Unfortunately, the SCF and the HFCS only provide very limited information on this issue. For the SCF, oversampling is done using seven strata based on a calculated wealth index derived from income tax data. For the Spanish HFCS, eight strata are constructed using tax wealth data. However, the oversampling rates are not made public. Both the SCF and the Spanish HFCS report that oversampling is done at progressively higher rates for higher strata. To approximate, in a simplified way, the SCF and the Spanish HFCS, we use four strata. We assume that the tail population is divided into four strata corresponding to the quartiles of the distribution. We assume again a sample size of 750 households with an increasing oversampling rate: out of the lowest stratum 75 households are sampled, with 150 out of the second, 225 out of the third, and 300 out of the last.[13]

Non-response correction of the weights for the no-oversampling case is only based on aggregate non-response rates. Survey weights are constructed for the responding households so that they sum up to 1 million. For instance, if all 750 households respond, the household weight for each individual would be equal to $10^6/750$. When fewer than 750 households respond, divide the 750 into non-responding, N_{nr}, and responding, N_r, households. Then each responding household is given a weight of $(10^6/750) * (750/N_r)$, so that household weights again sum up to 1 million. For the oversampling case, the survey weight correction uses the stratum non-response rate. For example, for the first stratum, divide the 75 households into non-responding, N_{nr}^1, and responding, N_r^1, households. Then each responding household is given a weight of $(10^6/(4*75)) * (75/N_r^1)$. This

[12]The aggregate expected non-response probability of this non-response mechanism can be found by taking the expectation of $0.097167 + 0.036594 * \ln(\text{Wealth})$ (where wealth has a Pareto distribution). This itself will depend on the threshold of 1 million and α. The formula for this expectation is $P(\text{non-response}) = 0.097167 + 0.036594 * \ln(10^6) + (0.036594/\alpha)$. This gives an aggregate non-response rate of between 62.1 percent (for $\alpha = 2$) and 63.6 percent (for $\alpha = 1.1$) for this tail population. This number looks high but is actually quite reasonable. For instance, the aggregate non-response rate in the German HFCS is 81.3 percent (HFCS, 2013) (this aggregate includes all households not just the tail), even higher than assumed in the Monte Carlo.

[13]Note that oversampling with an identical total sample size of the no-oversampling case will lead to a lower number of actual observations (as we sample more out of the higher non-response regions).

TABLE 6

Monte Carlo Estimates of the Pareto Tail Index Using Different Estimation Methods under Differential Unit Non-response

	No Oversampling				Oversampling of the Rich				
α	α_{ml}	α_{reg}	α_{regfor}	Obs	α_{ml}	α_{reg}	α_{regfor}	Obs	Rich List
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1.10	1.21	1.22	1.09	273	1.13	1.18	1.10	264	698
	0.07	0.10	0.01	13	0.03	0.08	0.01	13	26
1.20	1.31	1.32	1.19	275	1.23	1.28	1.20	267	360
	0.08	0.11	0.01	13	0.03	0.09	0.01	13	19
1.30	1.41	1.41	1.29	277	1.33	1.38	1.30	269	186
	0.08	0.12	0.02	13	0.03	0.09	0.02	13	14
1.40	1.51	1.51	1.39	278	1.43	1.47	1.40	271	96
	0.09	0.13	0.02	13	0.04	0.10	0.02	13	10
1.50	1.61	1.61	1.49	280	1.52	1.57	1.50	273	50
	0.10	0.13	0.03	13	0.04	0.11	0.03	13	7
1.60	1.71	1.71	1.60	281	1.62	1.67	1.60	275	26
	0.10	0.14	0.04	13	0.04	0.12	0.04	13	5
1.70	1.80	1.81	1.70	282	1.72	1.77	1.70	276	13
	0.11	0.15	0.06	13	0.04	0.12	0.06	13	4
1.80	1.90	1.91	1.80	283	1.82	1.87	1.80	277	7
	0.11	0.16	0.08	13	0.04	0.13	0.07	13	3
1.90	2.00	2.00	1.91	284	1.92	1.97	1.90	278	4
	0.12	0.17	0.10	13	0.05	0.14	0.09	13	2
2.00	2.10	2.10	2.01	284	2.02	2.07	2.00	279	2
	0.12	0.17	0.13	13	0.05	0.14	0.12	13	1

Notes: Reported are mean estimates of the Pareto tail index. Standard errors are reported in the line below the mean. Means and standard errors are derived from 20,000 Monte Carlo iterations. In each iteration 1 million households draw wealth from a Pareto distribution with true tail index given in column (1). From each population, a survey sample of 750 households is drawn. Each household drawn has a non-response probability equal to 0.097167+0.036594*log(wealth). Estimates of tail index using maximum likelihood are in columns (2) and (6). Estimates using the regression method excluding the rich list are in columns (3) and (7). Estimates using the regression method including rich list are in columns (4) and (8). Columns (5) and (9) report the mean number of respondent observations (and the standard error). Column (10) reports the mean number of observations on the rich list (and the standard error).

stratum-specific non-response correction is the key to reducing the bias caused by differential unit non-response.[14]

Table 6 presents the results of the Monte Carlo experiments. Reported are mean estimates and standard errors of the Pareto tail index under the two sampling scenarios, using the different estimation methods. Column (1) shows the true α, columns (2)–(5) show the results under the non-oversampling case, and columns (6)–(9) show the results under the oversampling case. Column (10) shows

[14]I experimented with different degrees of oversampling, that is, keeping the sample size constant but sampling progressively more out of the higher strata and less out of the lower strata. The results are available in Part IV of the Appendix (in the Online Supporting Information). These experiments show that the bias reduction from no oversampling to oversampling does not vary much across different degrees of oversampling. This shows that the bias reduction is mainly due to the stratum-specific non-response correction which oversampling makes possible, and not so much to the degree of oversampling itself.

the number of households on the rich list, that is, the number of households with wealth in excess of €740 million.

