HOW FAT IS THE TOP TAIL
OF THE WEALTH DISTRIBUTION?

Philip Vermeulen

**HOUSEHOLD FINANCE AND
CONSUMPTION NETWORK**

In 2014 all ECB
publications
feature a motif
taken from
the €20 banknote.

Household Finance and Consumption Network

This paper contains research conducted within the Household Finance and Consumption Network (HFCN). The HFCN consists of survey specialists, statisticians and economists from the ECB, the national central banks of the Eurosystem and a number of national statistical institutes.

The HFCN is chaired by Gabriel Fagan (ECB) and Carlos Sánchez Muñoz (ECB). Michael Haliassos (Goethe University Frankfurt), Tullio Jappelli (University of Naples Federico II), Arthur Kennickell (Federal Reserve Board) and Peter Tufano (University of Oxford) act as external consultants, and Sébastien Pérez Duarte (ECB) and Jiri Slacalek (ECB) as Secretaries.

The HFCN collects household-level data on households' finances and consumption in the euro area through a harmonised survey. The HFCN aims at studying in depth the micro-level structural information on euro area households' assets and liabilities. The objectives of the network are:

1) understanding economic behaviour of individual households, developments in aggregate variables and the interactions between the two;

2) evaluating the impact of shocks, policies and institutional changes on household portfolios and other variables;

3) understanding the implications of heterogeneity for aggregate variables;

4) estimating choices of different households and their reaction to economic shocks;

5) building and calibrating realistic economic models incorporating heterogeneous agents;

6) gaining insights into issues such as monetary policy transmission and financial stability.

The refereeing process of this paper has been co-ordinated by a team composed of Gabriel Fagan (ECB), Pirmin Fessler (Oesterreichische Nationalbank), Michalis Haliassos (Goethe University Frankfurt), Tullio Jappelli (University of Naples Federico II), Sébastien PérezDuarte (ECB), Jiri Slacalek (ECB), Federica Teppa (De Nederlandsche Bank), Peter Tufano (Oxford University) and Philip Vermeulen (ECB).

The paper is released in order to make the results of HFCN research generally available, in preliminary form, to encourage comments and suggestions prior to final publication. The views expressed in the paper are the author's own and do not necessarily reflect those of the ESCB.

Philip Vermeulen
European Central Bank; e-mail: philip.vermeulen@ecb.europa.eu

Abstract

The US Survey of Consumer Finances (SCF) and the Eurosystem's Household Finance and Consumption Survey (HFCS) provide evidence that wealth is heavily concentrated at the upper tail of the wealth distribution. A commonly cited number for the US is that 1 percent of the households hold 30 percent of total household wealth. I investigate the reliability of upper tail wealth estimates from household wealth surveys in the presence of survey differential non-response, i.e. the fact that richer households have lower response rates than poorer households. Differential non-response can often not be remedied by adjustment of survey weights, as wealth of the non-responding households remains unobserved. Differential non-response biases tail wealth estimates downwards. Monte Carlo evidence shows that such a bias can be quite substantial. I provide a method that greatly reduces the bias. The method combines survey data with data from rich lists and uses them jointly to estimate a Pareto (power-law) distribution for tail wealth. Using this method, the paper combines the SCF and HFCS data with Forbes World's billionaires data to provide new estimates of tail wealth. For surveys with low or no oversampling of the wealthy, these estimates tend to indicate a higher concentration of wealth at the top than those calculated from the wealth surveys alone.

Key words: wealth distribution; Survey of Consumer Finances ; Household Finance and Consumption Survey ; power law

JEL:D31

Non-technical Summary

One very well established fact about the distribution of wealth, one that is often recurring in the public debate, is its substantial positive skew. That is, only a small percentage of the population hold a considerable fraction of total wealth. For instance, in the US, the one percent richest households hold approximately 30 percent of total household wealth; the five percent richest households hold approximately 60 percent of wealth. Good measurement of these relative shares of wealth holdings are important. For instance, macro-economic models are sometimes calibrated using these shares. In addition, the wealth distribution is a topic of broad public interest.

Up to recently, relatively little was known about the wealth distribution of households for most countries in the euro area. That should not be too surprising. Knowledge of the wealth distribution of households is only possible if it is somehow measured. One way to measure wealth of individual households (and thereafter the distribution of the population of households) is to use extensive household surveys that ask questions of all asset holdings and debts of the household. In the US, the Survey of Consumer finances (SCF), sponsored by the Board of Governors of the Federal Reserve System, is widely conceived to be the best source of information on the US wealth distribution. This household survey has a long standing and much of what we know today about the US wealth distribution comes from this survey. In the euro area a new survey, similar to the SCF, the Household Finance and Consumption Survey (HFCS), from a joint project between the European Central bank, the Eurosystem and a number of national statistical institutes, has been recently released (April 2013). Both the SCF and HFCS are large surveys specifically designed to represent the entire population and to measure wealth at the household level.

This paper is concerned about measuring the share of wealth held by the top one and five percent of households. It starts with the premise that measuring wealth at the top is always difficult. Household surveys are widely believed to suffer from various degrees of non-response and differential non-response. Sampled households don't participate in surveys for numerous reasons: absence, lack of time, refusal to reveal sensitive information, etc. When these reasons are correlated with wealth itself serious biases can occur in wealth estimation and the shares of wealth can be biased too. Wealthier households are generally thought to have higher non-response rates. If non-responding households are having higher wealth in some systematic way, wealth estimates will be biased downwards, particularly estimates of wealth at the top of the distribution .

One way that survey analysts try to remedy this potential problem is by stratifying the sample and oversampling of the wealthy. This is usually done by drawing the rich housholds from a special sampling frame (e.g. tax records). This way more wealthy households will end up in the sample, and their sample weights can be better adjusted for non-response. This strategy goes a long way in minimizing the downward bias of wealth estimates at the top of the distribution.

However not all wealth surveys oversample the rich and not all oversampling strategies

are equally successful in obtaining rich households in the sample. This paper proposes a method to overcomes this bias, which is especially useful when the rich households are not or little oversampled. The sample data of the HFCS and the SCF are combined with Forbes World's billionaires data. Estimates of wealth at the top can then be obtained by estimating a Pareto distribution for the top wealth holders, as earlier research has shown that the Pareto distribution forms a good approximation of wealth held at the top.

In the paper it is then first shown that combining these sources of data when estimating such a distribution can improve on the estimates of the share of wealth held at the top of the distribution. Thereafter, the paper compares estimates of the upper tail of the wealth distribution using different methods. Direct estimates of tail wealth from the sample as well as estimates from an estimated Pareto distribution are compared. The Pareto distribution is estimated with and without adding Forbes individuals at the extreme tail.