The (pseudo-)maximum likelihood estimates α_{ml} are in columns (2) and (6). They are clearly different under the two sampling cases. As expected, under no oversampling, the estimates of α are significantly upward biased, indicating an estimated lower concentration of wealth in the tail than the true concentration. The bias is around 0.11 for all α's. The upward bias in the estimated α's is much reduced, to around 0.02 for the oversampling case (column (6)). Note that also the standard error is more than cut in half, which is due to the efficiency gained by stratified sampling.

The regression estimates, α_{reg}, using only the survey data, are in columns (3) and (7). For both the no-oversampling and oversampling cases, they show an upward bias. In the no- oversampling case, the estimates are practically identical to the (pseudo-)maximum likelihood estimates thereby showing the same upward bias of around 0.11. In the oversampling case, they show a lower upward bias than the no-oversampling case, but a higher upward bias than the (pseudo-)maximum likelihood estimates.

The regression estimates, α_{regfor}, derived from combining the survey data with the observations on the rich list, are reported in columns (4) and (8). The number of observations from the rich list are shown in column (10). Obviously, the number decreases as true α increases. When α is equal to 1.5, there are on average 50 observations on the rich list (with a standard deviation of 7) (out of, we recall, a population of 1 million: compare this with an actual number of 53 in the German case). This drops to only two observations when α is equal to 2. The improvement of the estimate of the tail index, in terms of a reduction in bias, under no oversampling is dramatic. Essentially, when including the rich list with the survey data in the regression method, the tail index is estimated with almost no bias (a tiny upward or downward bias of 0.01 occurs). It is also important to note that the reduction in the standard error is impressive. Again, as one should expect, the reduction in the standard error is much larger when the tail index is lower, that is, the number of observations on the rich list is higher. But strikingly, even when the rich list contains very few individuals, both the bias in the estimate of α almost disappears and the standard error is reduced. Similarly, combining the survey data with the observations on the rich list also helps in the oversampling case. The tail index becomes unbiased. So, importantly, the rich list is useful for both the no-oversampling case and the oversampling case.

Figure 1 shows the intuition for the reduction in bias, and lower standard error, when a rich list is added to the data. It shows the empirical CCDF of a Monte Carlo sample and the rich list, together with the true power law from which the Monte Carlo sample was drawn (for the no-oversampling case).[15] It also shows the power law implied by the three estimates of the tail index, the pseudo-maximum likelihood, and the two estimates using the regression method.

[15]The empirical CCDF is constructed as follows. The individuals on the rich list each have a weight of one. The weights of the survey observations have a weight of $(10^6/750) * (750/N_r)$. The empirical CCDF, $P(X > x)$, is then given by the sum of the weights of sample and rich list observations above wealth x, divided by 10^6.

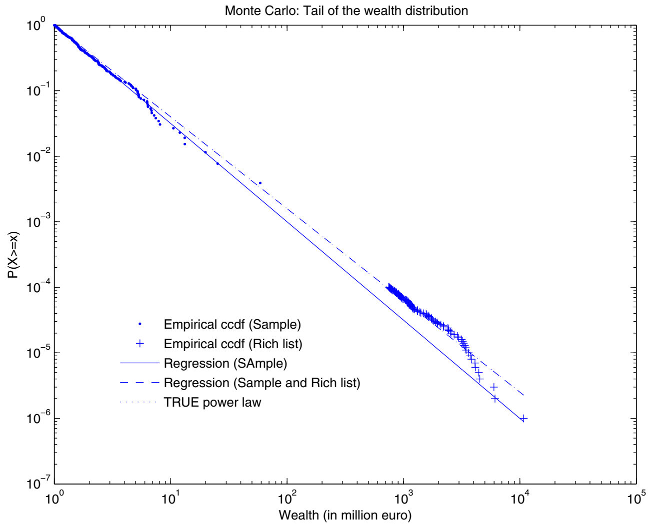

Figure 1. A Monte Carlo Example of the Tail of the Wealth Distribution [Color figure can be viewed at wileyonlinelibrary.com]

Due to the non-response, the empirical CCDF from the sample observations of wealthy households will be below the line implied by the true power law, that is, it will provide an underestimate of the relative frequency of the households that are richer. On the contrary, the households on the rich list will follow the true power law.[16] By adding the rich list to the survey sample, the regression line shifts to the right. Intuitively, by adding the rich list in the presence of differential non-response, the regression line gets "anchored." This will be reflected in a lower standard error of the slope of the regression line, and a lower (to almost no) bias.