This paper provides a set of estimates on the tail of the wealth distribution in the US and nine euro area countries. The results suggest that differential non-response problems, are particularly high in a number of euro area countries, leading to underestimation of the top wealth shares when using only the surveys to construct tail wealth or using estimates from a Pareto tail without extreme tail observations. When using the extra information provide by the Forbes World's billionaires list, estimates of tail wealth, and shares of the top one percent and five percent increase.

1 Introduction

One well established fact about the wealth distribution is its substantial positive skew. Davies and Shorrocks (1999) call it a stylized fact. Only a small part of the population hold a large fraction of wealth. For instance, one percent of the U.S. households at the top of the wealth distribution hold around one third of the total household wealth (see Wold, 2006 and Kennickell, 2007). While in other countries the share of the top one percent might be smaller than in the US, they are still believed to be considerable (Davies et al. 2010).

Policy makers and macroeconomic researchers are both interested in accurate measurement of the wealth distribution, and especially its top tail. While policy makers might be most interested in the issue of redistribution, macroeconomic research is more focused on the causes and consequences of the large heterogeneity in income and wealth. In an attempt to understand the sources and effects of income and wealth inequality, a growing macro-economics literature calibrates models to wealth distribution data. Quite some advances have been made in the development of models that match especially the upper tail of the wealth distribution.[1] Having trustworthy estimates of the wealth distribution, and in particular its upper tail, is therefore important. However, the skew of the wealth distribution makes accurate measurement of top tail wealth particularly challenging.

The difficulty rests in the fact that much of our knowledge of the wealth distribution is derived from household surveys. The Survey of Consumer finances (SCF), sponsored by the Board of Governors of the Federal Reserve System, is widely conceived to be the best source of information on the US wealth distribution. Up to recently, information on the wealth distribution in different European countries was relatively scarce.[2] The situation has changed, as in the euro area a new survey, similar to the SCF, the Household Finance and Consumption Survey (HFCS), from a joint project between the European Central bank, the Eurosystem and a number of national statistical institutes, has been recently released (April 2013).

Surveys such as the SCF and the HFCS serve multiple purposes (such as understanding asset portfolio choices, indebtedness patterns etc.) so that it should be clear at the outset that providing a precise estimate of top tail wealth might not necessarily be a goal of these surveys. Estimates of median wealth or other percentiles of the wealth distribution might be deemed more important goals. However, these surveys share a number of features that make them particularly useful to also analyze the top tail of the wealth distribution. Both the SCF and the HFCS are designed to capture the entire household population and to provide a complete picture of wealth of the households. Therefore, not only do the SCF and HFCS allow, in principle, measurement of the top tail of the household wealth distribution, they are probably the single most important source to do so. Few

---

[1]See e.g. Benhabib et al. (2011), Cagetti and DeNardi (2008), and Castañeda et al (2003) and references therein.

[2]There is much more information on the income distribution than on the wealth distribution.

other surveys or data sources come to mind that have the necessary household wealth information and have a scope as wide as the SCF and the HFCS.

However, measuring wealth at the top is always difficult with household surveys, as these are widely believed to suffer from various degrees of non-response and differential non-response. Survey non-response is the non-participation of a sampled household in the survey, whereas differential non-response refers to non-response that differs across various groups in the population.[3] Sampled households don't participate in surveys for numerous reasons: absence, lack of time, refusal to reveal sensitive information, etc. When these reasons are correlated with wealth itself serious biases can occur in wealth estimation.

A bias will occur in wealth estimates if the non-participating households are not entirely randomly distributed. The bias will depend on many factors, but particularly important in the context of wealth distribution estimation is when non-response bias correlates with household wealth. Such a correlation creates effectively a *differential* non-response in the survey population, with wealthier households having higher non-response rates. If non-responding households are having higher wealth in some systematic way, wealth estimates will be biased downwards, particularly estimates of tail wealth.[4]

The purpose of this paper is threefold. First, I provide new insights on the importance of non-response and in particular the *differential* non-response of the wealthy in the SCF and HFCS. The main emphasis is on how differential non-response influences the accuracy of estimates of the top of the wealth distribution. Second, I propose a method to lessen the effect of differential non-response on the estimates of the tail of the wealth distribution. Third, I provide new estimates of the share of wealth held by the top one and top five percent households.

With respect to the first purpose, I begin by documenting that the SCF and the HFCS wealth surveys suffer to a different degree from (differential) non-response at the tail of the wealth distribution. There are systematically "missing rich" in all those surveys. Where the very rich, that is billionaires, are missing from all surveys, some of the HFCS surveys in particular, suffer from various degrees of "missing rich", which are significantly larger than the SCF. The HFCS surveys differ substantially across countries in the methods used to oversample the rich. Across countries, there is a positive correlation between the degree to which the "rich" are missing from the sample and the method used to oversample the rich.

Using the insights of the power law literature, I address the problem of the "missing rich" for wealth distribution estimation. The literature on the wealth distribution seems to have converged on the idea that the top of the wealth distribution can be well described

---

[3]Not to be confused with item non-response, which is the absence of an answer to a particular question. Item non-response in the SCF and HFCS are dealt with using multiple imputation techniques.

[4]Another source of potential bias is underreporting of assets for the participating households. To the extent that underreporting is homogeneous across the population, the share of wealth of the tail should be little affected. When underreporting is positively correlated with wealth, wealth shares of the top would be biased downwards. Effectively, there is relatively little detailed information about underreporting and differential underreporting.

by a Pareto distribution (Davies and Shorrocks, 1999). While estimates of the top of the wealth distribution can be obtained from the survey sample directly, estimates of wealth at the top can also be obtained by estimating a Pareto distribution for the top wealth holders of a survey sample. However, a detailed investigation of how the wealth estimates of the tail of the distribution using different methods are affected by the presence of differential non-response is missing from the literature. This paper attempts to fill that gap. I show, using Monte Carlo simulation, that the wealth in the tail can be seriously underestimated under reasonable differential non-response assumptions. Underestimation of tail wealth is true for both direct estimates of tail wealth from the sample as well as estimates from an estimated Pareto distribution. I then show that adding observations of individuals at the extreme tail, even if only a few observations are available, can improve estimation of the Pareto distribution dramatically. Underestimation practically disappears.

I use these insights to compare estimates of the upper tail of the wealth distribution using the different methods. Direct estimates of tail wealth from the sample as well as estimates from an estimated Pareto distribution are compared. The Pareto distribution is estimated with and without adding individuals at the extreme tail. Adding observations of individuals at the extreme tail is done by adding the Forbes World's billionaires list to the SCF and HFCS observations. This provides a set of estimates on the tail of the wealth distribution in the US and nine euro area countries. The results suggest that differential non-response problems are particularly high in a number of euro area countries, leading to underestimation of the top wealth shares when using only the surveys to construct tail wealth or using estimates from a Pareto tail without extreme tail observations. When using the extra information provide by the Forbes World's billionaires list, estimates of tail wealth, and shares of the top one percent, increase substantially. Estimates of the shares of the top five percent also increase, but less so.

The remainder of the paper is structured as follows. Section 2 describes the data used, the SCF, 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. It also contains a Monte Carlo study, illustrating that information from rich lists can improve Pareto estimates in the presence of differential non-response. Section 4 provides new estimates of the share of wealth held by the top one and five percent households. Section 5 concludes.

2 The data

2.1 The US SCF and the Eurosystem HFCS

This paper combines the 2010 wave of the US Survey of Consumer finances (SCF), the first wave of the European HFCS, and the Forbes World's billionaires list to estimate wealth at the upper tail of the distribution. The SCF is a triannual survey of US household wealth,

sponsored by the Board of Governors of the Federal Reserve System. It provides the most comprehensive source of wealth information of US households. The HFCS, building very much on the SCF in terms of methodology, provides detailed information on household assets and liabilities of individual households in fifteen euro area countries. In total, there are more than 62000 households in the dataset.

I use the HFCS data for Germany, France, Italy, Spain, Belgium, Portugal, Austria and Finland. I drop Greece, Cyprus, 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.[5]

Both the SCF 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. Importantly, by design, sample inclusion probabilities are different for different households. To account for that, and other features of the complex survey design, survey weights are provided for each sample observation, so that, in principle, an unbiased representation of the survey population can be obtained. Each sample weight signifies the number of households in the population that the sample point represents. The total sum of weights for each country therefore is equal to the total number of households in the population.

Usually in survey settings some participating households leave some questions unanswered (but provide an answer to the bulk of the questions). Excluding the survey responses of those households (and seeking replacement households) simply on the basis of a few unanswered questions is generally considered too costly or impractical. It is customary to deal with missing observations using multiple imputation. This is also the case for the SCF and the HFCS. For each household observation there are five implicates. 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. A more detailed description of the SCF and HFCS methodologies can be found in Kennickell (2000) and HFCN (2013). For comparison purposes, the US SCF data are converted into euro using the dollar/euro exchange rate of 12 feb 2010, 1.3572.

---

[5]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, sight 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 labour income.

2.2 Oversampling the wealthy and non-response

A problem that any survey of wealth faces is that wealth is concentrated at the top tail of the distribution. Using simple random sampling, it would be close to impossible to have a good representation of that top, unless if the sample was very large. For that reason, 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 to be on average wealthier. Oversampling of the wealthy is the case for the SCF and for some, but not all, country surveys in the HFCS. Oversampling of the wealthy serves two purposes. On the one hand, for efficient estimates of wealth held in the tail of the wealth distribution one needs enough observations of wealthy households. One the other hand, non-response, particularly of the wealthy, is a serious problem and can be partially addressed by oversampling.

Differential non-response in a wealth survey is a serious issue. Wealth estimates can be biased, especially if non-response is correlated[6] with wealth. The very long tail of the wealth distribution, with a small fraction of the population holding a large fraction of the wealth, exacerbates the problem. Note that differential non-response cannot be addressed by simply increasing the sample. Imagine non-responding households being replaced by newly drawn households of the population. Although it avoids a reduction in sample size, it does not address the fundamental problem of non-response correlated with wealth. In case of such correlation, a rich non-responding household is therefore more likely to be replaced with a poorer responding household.

Non-response, both of the general type and differential one, generally leads to a re-adjusting of the weights given to the sample observations. The respondent households weights are scaled up, making up for the non-responding households. However, differential non-response is only addressed effectively if for non-responding wealthy households the other wealthy responding households get higher weights. In other words, re-adjustment of the weights has to be selective. Selective readjustment is only possible if the survey is designed in such a way as to provide (at the minimum some partial) knowledge of the wealth of the non-responding household, so that the weights of particular responding households can be scaled up. For instance, if wealthy households are being selected from a special sampling frame, maybe based on income tax or wealth tax data, such a correction can be done. However, the reweighing will only be as good as the ability to identify the non-responding household as being wealthy.

The case of the SCF is interesting in that regard. The SCF uses a dual frame to sample households. On the one hand, a representative area probability sample is drawn. On the other hand, a high-income sample is drawn using Federal tax returns to construct the

---

[6]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).

sample frame. The high-income sampling frame allows to construct different strata, with higher strata having higher income (and higher expected wealth) and higher oversampling rates. Details are provided in Kennickell (2007). The different strata from the high-income frame allows to address non-response problems in a selective way, as described above.

Outside of the SCF, relatively little is known in practice of the degree in which non-response is correlated with wealth. The SCF provides the most evidence on this issue. Evidence presented in Kennickell and Woodburn (1997) shows that when using the high-income frame to construct a wealth index (essentially an estimate of wealth based on income information), and thereafter sorts sampled individuals into bins of the wealth index, the bin of 1 million to 2,5 million dollar has a response rate of 34 percent, whereas the bin of 100 million to 250 million has a response rate of 14 percent. Such information is used to adjust the weights of the respondent households.

As Kennickell (2007) observes *In the stratum of the SCF list sample that contains the respondents likely to be the wealthiest, the overall response rate is only 10 percent. The survey has often been critized for this low cooperation rate. Regrettable as this rate is, the fact that it is known is actually a strength of the survey. Presumably, other surveys also have a similar problem, but without some means of identifying it, they will fail to correct for an important source of bias in the estimation of wealth. In the SCF the original frame data for the list sample provides a rich basis to use for adjusting the sampling weights to compensate for nonresponse.* In the case of the SCF, oversampling of the wealthy implies the ability to also adjust the weights for non-response, as the wealthy are drawn from a special frame.

However, the degree and methods of oversampling of the wealthy differ dramatically across surveys, and therefore also the possibility to adjust selectively the weights for non-response. Arguably, having wealth tax data to identify 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. The HFCS differs across countries, both in the degree of oversampling of the rich households, with a few countries not doing oversampling at all, and the method to identify the wealthy.

Table 1 provides an overview of the different methods used to oversample the wealthy. Clearly, the construction of sampling strata that use individual information such as taxable wealth or income are likely to, not only effectively sample more out of the tail of the distribution, but also to provide means to selectively re-adjust weights in the presence of differential nonresponse. Oversampling using information at the individual level of wealth or income is done in the US, Spain, France and Finland. When individual level information is not available, oversampling can be done using regional income information. Regional income information is used in Germany and Belgium. When such information is not available, one can oversample simply based on regional criteria, with the idea that

the wealthy are most likely to live in the capital or in large cities. This is the strategy used in Austria and Portugal. Finally, in the Netherlands and Italy no oversampling is done.