The ultimate interest in the estimation of the power law is to provide an estimate of total wealth in the tail. Total wealth can be directly calculated from the estimated power law. Alternatively, total wealth in the tail can be calculated from the survey directly as the weighted sum of wealth of the sample; we recall that survey weights sum up to population totals. To see how far off estimated wealth is from the truth, Table 7 shows the total wealth in the population estimated from the survey sample and from the estimated power laws, as a ratio to the true total wealth in the population.[17] A ratio of 1 signifies no bias in estimated wealth.

Under no oversampling, there is a downward bias in the estimated wealth from the survey sample. The size of the bias depends very much on the level of the tail index. The intuition is clear: with higher tail indexes, the bias gets smaller. A higher tail index indicates a lower degree of wealth concentration at the top, so

[16]Note that even in a population of 1 million individuals drawn from a Pareto distribution there is sampling variation, that is, the number of individuals on the rich list (and their wealth) will vary. So although the extreme rich are drawn from the true Pareto distribution, they do not position themselves exactly on the true CCDF. This can be seen in Figure 1, where the *richest* of the rich are below the true Pareto CCDF.

[17]The true total wealth in the population is simply the sum of the wealth of the 1 million households.

TABLE 7
Monte Carlo Estimates of the Tail Wealth as a Proportion of the Actual Tail Wealth

	No Oversampling				Oversampling of the Rich			
α	Survey Estimate	α_{ml}	α_{reg}	α_{regfor}	Survey Estimate	α_{ml}	α_{reg}	α_{regfor}
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1.10	0.53	0.76	1.00	1.39	0.59	1.07	1.13	1.30
	0.33	0.79	5.63	0.23	0.32	0.32	6.21	0.20
1.20	0.67	0.78	0.83	1.09	0.73	0.96	0.91	1.06
	0.32	0.21	0.69	0.10	0.29	0.15	0.93	0.10
1.30	0.77	0.83	0.86	1.04	0.82	0.96	0.91	1.02
	0.35	0.15	0.30	0.06	0.27	0.09	0.26	0.06
1.40	0.83	0.87	0.89	1.02	0.87	0.97	0.93	1.01
	0.37	0.12	0.19	0.04	0.23	0.07	0.39	0.04
1.50	0.87	0.90	0.91	1.01	0.91	0.97	0.94	1.01
	0.25	0.10	0.15	0.04	0.23	0.05	0.14	0.04
1.60	0.90	0.92	0.93	1.01	0.93	0.98	0.95	1.00
	0.15	0.08	0.12	0.04	0.12	0.04	0.11	0.04
1.70	0.92	0.93	0.94	1.01	0.95	0.98	0.96	1.00
	0.13	0.07	0.11	0.04	0.28	0.03	0.10	0.04
1.80	0.93	0.95	0.95	1.00	0.96	0.99	0.97	1.00
	0.10	0.06	0.09	0.05	0.08	0.03	0.09	0.05
1.90	0.94	0.95	0.96	1.00	0.97	0.99	0.97	1.01
	0.10	0.06	0.09	0.06	0.07	0.03	0.08	0.05
2.00	0.95	0.96	0.96	1.00	0.97	0.99	0.98	1.01
	0.07	0.05	0.08	0.06	0.06	0.02	0.07	0.06

Notes: Reported are means of the ratio of the estimated tail wealth to the actual tail wealth. Standard errors are reported in the line below. Means and standard errors are derived from 20,000 Monte Carlo iterations as described in the footnote to Table 6. The estimated tail wealth used to construct the ratio in columns (2) and (6) is calculated from survey only. The estimated tail wealth used to construct the ratio in columns (3)–(5) and (7)–(9) is constructed using the estimated Pareto tail index.

that differential non-response is less of a problem (with a higher tail index, the very wealthy are much less numerous). The wealth estimate using the survey sample is expected to be 13 percent lower at a tail index level of 1.5 (the level mentioned by Gabaix, 2009) or even lower, in the case of power laws with low tail indexes. However, the striking feature of the ratio of estimated wealth from the survey to true wealth (column (2)) is not so much the bias, but its large standard error. Estimating total tail wealth from the survey directly implies having a very imprecise estimate! This obviously depends on sample size. Note, however, that the Monte Carlo has around 280 observations of wealth above 1 million. This is a larger number than the observations above 1 million for the HFCS surveys of Germany (246), Belgium (207), Austria (113), Portugal (87), and the Netherlands (32). So precision for some surveys in reality, compared to the Monte Carlo, is likely even worse. Estimating a power law and then calculating the wealth using the estimated law reduces the standard error enormously. The biggest reduction is when using the regression method including the rich list. For example, the reduction in standard error is by a factor of 6 in the case of $\alpha=1.5$. Although that leads to a very small upward bias of wealth estimates, the reduction in variability of the estimate is quite substantial.

For the oversampling case, the biases are reduced but still notable. For example, the wealth estimate using the survey sample is still expected to be 9 percent too low at a tail index level of 1.5. It is again preferable to add a rich list. The bias practically disappears and the standard error is reduced. An exception occurs when α is small, around 1.2. Note that biases and standard errors are generally large for such low α. This is not surprising as α approaches 1, the mean of the Pareto distribution approaches infinity. In any case, such low α are not commonly found anyway.