TABLE 1

Oversampling method in SCF and HFCS

Using individual information	
USA	list based on income tax information
Spain	list based on taxable wealth information
France	list based on taxable wealth information
Finland	income information from register
Using geographic income information	
Belgium	average regional income
Germany	taxable income of regions
Using geographic information	
Austria	Vienna oversampled
Portugal	Lisbon and Porto oversampled
No oversampling	
Netherlands	No oversampling
Italy	No oversampling

Source: Own construction based on Kennickell (2009) and HFCN (2013).

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 euro, 1 million euro, and 500 thousand euro. 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 euro. This is not just a reflection of the presence of higher wealth in the US, but rather is indicative of the very high degree of oversampling in the SCF. In Spain and France, two other countries using individual information to oversample the wealthy, respectively 9 and 4 percent of the sample are households with wealth above 2 million euro. The two countries using geographic income information, Belgium and Germany, have respectively 3 and 2 percent of the sample with wealth above 2 million euro. Clearly, 2 to 3 percent are much smaller percentages compared to 4 to 9. The countries for which only geographic information is used, Portugal and Austria, only have respectively 2 and 1 percent of the sample in the highest wealth category. The case of no-oversampling, Italy and the Netherlands, have respectively a small 1 and 0 percent. 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.

TABLE 2

Summary statistics

Number of wealthy households in the survey samples

	Sample size	Absolute number			Pct of sample		
		> 2 million	> 1 million	> 500TH	> 2 million	> 1 million	> 500TH
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
USA	6482	965	1259	1692	0.15	0.19	0.26
Germany	3565	85	246	654	0.02	0.07	0.18
France	15006	638	1712	3522	0.04	0.11	0.23
Italy	7951	78	300	1075	0.01	0.04	0.14
Spain	6197	544	1129	2086	0.09	0.18	0.34
Netherlands	1301	2	32	172	0.00	0.02	0.13
Belgium	2327	71	207	599	0.03	0.09	0.26
Austria	2380	47	113	271	0.02	0.05	0.11
Finland	10989	59	296	1233	0.01	0.03	0.11
Portugal	4404	24	87	252	0.01	0.02	0.06

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 households that those observations in the tail represent (i.e. their weight). For instance, for the category above 2 million euro, Spain has 544 sample observations (Table 2) representing 139539 households (Table 3). Whereas Germany has a sample of 85, representing almost three times as much households. The Netherlands, with no oversampling, only has 2 households in the sample above 2 million euro. One immediately observes how efficiency of tail estimation will dramatically be effected by the different degree of oversampling.

TABLE 3

Summary statistics

Number of wealthy households in the population

(estimates derived from the survey samples)

	households	Absolute number			Pct of population		
		HH > 2 million	HH > 1 million	HH > 500TH	> 2 million	> 1 million	> 500TH
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
USA	117609217	3661191	8407106	15311762	0.03	0.07	0.13
Germany	39673000	368693	1051250	3261600	0.01	0.03	0.08
France	27860408	209668	830661	2891897	0.01	0.03	0.10
Italy	23817962	265782	901176	3100288	0.01	0.04	0.13
Spain	17017706	139539	621067	2299825	0.01	0.04	0.14
Netherlands	7386144	2895	83813	508482	0.00	0.01	0.07
Belgium	4692601	85386	264728	890283	0.02	0.06	0.19
Austria	3773956	70939	174550	427248	0.02	0.05	0.11
Finland	2531500	6555	34632	158436	0.00	0.01	0.06
Portugal	3932010	14141	64443	185746	0.00	0.02	0.05

2.3 Forbes data

Journalists lists of wealthy individuals is another source of information on the wealth of the very top of the distribution. The SCF 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, 2009). One notorious list is the annual Forbes World's billionaires list, which is measured in US dollar. An individual is on the Forbes billionaire list if his or her wealth is estimated to be above 1 billion dollar. From all existing rich lists, the Forbes billionaire list seems to be the best researched. Not only does Forbes have a long tradition, and therefore experience, in constructing such a list, but also some individual billionaires seem to co-operate in the construction of it. The methodology is explained in more detail on the Forbes website. Ultimately, of course, it has to be kept in mind that the Forbes wealth figures are also estimates. For the purpose of this paper, the wealth of individuals on the list is recalculated in euro.[7]

Table 4 provides an overview of the number of individuals on the Forbes billionaires list, the total wealth they have, and their wealth as a percentage of total household wealth of the country. Note that the SCF and HFCS surveys differ slightly with respect to the reference years, which range depending on the country from 2009 to 2011. Therefore, I match the year of the survey with the year of the Forbes list. As the largest country, the US has the most individuals on the list, with Germany and Italy second and third. Note that the individuals on the Forbes list can add significant information on the tail. For instance, the HFCS survey in Germany only has 85 individuals with wealth above 2 million euro, whereas there are 52 individuals on the Forbes billionaires list. For Italy, these numbers are is 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 euro, namely only two.

---

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

TABLE 4
The Forbes billionaires list
Number of people and Wealth

	Number of individuals	Total wealth	Percentage of country wealth
USA	396	978.6	2.3
Germany	52	183.3	2.4
France	11	60.1	0.9
Italy	14	46.6	0.7
Spain	12	28.3	0.6
Netherlands	3	4.8	0.4
Belgium	1	1.9	0.1
Austria	5	13.0	1.2
Finland	1	1.0	0.2
Portugal	2	4.1	0.7

Source: own calculations based on Forbes, HFCS and SCF. Total wealth in billion euro.

In principle, the HFCS covers all households resident in the country, thus also potentially the individuals on the Forbes billionaires list. In practice, Forbes billionaires are obviously not covered. Table 5 compares the maximum wealth found in the SCF and HFCS with the minimum wealth of a person on the Forbes Word's billionaires list. Starting with the SCF, there are sample observations that have higher wealth than the "poorest" Forbes billionaire. The very high oversampling rate of the wealthy in the SCF clearly is very effective. Contrary to the SCF, there is a serious gap between the richest household in the HFCS and the poorest person on the Forbes list. Such a gap can be found in all euro area countries. So the first observation is that none of the households in the HFCS 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 HFCS household is very large. Households that fall in between the richest HFCS household and the poorest Forbes billionaire are not in the sample.[8] Note that in Spain, whose survey among the other HFCS ones arguable does the best job in oversampling the rich (using tax reports), the maximum wealth in the HFCS is 401 million euro, whereas it is only 76 in Germany. Indeed in Spain the gap is the lowest, but it is still significant between 401 million and 769 million.

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, two countries where oversampling

---

[8]One possibility one could entertain is that these households don't exist. However, that thought seems clearly absurd. Such a reasoning would imply that in Germany there would be no household that has a net wealth between 76 million, the wealthiest household in the HFCS sample, and 818 million, the poorest German dollar billionaire.

is done based on individual tax records of wealth. The Netherlands, with no oversampling, has a rather low value of the maximum of wealth, namely 5 million euro. The other country with no oversampling, Italy, also has a low maximum value of wealth (26 million euro). Also, using only geographic information, which is the case of Portugal and Austria, or geographic income information, the cases of Belgium and Germany, does not guarantee to observe a high maximum of wealth.

The conclusion is clearly that the very rich households are not in the HFCS sample because of a combination of non-response and lack of effective oversampling, with the effectiveness greatly varying across countries. The few wealthy households at 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 done to oversample, such as in the SCF, and France and Spain for the HFCS, one can observe larger maximum of wealth.

TABLE 5
The GAP: Maximum wealth vs minimum at Forbes
Million euros

	Maximum wealth SCF/HFCS	Minimum wealth Forbes
US	806	737
Germany	76	818
France	153	810
Italy	26	893
Spain	409	780
Netherlands	5	958
Belgium	8	1920
Portugal	27	1110
Austria	22	1560
Finland	15	958

Source: own calculations based on Forbes World's Billionaires, SCF and HFCS.

Maximum is over all five implicates.

3   A Pareto law for the tail of the wealth distribution

3.1   The Pareto distribution

Davies and Shorrocks (1999) call 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'*. This last feature has been confirmed by a number of studies on the wealth distribution, using different countries and episodes. Some recent evidence is provided by Ogwang (2011), who estimates power laws for the 100 wealthiest Canadians for the

years 1999-2008, Levy and Solomon (1997), who estimate a Pareto law for the Forbes 400 wealthiest people in the US for the year 1996, and Klass et al. (2006), who estimate pareto laws also using the Forbes 400 in the US for the period 1988-2003.

The Pareto distribution has the following complementary cumulative distribution function (ccdf)[9]:

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

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 tail index, determines the fatness of the tail. The lower α, the fatter 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}.[10] 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)[11] power law if it is distributed according to the following relationship:

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

Simply stated, this relationship implies that the fraction (or empirical relative frequency) of households with wealth at or above w_i follows the regularity of a power function. The power law essentially mimics (1), where the probability $P(W > w)$ is replaced by the empirical relative frequency $\frac{N(w_i)}{N}$. In other words, a population is said to follow a power law if the empirical ccdf of the population is well approximated by the ccdf of a Pareto distribution. In such a finite population, a Pareto distribution will be the natural continuous approximation of the discrete distribution of wealth. Alternatively, one can say that a power law in a finite population is the likely outcome of a process where each household i in this population has drawn its wealth w_i from a Pareto distribution with parameters w_{min} and α.

The mean of the Pareto distribution with tail index $\alpha(> 1)$ and lower bound w_{min} is given by $w_{min}\frac{\alpha}{\alpha-1}$, so that total wealth of the population has an expected value of

$$N * w_{min}\frac{\alpha}{\alpha - 1}, \tag{3}$$

---

[9]In line with the literature, when discussing the Pareto distribution, it is much easier to use the ccdf than to use the cdf.

[10]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, i.e the N richest households.

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

which is the expectation of the sum of N i.i.d draws of a Pareto distribution.

3.2   Estimation of the Power law

3.2.1   Estimation on simple random samples versus samples from complex survey designs

There exists a large literature on the estimation of power laws, so that it suffices to be brief. For detail on different methods, the interested reader is referred to Gabaix (2009) and Clauset et al. (2009). However, there are a number of particularities to the estimation of power laws on samples from complex survey designs that have not received a lot of attention in the literature. Those will be emphasized in what follows.

The density of the Pareto distribution is given by:

$$f(w) = \frac{\alpha w_{min}^{\alpha}}{w^{\alpha+1}}, \tag{4}$$

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

$$\tilde{\alpha_{ml}} = [\sum_{i=1}^{n} \frac{1}{n} \ln(\frac{w_i}{w_{min}})]^{-1} \tag{5}$$

Without some adjustment, the maximum likelihood estimator should not be used on complex survey data. The sampling observations of the SCF and 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 (both SCF and HFCS only provide weights, but not the exact sampling detail to preserve confidentiality) a true likelihood cannot be constructed. Due to stratification and clustering and possible oversampling some observations will have a much higher likelihood to occur in the sample than others. Using a maximum likelihood estimator on such samples would clearly lead to onerous results.

Remember that in the SCF and the HFCS the 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 $w_1, w_2, w_3, ...$ 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

$$\tilde{\alpha_{pml}} = [\sum_{i=1}^{n} \frac{N_i}{N} \ln(\frac{w_i}{w_n})]^{-1} \tag{6}$$

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 also has given rise heuristically to an alternative estimation method in simple random samples. Imagine again a population of N household that follows a power law as in (2). Assume that a simple random sample $\{w_i, i = 1, ...n\}$ is drawn from the population. Denote by $n(w_i)$ 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 relative frequency in the sample provides an estimate of the relative frequency in the population, i.e.:

$$\frac{n(w_i)}{n} \cong \frac{N(w_i)}{N}, \forall w_i \tag{7}$$

As the sample gets larger, the estimate will obviously become closer to the true population frequency. Combining this with the power law relationship in the population we get

$$\frac{n(w_i)}{n} \cong (\frac{w_{min}}{w_i})^{\alpha}, \forall w_i \tag{8}$$

Taking logs on both sides, we have that the log of the relative frequency (or empirical ccdf) in the sample is a downward sloping linear function of the log of the observation scaled by the threshold w_{min}.

$$\ln(n(w_i)/n) = -\alpha \ln(w_i/w_{min}) \tag{9}$$

In the literature, one can find a number of variants of equation (9). For instance, α can be estimated using a linear regression of the log of the rank on the log of the observation, also known as log rank-log size relationship.

$$\ln(n(w_i)) = C - \alpha \ln(w_i), \tag{10}$$

with $C = \ln(n) + \alpha \ln(w_{min})$, where both α and C are estimated. However, if w_{min} is

known, or somehow set beforehand, (9) can be estimated without constant term. Estimation of (9) or (10) is called the regression method in estimating a power law.

In a complex survey sample, again the survey weights have to be taking into account to construct the empirical ccdf. Taking into account survey weights can be done the following way. Recall that a survey weight represents the number of households that a sample point represents. Imagine a survey sample from a complex survey design. First, rank the sample households according to wealth. That is, the wealthiest household has wealth w_1 and a survey weight of N_1, and the second wealthiest household has wealth w_2 and survey weight of N_2, and so on. The relative frequency represented by the first household is $\frac{N_1}{N}$, by the second household it is $\frac{N_1+N_2}{N}$, and so on. So that for a complex survey sample in (9) and (10) the rank $n(w_i)$ can be replaced by the sum of all survey weights of sample observations with wealth at least as large as w_i, which is $N_1+N_2+...N_i$. The sample size n is replaced by the population size N. The population size is then the sum of all survey weights of sample points with wealth above w_{min}.