Combining all these results, the Monte Carlo shows that differential unit non-response clearly biases top tail wealth estimates downward. Under the assumptions of the Monte Carlo, the bias could easily be more than 10 percent. It also shows that when a rich list is available, adding it to the survey data and estimating wealth through the estimated power law is a reasonable idea; it removes the downward bias caused by differential unit non-response and reduces the variance of the estimated wealth. When no rich list is available, the difference in downward bias between the (pseudo-)maximum likelihood and regression method are on average not that large, except that the (pseudo-)maximum likelihood estimates show lower variance. These ideas are taken up in the next section, where the results of power law estimation are shown.

4. Estimation Results

This section provides new estimates of the share of wealth held by the top 1 and 5 percent derived from the U.S. SCF, the U.K. WAS, and the HFCS.[18] A detailed set of estimation results is tabulated and available in Part I of the Appendix (in the Online Supporting Information). Here, only key results are discussed. The emphasis is on documenting the downward bias when estimates are based on surveys only and its remedy, including extreme observations and estimating a Pareto tail using the regression method.

Percentage shares for the 1 and 5 percent richest households for various countries can also be found elsewhere, although they are somewhat scattered in the literature. Summaries can be found in Davies *et al*. (2010) and, more recently, in Roine and Waldenström (2014). Using the SCF, Wolff (2006) provides historical U.S. series for the 1980s and 1990s, and Kennickell (2009b) calculates series for the period 1989–2007. Saez and Zucman (2014) construct historical series for the U.S. based on capitalized income tax data, and Kopczuk and Saez (2004) provide top wealth shares for the U.S. based on estate tax returns. Piketty (2014) discusses the evolution of top wealth shares going back as far as the eighteenth century, using various data sources.

Estimates of top wealth shares are constructed either as direct estimates from the surveys or by replacing the tail observations of the surveys with the estimated Pareto distribution. The Pareto distribution is estimated either using the pseudo-maximum likelihood method or the regression method. For this last method, estimates using the survey only and using the survey combined with the Forbes

[18]For the HFCS and SCF data, the estimates are based on all five implicates of the multiple imputed datasets and standard errors are provided using the bootstrap weights.

World's billionaire list are given. As it is unclear exactly where the tail starts, and to investigate the variability of tail estimates depending on the level of wealth where the tail starts, estimates are presented for up to six different threshold levels (€10 million, €5 million, €3 million, €2 million, €1 million, and €500,000). Estimates for all six thresholds are provided for the U.S., the U.K., France, and Spain. For the other countries, due to too few survey observations at the top, estimates are provided for the three lower thresholds (€2 million, €1 million, and €500,000). Using lower thresholds increases the sample size over which the Pareto distribution is estimated. However, there is a tradeoff. On the one hand, the increased sample size leads to more precise tail index estimates, but on the other it also includes observations that potentially do not obey the Pareto tail behaviour. This itself might lead to biased estimates. The use of a high level of the threshold leads to fewer observations and hence more imprecise estimates, but is more likely to restrict the estimation on a sample that truly follows Pareto tail behaviour.

Alternatively, a "best-fit" threshold can be found using a methodology developed in Clauset *et al*. (2009). First, the Pareto tail is estimated on different threshold levels. Second, at each threshold level, the fit of the Pareto tail is tested using a Kolmogorov–Smirnov test. The Kolmogorov–Smirnov test statistic measures the maximum distance between the CDF of the data and the CDF of the estimated Pareto distribution. The "best-fit" threshold is the one which leads to the smallest maximum distance, therefore providing the threshold where the Pareto tail has the best fit. Following this methodology, the Pareto tail was estimated on a fine grid of thresholds between €100,000 and €10 million (varying the threshold each time by by €25,000). Detailed results of the Pareto tail index and the "best-fit" threshold are given in Part V of the Appendix (in the Online Supporting Information). However, a word of caution is needed. The Clauset *et al*. (2009) methodology was developed with simple random samples in mind and not for complex survey data that have differential non-response problems. Therefore, just like a Pareto tail index that might be biased using such data, a "best-fit" threshold might not necessarily coincide with the "true" threshold. Notwithstanding this caveat, the estimates of top wealth shares obtained using the Pareto tails at the "best-fit" thresholds fall pretty much within the intervals provided by the six threshold levels considered above. The discussion below considers the estimates of the wider set of thresholds considered above.

The Monte Carlo results showed that in the presence of differential unit non-response, Pareto tail index estimates from the survey data only are biased upward. Including extreme observations, the tail index estimates should drop and become unbiased. The drop should be highest when there is no oversampling. Table 8 shows the average[19] Pareto tail index estimates for the regression method both when excluding or including the Forbes billionaires. For all countries, except Portugal, the Pareto tail index drops when adding the Forbes billionaires. The drop is the largest for the Italian and Dutch surveys, the two surveys that do not oversample the rich. The lowest drop is observed for the U.S. SCF and the Spanish survey, both of which use heavy oversampling.

[19]The average is taken over the three estimates corresponding to the €2 million, €1 million, and €500,000 thresholds, as these are available for all countries.