3.2.2 Combining survey with Forbes data

As discussed above, both the SCF 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. To estimate a power law on the pooled survey-Forbes dataset one only needs the rank of the household (where households are ranked from highest to lowest wealth.) and its net wealth. Net wealth is provided in the HFCS data and the Forbes data. It is straitforward to rank individuals on the Forbes list (1,2,...). The richest Forbes individual has therefore a relative frequency of 1/N, the second of 2/N and so forth.

Equation 9 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. Figure 1 shows for the SCF and HFCS the empirical ccdf and wealth on a log-log scale for the tail of the data.[12] The tail is assumed to start at a value of 1 million euro (i.e. $w_{min} = 10^6$) so that a value of 1 on the x-axis corresponds to 1 million euro in wealth. The crosses represent the Forbes billionaires, the dots represent the survey households. One first observes the finding of table 5, namely that there is a substantial gap between the highest ranked survey household and the lowest ranked Forbes individual for the HFCS, but not so for the SCF. The graphs also show that most survey sample observations fall in the range of [0.01, 1] for the empirical ccdf (Shown on the graphs by the two horizontal lines). Otherwise said there are relatively few sample points at the top 1 percent of the tail of the wealth distribution. The exceptions are the SCF, and Spain and France. Both the dots and the crosses seem to closely follow a linear relationship, suggestive of a potential good fit by a Pareto power law.

---

[12]To draw the graph the first implicate is used. Other implicates lead to very similar graphs.

Figure 1: *Empirical CCDF on log-log scale*

3.3 Monte Carlo results: power law when survey data has non-response

The presence of differential non-response positively correlated with wealth will cause the empirical distribution of the tail from a survey sample to systematically differ from the actual tail distribution in the population. As wealthy households are responding less frequently when being sampled than less wealthy ones the tail in the survey sample will likely be thinner than reality. This will cause the tail index to be biased upward, i.e. showing a lower degree of wealth concentration.

How do the two methods, pseudo-maximum likelihood and regression method, to estimate the tail index perform in the presence of non-response? How large is the bias, and how precise are tail index estimates? Can extra observations of wealthy individuals from rich lists in the regression method reduce bias and increase precision? How different are estimates of tail wealth when constructed from the surveys only versus constructed from the estimated tail index (as in using equation (3))? 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 provides potentially a method to improve 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 that has no oversampling of the wealthy in the presence of non-observed differential non-response. The Monte Carlo experiment is as follows. Consider a large country with a tail population consisting of 1 million households, each with wealth above 1 million euro. 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. According to the HFCS survey results in Germany, about 1 million households have wealth above 1 million euro; in France, this is about 800000 households. See Table 3 for details. As we are only interested in the tail, the Monte Carlo only models the tail of the distribution.

Next, imagine that a survey sample is drawn from this tail population, with a sample size of 750 households. Some households respond to the survey, others don't. Survey weights are constructed for the responding households so that they sum up to 1 million. Imagine that only the aggregate non-response rate is observed; when constructing the weights there is no information available on differential non-response. Non-response correction of the weights is only based on aggregate non-response rates. For instance, if all 750 household would respond, the household weight for each individual would be equal to $10^6/750$. When less than 750 households respond, divide the 750 into non-responding N_{nr} and responding households N_r. Then each responding household gets a weight of $(10^6/750)*(750/N_r)$, so that household weights again sum up to 1 million. In the absence of differential non-response information, that is the best non-response correction that is possible.

Imagine further that all households with wealth above 740 million euro are also on a journalist rich list, say a dollar billionaires list. It is assumed that the rich list is exhaustive. From the sample of survey respondents, the tail index is estimated using the two estimation methods. 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 10000 times; i.e. 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, the tail index is estimated from the respondents (with or without the rich list).

To answer the questions above, the experiment is performed for 10 different α's (i.e. $\alpha = 1.1, 1.2, ...2.0$). According to Gabaix (2009), the tail exponent of wealth found in earlier studies is around 1.5, so that the interval of α's considered here should suffice. Each experiment for a given α is also performed for two different non-response mechanisms.

The first non-response mechanism attempts to model a reasonable relation of wealth with non-response in the population, i.e. a differential 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 also likely of surveys in other countries. In the 1992 SCF, individuals with a wealth index between 1 million and 2,5 million dollar have a response rate of 34.4 percent. This rate gradually declines to 14.3 percent for individuals with a wealth index between 100 and 250 million dollar.[13] Households non-response probability as a function of wealth is then calibrated to mimic the non-response rate as a function of the wealth index in the 1992 SCF. This is done the following way. The non-response rate 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(non-response)=0.097167+0.036594*ln(Wealth). This relationship is our first non-response mechanism.

The second non-response mechanism is a simple constant non-response probability that is set equal to the aggregate expected non-response rate of the first non-response mechanism. That is, each household has the same non-response probability. The aggregate expected non-response probability of the first non-response mechanism can be found by taking the expectation of 0.097167+0.036594*ln(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(non-response)=$0.097167 + 0.036594 * \ln(10^6) + (0.036594/\alpha)$. This gives a constant non-response rate between 62.1 percent (for $\alpha = 2$) to 63.6 percent (for $\alpha = 1.1$).

Both non-response mechanisms lead therefore to the same expected number of respondents out of a sample of 750 households. The combination of a sample of 750 households with the non-response functions 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 euro. Note that the aggregate non-response rate in the German HFCS is 81.3 percent (HFCN,2013), higher than assumed in the Monte Carlo.

Table 6 presents the results of the Monte Carlo. Shown are the estimates of the Pareto tail index under the two scenarios of non-response, using the different estimation methods. Column (1) shows the true α, columns (2) to (5) show the results under the constant non-response mechanism, and columns (6) to (9) the results under the differential non-response mechanism. Column (10) shows the number of households on the rich list, i.e. the number of households with wealth higher than 740 million euro. The (pseudo) maximum likelihood estimates α_{ml} are in columns (2) and (6). They are clearly different under the two non-response mechanisms. Under constant non-response, these estimates show a very small

---

[13]Note that in the SCF, this information can be used to correct the weights of the responding households to correct for the non-response of the households in these different strata. Such a correction can not be done however in surveys which do not use individual household data to oversample the wealthy such as most of the HFCS surveys.

upward bias, of 0.01 points for all α's, except when α is equal to 1.1 or 1.5, then the Monte Carlo indicates no bias. Of course, some slight variation is due to the Monte Carlo itself. This needs to be kept in mind in everything stated below as well. Under differential non-response, 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. So for the (pseudo) maximum likelihood estimator, non-response per se is not a problem, but "differential" non-response clearly is. The regression estimates, α_{reg}, using only the survey data are in columns (3) and (7). Under the constant non-response mechanism, the estimates of the tail index show a small downward bias, around 0.03. Again, under the differential non-response, the bias is upwards and is relatively large, around 0.07.