TABLE 8

The Average Pareto Tail Index

	Regression Excluding Forbes	Regression Including Forbes	Δ
	(1)	(2)	(3)=(1)−(2)
Countries using individual information to oversample			
U.S.	1.59	1.52	0.07
U.K.	2.05	1.74	0.31
France	1.76	1.62	0.14
Spain	1.77	1.69	0.07
Finland	2.11	1.88	0.23
Countries using geographic income to oversample			
Germany	1.68	1.39	0.29
Belgium	2.18	1.87	0.31
Countries using geographic information to oversample			
Austria	1.65	1.46	0.20
Portugal	1.45	1.47	−0.02
Countries using no oversampling			
Italy	2.02	1.58	0.44
Netherlands	5.09	1.48	3.61

Source: Author's calculations based on the SCF, the HFCN, the WAS, and Forbes World's Billionaires.

Notes: Column (1) provides average of estimated Pareto tail indexes using the regression method on the survey data at three thresholds: €500,000, €1 million, and €2 million. Column (2) provides the same as column (1) when adding Forbes billionaires to the survey sample. Column (3) shows the average reduction in the Pareto tail index when Forbes billionaires are added.

Table 9 shows a summary of the estimates for the top 1 percent shares. The first column shows the estimates directly calculated from the surveys. In the presence of differential non-response, these should be biased downward. Again, for surveys using oversampling of the wealthy, the bias should be smaller. The second column shows the range of estimates when the tail observations are replaced by an estimated Pareto distribution using the regression method applied on the tail survey observations including the Forbes data.[20] As expected, estimates of the percentage wealth share of the top 1 percent of households are affected the most for countries with no or low oversampling, the Netherlands and Italy. Indeed, a direct sample calculation for the Netherlands results in an estimate of a percentage wealth share of 9 percent, the lowest across all countries. Including the three Forbes observations in the regression method, the wealth share of the top 1 percent is estimated between 10 and 19 percent. Such an increase in the estimated percentage suggests that 9 percent is a severely downward-biased estimate of wealth at the tail in the Netherlands. Likewise for Italy, the top 1 percent share calculated directly from the survey is 14 percent. Including the Forbes data and

[20]Note that in this case, total wealth is estimated using the estimated Pareto tail, that is, using the survey to calculate the sum of wealth below the Pareto threshold and adding to this the wealth in the estimated Pareto tail.

TABLE 9
THE PERCENTAGE WEALTH SHARE OF THE TOP 1 PERCENT OF HOUSEHOLDS

	Survey	Regression Including Forbes	Δ
	(1)	(2)	(3)=(2)−(1)
Countries using individual information to oversample			
U.S.	34	31–37	−3 to +3
U.K.	13	14–18	+1 to +5
France	18	19–21	+1 to +3
Spain	15	15–17	+0 to +2
Finland	12	13–15	+1 to +3
Countries using geographic income to oversample			
Germany	24	32–34	+8 to +10
Belgium	12	15–16	+3 to +4
Countries using geographic information to oversample			
Austria	23	31–32	+8 to +9
Portugal	21	23–27	+2 to +6
Countries using no oversampling			
Italy	14	20–21	+6 to +7
Netherlands	9	10–19	+1 to +10

Source: Author's calculations based on the SCF, the HFCN, the WAS, and Forbes World's Billionaires.

Notes: Column (1) provides the top 1 percent share of wealth directly derived from the surveys. Column (2) provides the range of estimates when the tail is replaced by the estimated Pareto distribution using the sample and Forbes data. The Pareto distribution is estimated at thresholds of €500,000, €1 million, €2 million, €3 million, €5 million, and €10 million for the U.S., the U.K., France, and Spain, and at thresholds of €500,000, €1 million, and €2 million for other countries.

estimating a power law, the share rises to a range between 20 and 21 percent. The top 1 percent share in Italy is therefore relatively insensitive to the threshold.

From the SCF, the wealth share in the U.S. calculated from the survey is 34 percent, while it is estimated to be between 31 and 37 percent when including the Forbes data and estimating a power law. It is interesting to note that the low estimate of 31 percent only occurs if the tail is estimated with a threshold of €500,000. Likely, this threshold is too low.[21] Replacing SCF observations with a Pareto tail from €10 million onwards leads to the higher estimate of 37 percent, three percentage points higher than the SCF. Note that the SCF explicitly excludes the Forbes 400, which have an estimated wealth of 2.3 percent of total household wealth. The discrepancy between the SCF survey estimate, 34 percent, and the estimate of 37 percent can therefore largely be explained by the addition of the Forbes billionaires. Saez and Zucman (2014) obtain an estimate of 39.5 percent.[22] This number is hard to compare, however, as they are using a completely

[21]The "best-fit" threshold for the SCF is €3.1 million. Using the Pareto tail at this threshold, the top 1 percent wealth share is 37 percent.

[22]This figure is for the year 2010. See Table B1 in the Appendix to Saez and Zucman (2014), available at http://eml.berkeley.edu/saez/.

TABLE 10

The Percentage Wealth Share of the Top 5 Percent of Households

	Survey	Regression Including Forbes	Δ
	(1)	(2)	(3)=(2)−(1)
Countries using individual information to oversample			
U.S.	61	53–63	−8 to +2
U.K.	30	31–35	+1 to +5
France	37	38–39	+1 to +2
Spain	31	31–33	+0 to +2
Finland	31	32–33	+1 to +2
Countries using geographic income to oversample			
Germany	46	51–54	+5 to +8
Belgium	31	33–34	+2 to +3
Countries using geographic information to oversample			
Austria	48	52–54	+4 to +6
Portugal	41	42–45	+1 to +4
Countries using no oversampling			
Italy	32	37–38	+5 to +6
Netherlands	26	27–36	+1 to +10

Source: Author's calculations based on the SCF, the HFCN, the WAS, and Forbes World's Billionaires.

Notes: Column (1) provides the top 5 percent share of wealth directly derived from the surveys. Column (2) provides the range of estimates when the tail is replaced by the estimated Pareto distribution the using sample and Forbes data. The Pareto distribution is estimated at thresholds of €500,000, €1 million, €2 million, €3 million, €5 million, and €10 million for the U.S., the U.K., France, and Spain, and at thresholds of €500,000, €1 million, and €2 million for other countries.

different methodology and dataset, that is, the capitalization of capital income tax data. A further major difference is that the SCF uses households, whereas Saez and Zucman (2014) use tax units.

Note that relative to the direct survey estimate, also in France and Spain, with heavy oversampling, the estimate using the Forbes data is not that different, adding only one to three percentage points. Obviously, such increases in the estimates are still non-negligible, but much smaller than the adjustments for the other countries. For the other countries without strong oversampling, Germany, Belgium, Austria, and Portugal, the survey estimate is also well below the regression estimate using the Forbes data, indicating the downward bias caused by differential unit non-response. Note that adjustments can be quite large but simultaneously not very sensitive to the chosen threshold of where the Pareto tail starts. For instance, for Germany, the top 1 percent of households hold 24 percent according to the direct estimate from the survey sample, but hold between 32 and 34 percent when replacing the survey sample tail observations by the estimated Pareto tail. According to this top tail measure, such a large adjustment indicates that German wealth is as unequally distributed as in the U.S., something which might have escaped attention if only the survey estimates of the German HFCS and SCF were compared. Indeed, a key lesson is that a comparison across countries of wealth inequality based on top wealth figures derived from surveys is a treacherous exercise. The data user might not be aware that a technical decision in the background, such as which oversampling method was used, has such large effects.

Table 10 shows the percentage wealth share of the top 5 percent of households.[23] Similarly here, the direct survey estimates are biased downward. For most surveys, including the Forbes data and replacing the survey sample tail by the estimated Pareto tail increases the percentage wealth share by multiple percentage points. Also here, the number calculated directly from the SCF, 61 percent, is within the bounds of the estimation with the Forbes data, 53–63 percent, when taking into account estimates at all thresholds. However, taking only account of the highest thresholds from €3 million onwards, the estimates range from 62 to 63 percent, a small 1–2 percent higher than the SCF survey estimate. Again, the adjustments are largest for countries that either do not oversample or that only use geographic income or geographic information to oversample the wealthy.

When interpreting the results, it has to be kept in mind that they are obtained under the implicit assumption that the survey responses and Forbes data are reasonably accurate. Forbes does not provide enough information to validate the data. However, the consensus seems to be that the numbers are reasonably accurate. A more serious concern is that respondents in surveys might underreport holdings and values of assets and liabilities. Underreporting in wealth surveys could lead to a different set of biases, as discussed above. For instance, if wealth in surveys is underreported, combining it with Forbes data might lead to overestimation of the degree of inequality in the tail.

The existence of underreporting problems in wealth surveys is demonstrated by comparisons of the aggregate wealth estimates obtained by household surveys with the wealth figures from the national Household Balance Sheet (HBS) (which is part of the system of national accounts). Those comparisons suggest that underreporting problems are unfortunately quite common. HFCS (2013) discusses in some detail a comparison of aggregate wealth estimates using the HFCS survey versus HBS estimates. The ratio of aggregate survey wealth on aggregate national wealth based on HBS ranges from 0.53 in the Netherlands to 0.94 in Belgium. Henriques and Hsu (2014) discuss a similar SCF–HBS comparison. They estimate that the aggregate wealth estimate of the 2010 SCF is actually 21 percent larger than the national wealth estimate based on the HBS. This would suggest an over-reporting problem in the SCF instead of the underreporting problem in the HFCS.

Both HFCS (2013) and Henriques and Hsu (2014) argue convincingly that a comparison of wealth survey data with national accounts HBS data is far from straightforward. There are serious comparability issues when comparing national accounts data and survey data. One important issue is that the target population of the surveys (households) is not identical to what is reported in the national

[23]Note that the top 1 and 5 percent wealth shares calculated using the Pareto tail at the “best-fit” threshold, available in Part V of the Appendix (in the Online Supporting Information), are almost always within the ranges given in Tables 9 and 10.

accounts. Namely, in national accounts, households are, in many countries, reported upon jointly with non-profit institutions serving households (also known as NPISHs, such as churches, labor unions, and so on). National accounts often do not provide a separate estimate of the wealth held by these non-profit institutions. This is, however, needed if one would want to compare the HBS with surveys: optimally, one would want to exclude the wealth of NPISHs from the HBS. Also, the definition, valuation, and recording dates of different items in national accounts and surveys are generally not identical. Correcting for these differences when comparing HBS and surveys is not trivial. Notwithstanding comparability issues, a comparison of carefully adjusted HBS aggregate numbers with aggregate wealth estimates of household surveys seems to be a fruitful avenue to investigate potential problems with the micro-data. Note that this takes the view that national account numbers are closer to the "truth" than survey numbers, which seems to be a most reasonable assumption.