The regression estimates 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. There are on average 698 observations on the rich list (with a standard deviation of 26) (remember out of a population of 1 million) when α is equal to 1.1. This drops to only 2 observations when α is equal to 2. The improvement of the estimate of the tail index, in terms of a reduction in bias, under differential non-response is dramatic. Essentially, when including the rich list with the survey data in the regression method, the tail index is estimated without bias for all α's in the range from 1.1 to 1.6, and with a tiny downward bias otherwise. Also important, the reduction in standard error is impressive. Again, as one should expect, the reduction in the standard error is much larger when the tail index is lower, i.e. the number of observations on the rich list is higher. But even when the rich list contains very few individuals, two in the case of α equal to 2, both the bias in the estimate of α almost disappears, and the standard error is reduced.

Figure 2 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 sample and the rich list, together with the true power law. 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. Due to the non-response, the estimates of the ccdf from the sample observations of wealthy households will most likely be below the line implied by the true power law, i.e. 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. By adding the rich list to the survey sample the line shift 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.

TABLE 6

Monte Carlo estimates of Pareto tail index
using different estimation methods
Two scenarios for non-response

	Constant non-response				Differential non-response				
α	α_{ml}	α_{reg}	α_{regfor}	Resp obs	α_{ml}	α_{reg}	α_{regfor}	Resp obs	Rich list obs
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
1.10	1.10	1.07	1.10	273	1.22	1.18	1.10	273	698
	0.07	0.07	0.01	13	0.07	0.08	0.01	13	26
1.20	1.21	1.17	1.20	275	1.31	1.28	1.20	275	361
	0.07	0.08	0.01	13	0.08	0.08	0.01	13	19
1.30	1.31	1.27	1.30	277	1.41	1.38	1.30	277	186
	0.08	0.09	0.01	13	0.08	0.09	0.01	13	13
1.40	1.41	1.37	1.39	278	1.51	1.48	1.40	279	96
	0.08	0.09	0.02	13	0.09	0.10	0.02	13	10
1.50	1.50	1.46	1.49	280	1.61	1.57	1.50	280	50
	0.09	0.10	0.03	13	0.10	0.10	0.03	13	7
1.60	1.61	1.57	1.58	281	1.71	1.67	1.60	281	26
	0.10	0.11	0.04	13	0.10	0.11	0.04	13	5
1.70	1.71	1.66	1.67	282	1.81	1.77	1.69	282	13
	0.10	0.11	0.05	13	0.11	0.12	0.05	13	4
1.80	1.81	1.76	1.75	283	1.91	1.86	1.79	283	7
	0.11	0.12	0.06	13	0.11	0.12	0.07	13	3
1.90	1.91	1.86	1.84	284	2.01	1.96	1.89	284	4
	0.11	0.13	0.08	13	0.12	0.13	0.09	13	2
2.00	2.01	1.96	1.93	284	2.11	2.06	1.99	284	2
	0.12	0.13	0.10	13	0.12	0.14	0.11	13	1

Notes: Reported are mean estimates of Pareto tail index under two non-response scenarios. Standard errors are reported in the line below the mean. Means and standard errors are derived from 10000 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 constant non-response probability in scenario 1 and a non-response probability equal to 0.097167+0.036594*ln(wealth) in scenario 2. Estimates of tail index using maximum likelihood are in columns (2) and (6). Estimates using regression method excluding rich list are in columns (3) and (7). Estimates using regression method including rich list are in columns (4) and (8). Columns (5) and (9) report the mean number of respondent obervations (and standard error). Column (10) reports the mean number of observations on the rich list (and standard error).

The case for adding rich list data to the survey in the case of constant non-response, is more differentiated, depending on the true α. For low α, in the range between 1.1 and 1.3, the bias completely disappears, and the standard error is reduced significantly. So adding the rich list, which ranges from almost 700 observations to 186 observations, improves estimation of the tail index. For intermediate α, in the range between 1.4 and 1.7, including the rich list leads to estimates of α that are slightly biased downwards, which is an improvement on the regression estimates without the rich list, which are more biased downwards. For larger α, i.e above 1.8, the downward bias when including the rich list becomes larger. However, in al cases, the standard error reduces significantly.

The ultimate interest in the estimation of the power law is to provide an estimate of total wealth in the tail. Wealth estimated under the power law can be calculated as in equation (3). Alternatively, total wealth in the tail can be calculated from the survey directly as the weighted sum of wealth of the sample; remember that survey weights sum

Figure 2: *Monte Carlo Example of Tail of the wealth distribution*

up to population totals. To see how far off estimated wealth is from the truth, Table 7 shows total wealth in the population estimated from the survey sample and from the estimated power laws, as a ratio to true total wealth in the population.[14] A ratio of 1 signifies no bias in estimated wealth.

---

[14]True total wealth in the population is simply the sum of wealth of the 1 million households.

Monte Carlo estimates of tail wealth
as a proportion of actual tail wealth

α	Constant non-response				Differential non-response			
	survey est.	α_{ml}	α_{reg}	α_{regfor}	survey est.	α_{ml}	α_{reg}	α_{regfor}
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1.10	0.94	3.00	5.27	1.28	0.53	0.75	1.10	1.27
	3.55	33.80	72.66	0.20	0.33	0.73	7.90	0.20
1.20	0.98	1.21	1.89	1.06	0.67	0.77	0.87	1.05
	3.54	1.10	7.83	0.10	0.28	0.19	0.36	0.10
1.30	1.04	1.06	1.24	1.02	0.77	0.82	0.89	1.02
	4.43	0.27	0.65	0.06	0.41	0.14	0.20	0.06
1.40	1.00	1.03	1.13	1.02	0.83	0.87	0.92	1.01
	1.88	0.18	0.27	0.04	0.29	0.11	0.15	0.04
1.50	1.01	1.02	1.09	1.02	0.87	0.90	0.94	1.01
	0.79	0.13	0.19	0.03	0.20	0.09	0.12	0.03
1.60	1.00	1.01	1.06	1.02	0.90	0.91	0.95	1.01
	0.90	0.11	0.14	0.04	0.16	0.08	0.10	0.03
1.70	1.00	1.01	1.05	1.03	0.92	0.93	0.96	1.01
	0.26	0.09	0.12	0.04	0.18	0.07	0.09	0.04
1.80	1.00	1.01	1.04	1.04	0.93	0.94	0.97	1.01
	0.19	0.08	0.10	0.05	0.10	0.06	0.08	0.05
1.90	1.00	1.01	1.04	1.04	0.94	0.95	0.98	1.01
	0.14	0.07	0.09	0.05	0.09	0.06	0.07	0.05
2.00	1.00	1.00	1.03	1.05	0.95	0.96	0.98	1.01
	0.10	0.06	0.08	0.06	0.08	0.05	0.06	0.05

Notes: Reported are means of the ratio of estimated tail wealth on actual tail wealth under two non-response scenarios. Standard errors are reported in the line below. Means and standard errors are derived from 10000 Monte Carlo iterations as described in footnote to table 6. Estimated tail wealth used to construct ratio in columns (2) and (6) is calculated from survey only. Estimated tail wealth used to construct ratio in columns (3),(4),(5),(7),(8),(9) is constructed using the estimated Pareto tail index.