Davies *et al*. (2014) describe a simple strategy to deal with the underreporting problem in household surveys, which they then use when estimating the global distribution of wealth (see Davies *et al*., 2014, subsections 1.7 and 3.2). Before grafting a Pareto tail onto survey data, the survey numbers are scaled up or down to ensure that the newly estimated aggregate wealth estimate (that is, survey plus Pareto tail) matches the HBS aggregate. Such a strategy seems reasonable; however, as indicated above, comparing household survey data with HBS data is far from trivial, so that it is unclear what scaling factor one should use. In other words, it is unclear if and by how much HBS data should be adjusted before one can construct such a scaling factor. Such a strategy also imposes the implicit assumption that underreporting is a uniform percentage for each household. There is certainly no guarantee that this is the case. Rather, if underreporting is more likely in financial assets such as stocks and bonds, this would imply that underreporting is more severe for the wealthy. If richer households underreport more percentage-wise, survey observations of these households should probably be scaled up more than the observations of poor households. Unfortunately, not much is known about the degree to which households differ in underreporting. A further analysis of the degree and distribution of underreporting among households remains an important avenue for further research.

5. Conclusion

The wealth distribution is an important variable for researchers, policy-makers, and society at large. Many analyses of redistribution or tax policy in general will be sensitive to the concentration of wealth at the top. Yet, our knowledge of the wealth distribution is less than perfect. This paper has investigated how differential unit non-response in household wealth surveys affects tail wealth estimates. The results are striking. Survey wealth estimates are very likely to underestimate wealth at the top, and this often by multiple percentage points. Countries that seem to have a more equal wealth distribution might not be so upon closer scrutiny.

This paper has investigated underestimation of wealth at the top in household surveys caused by differential unit non-response that is not remedied by appropriate reweighting of sample observations, because by its very nature the wealth of the non-respondents is unobserved. A first lesson learned is that survey estimates of top wealth are best seen as lower bounds. A second lesson is that the truncation at the top caused by differential unit non-response cannot be remedied by a simple interpolation of the survey by estimating a power law using survey data *only*. The presence of differential unit non-response leads to upward-biased tail index estimates and therefore too low total tail wealth estimates. Rather, a main result of this paper is that under the assumption of a true Pareto distribution for tail wealth, the Monte Carlo evidence shows that even very few extreme observations of wealth are sufficient to largely eliminate the serious downward bias in the Pareto tail index caused by differential unit non-response in wealth surveys, while substantially reducing the variance of the wealth estimates.

Therefore, rich lists such as the Forbes World's billionaires can help in dramatically improving top wealth estimates. This is not so much because of the wealth numbers of these billionaires themselves; rather, the combination of survey data and rich list leads to unbiased estimates of the Pareto tail index. Obviously, this is all true under the assumption that the extreme tail follows the same distribution as the wealthy just below. This need not be true. However, the fact that tail wealth estimates of surveys that do oversample the wealthy (such as the U.S., France, and Spain) all change relatively little when the surveys are combined with the Forbes data suggest that this assumption is a reasonable starting point.

Of course, as the evidence related to the SCF, and the French and Spanish HFCS, shows, improvement in terms of oversampling, combined with appropriate reweighting of the wealthy, will yield major benefits in terms of estimation of the tail of wealth. Ideally, wealth surveys should therefore follow this practice in *identifying the wealthy a priori*, thereafter heavily *oversampling* them and thereafter *adjusting the weights for differential unit non-response*. In that case, rich lists such as the Forbes World's billionaires would add little to the estimation of tail wealth.

In the meantime, however, researchers should be warned of top wealth estimates based on surveys alone, or on simple interpolations of the survey data, if there is evidence that differential unit non-response problems are serious and have not been completely addressed by readjustment of the survey weights and oversampling of the wealthy is limited. In those cases, combining survey data with data from rich lists could, at the minimum, provide a check of the robustness of the tail wealth estimates.

References

Aigner, D. J. and A. S. Goldberger, "Estimation of Pareto's Law from Grouped Observations," *Journal of the American Statistical Association*, 65, 712–23, 1970.

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

Avery, R. B., G. E. Elliehausen, and A. B. Kennickell, "Measuring Wealth with Survey Data: An Evaluation of the 1983 Survey of Consumer Finances," *Review of income and wealth*, 34, 339–69, 1988.

Brzezinski, M., "Do Wealth Distributions follow Power Laws? Evidence from Rich Lists," *Physica A: Statistical Mechanics and its Applications*, 406, 155–62, 2014.

Capehart, K. W., "Is the Wealth of the World's Billionaires Not Paretian?" *Physica A: Statistical Mechanics and its Applications*, 395, 255–60, 2014.

Clauset, A., C. R. Rohilla, and M. E. J. Newman, "Power-Law Distributions in Empirical Data," *SIAM Review*, 51, 661–703, 2009.
Davies, J. B. and A. Shorrocks, "The Distribution of Wealth," in A. B. Atkinson and F. Bourguignon(eds), *Handbook of Income Distribution: Volume 1, 605–75*, North-Holland, Amsterdam, 1999.
Davies, J. B., S. Sandström, A. Shorrocks, and E. N. Wolff, "The Level and Distribution of Global Household Wealth," *The Economic Journal*, 121, 223–54, 2010.
Davies, J. B., R. Lluberas, and A. Shorrocks, *Global Wealth Databook 2014*, Credit Suisse Research Institute, Zurich, 2014.
Dell, F., T. Piketty, and E. Saez, "Income and Wealth Concentration in Switzerland over the Twentieth Century" in A. B. Atkinson and T. Piketty (eds), *Top Incomes over the Twentieth Century*, Oxford University Press, Oxford, 2007.
HFCS (Eurosystem Household Finance and Consumption Network), "The Eurosystem Household Finance and Consumption Survey: Methodological Report for the First Wave," ECB Statistical Paper Series, 1, 2013.
Gabaix, X., "Power Laws in Economics and Finance," *Annual Review of Economics*, 1, 255–93, 2009.
Gabaix, X. and R. Ibragimov, "Rank $-1/2$: A Simple Way to Improve the OLS Estimation of Tail Exponents," *Journal of Business and Economic Statistics*, 29, 24–39, 2011.
Henriques, A. M. and J. W. Hsu, "Analysis of Wealth Using Micro- and Macrodata: A Comparison of the Survey of Consumer Finances and Flow of Funds Accounts," in D. W. Jorgenson, J. S. Landefeld, and P. Schreyer (eds), *Measuring Economic Sustainability and Progress*, 245–74, The University of Chicago Press, Chicago, 2014.
Kennickell, A., "Multiple Imputation in the Survey of Consumer Finances," Paper prepared for the August 1998 Joint Statistical Meetings, Dalles, TX, 1998.
———, "Wealth Measurement in the Survey of Consumer Finances: Methodology and Directions for Future Research," Paper prepared for the May 2000 Annual Meetings of the American Association for Public Opinion Research, Portland, OR, 2000.
———, "The Role of Over-Sampling of the Wealthy in the Survey of Consumer Finances," Federal Reserve Board Survey of Consumer Finances Working Papers, 2007.
———, "Getting to the Top: Reaching Wealthy Respondents in the SCF," Paper presented at the 2009 Joint Statistical Meetings, Washington, DC, 2009a.
———, "Ponds and Streams: Wealth and Income in the U.S., 1989 to 2007," Federal Reserve Board Finance and Economics Discussion Series, 13, 2009b.
Kennickell, A.and D. McManus, "Sampling for Household Financial Characteristics Using Frame Information on Past Income," Paper presented at the 1993 Joint Statistical Meetings, Atlanta, GA, 1993.
Kennickell, A.and R. L. Woodburn, "Consistent Weight Design for the 1989, 1992, and 1995 SCFs, and the Distribution of Wealth," Federal Reserve Board Survey of Consumer Finances Working Papers, 1997.
Klass, O. S., O. Biham, M. Levy, O. Malcai and S. Solomon, "The Forbes 400 and the Pareto Wealth Distribution," *Economics Letters*, 90, 290–5, 2006.
Kopczuk, W. and E. Saez, "Top Wealth Shares in the United States, 1916–2000: Evidence from Estate Tax Returns," *National Tax Journal*, 75, 445–88, 2004.
Levy, M. and S. Solomon, "New Evidence for the Power-Law Distribution of Wealth," *Physica A: Statistical Mechanics and its Applications*, 242, 90–4, 1997.
ONS (Office for National Statistics), *Wealth and Assets Survey Review Report*, ONS, London, 2012.
Ogwang, T., "Power Laws in Top Wealth Distributions: Evidence from Canada," *Empirical Economics*, 41, 473–86, 2011.
Piketty, T., *Capital in the Twenty-First Century*, Harvard University Press, Cambridge, MA, 2014.
Piketty, T., G. Postel-Vinay, and J. L. Rosenthal, "Wealth Concentration in a Developing Economy: Paris and France, 1807–1994," *American Economic Review*, 96, 236–56, 2006.
Roine, J. and D. Waldenström, "Long Run Trends in the Distribution of Income and Wealth," in A. B. Atkinson and F. Bourguignon (eds), *Handbook of Income Distribution: Volume 2A*, 469–592, North-Holland, Amsterdam, 2014.
———, "Wealth Concentration over the Path of Development: Sweden, 1873–2006,"*Scandinavian Journal of Economics*, 111, 151–87, 2009.
Rytgaard, M., "Estimation in the Pareto Distribution," *Astin Bulletin*, 20, 201–16, 1990.
Saez, E. and G. Zucman, "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," NBER Working Paper, 20625, 2014.
von Kalckreuth, U., J. Le Blanc, T. Schmidt, and J. Zhu, "The PHF: A Comprehensive Panel Survey on Household Finances and Wealth in Germany," Deutsche Bundesbank Discussion Paper, 13, 2012.

Wolff, E. N., "Methodological Issues in the Estimation of the Size Distribution of Household Wealth," *Journal of Econometrics*, 43, 179–95, 1990.
———, "Changes in Household Wealth in the 1980s and 1990s in the U.S.," in E. N. Wolff (ed.), *International Perspectives on Household Wealth*, Edward Elgar, Cheltenham, 2006.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publishers web-site:

Appendix
Part I: Pareto tail index and wealth shares
Part II: Tail of the wealth distribution
Part III: Derivation log-rank log-size regression
Part IV: Monte Carlo results
Part V: Best-fit thresholds