Under constant non-response, there is almost no bias in the estimated wealth. However, for the constant non-response the striking feature of the ratio of estimated wealth from the survey to true wealth (column 2) is not so much the absence of bias, but its large standard error. Estimating total wealth from the survey directly implies having a very imprecise estimate! Estimating a power law and then calculating the wealth using the estimated law clearly reduces the standard error enormously. The biggest reduction is when using the regression method including the rich list. Although that leads to a small upward bias of wealth estimates, the reduction in variability of the estimate is clearly worth it.

Under differential non-response the estimate of wealth using the survey (column 6) is, as expected, biased downwards. 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 lower degree of wealth concentration at the top, so 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 too low at a tail index level of 1.5 (the level mentioned by Gabaix (2009)) or even lower, in case of power laws with low tail indexes. Again the standard errors are relatively large, although much reduced compared to the constant non-response. The

higher probability of the very rich to not enter the sample clearly reduces the variance. Both bias and standard error can be reduced when estimating a power law. Again the regression method including the rich list performs the best. An exception occurs when α is very low at 1.1. 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 likely not commonly found anyway.

Combining all these results, the Monte Carlo seems to show that adding a rich list to the survey data and estimating wealth through the estimated power law is a reasonable idea. This idea is taken up in the next section where the results of power law estimation are shown.

4   Estimation results

How much of total household wealth is held by the one percent richest households? And how much by the five percent richest households? As the argumentation thus far shows, these seemingly simple questions are hard to answer precisely. Given the long history of the SCF, most earlier findings refer to the US wealth distribution. Historical estimates for the years 1983, 1989, 1992, 1995, 1998 and 2001 based on the US SCF of the percentage share of wealth held by the top one and five percent households in the US are provided in Wolff (2006). Estimates for the top one percent are in the order of one third of total wealth. For the top five percent the estimates are all around 59 percent. Davies et al. (2010) collect information from earlier studies and provide a list of estimates of the top one percent share for 11 countries, of which France for 1994 (21.3 percent), Italy for 2000 (17.2 percent), Spain 2002 (18.3 percent) and the US for 2001 (32.7 percent). They also provide a list of estimates of the top five percent share for 10 countries, of which Italy for 2000 (36.4 percent) and the US for 2001 (57.7 percent).

This section provides new estimates of the concentration of wealth in the tail. Estimates are based on all five implicates of the multiple imputed HFCS and SCF data. Table 8 provides estimates of the tail index using the pseudo-maximum likelihood method and the regression method. For this last method, estimates using the survey only and using the survey combined with the Forbes World's billionaire list are given. As it is unclear where the tail exactly starts, and to provide some idea of the variability of tail estimates depending on the level of wealth where the tail starts, estimates are given for three different threshold levels (2 million euro, 1 million euro and 500 thousand euro). Using a lower threshold (say 500 thousand relative to 2 million) increases the sample size. However, there is a tradeoff. On the one hand, the increased sample size should lead to more precise estimates, but on the other hand it also includes observations that potentially do not obey the Pareto tail behaviour. This itself might lead to biased estimates. Using a high level of the threshold certainly leads to fewer observations, but is more likely to restrict

the estimation on a sample that truly follows the Pareto tail behaviour.[15]

TABLE 8

Estimates of Pareto tail index

| | Pseudo max.likelihood | | | Regression method | | | | | |
| | | | | excluding Forbes | | | including Forbes | | |
	≥ 2M	≥1M	≥500k	≥ 2M	≥1M	≥500k	≥ 2M	≥1M	≥500k
USA	1.26	1.21	1.02	1.55	1.46	1.33	1.51	1.47	1.39
	0.05	0.04	0.03	0.03	0.02	0.02	0.01	0.01	0.01
Germany	1.43	1.43	1.61	1.59	1.50	1.56	1.37	1.39	1.42
	0.26	0.17	0.10	0.30	0.18	0.10	0.02	0.02	0.01
France	1.65	1.84	1.75	1.65	1.80	1.80	1.53	1.70	1.75
	0.09	0.08	0.04	0.10	0.07	0.04	0.03	0.04	0.03
Italy	2.00	1.85	1.79	2.03	1.86	1.81	1.54	1.59	1.65
	0.21	0.12	0.06	0.35	0.15	0.07	0.02	0.02	0.02
Spain	1.72	2.06	1.85	1.67	1.85	1.88	1.62	1.79	1.84
	0.28	0.18	0.08	0.15	0.11	0.07	0.08	0.07	0.05
Netherlands	1.40	3.55	2.61	0.59	2.97	2.68	1.11	1.53	1.80
	0.11	0.62	0.31	2.74	1.57	0.40	0.50	0.07	0.09
Belgium	2.22	1.79	1.77	2.28	1.87	1.80	1.88	1.81	1.78
	0.25	0.13	0.08	0.29	0.13	0.07	0.09	0.08	0.05
Austria	1.71	1.43	1.35	1.65	1.46	1.36	1.44	1.42	1.38
	0.44	0.30	0.16	0.55	0.34	0.20	0.05	0.10	0.10
Finland	2.04	2.48	2.26	1.85	2.25	2.25	1.59	2.03	2.19
	0.23	0.18	0.06	0.41	0.19	0.07	0.16	0.13	0.07
Portugal	1.28	1.84	1.59	1.10	1.56	1.57	1.29	1.48	1.52
	0.23	0.18	0.09	0.21	0.17	0.10	0.05	0.05	0.05

Notes: Mean estimate using all five implicates. Standard errors below mean estimate.

A first observation from Table 8 is that indeed the estimated tail index varies depending on where the tail is thought to start. Unfortunately, there seems to be no general rule that describes the variation of the estimate of the tail index as a function of the threshold. For some countries it increases as one takes a larger threshold, for some it decreases. The standard error of the estimates tends to increase as the threshold level increases.[16] This is especially the case for the pseudo maximum likelihood estimates and the estimates of the regression method without the Forbes data. This is mainly due to the reduction in sample size as one increases the threshold.

Next, I consider the different estimation methods. The pseudo maximum likelihood estimates can be quite different from the regression method estimates. A large difference between those two estimates is particularly the case for the SCF, but less so for the HFCS. Probably the pseudo maximum likelihood estimate is influenced by the degree of oversampling. The SCF has the highest degree of oversampling, so that survey weights can differ quite a bit for different sample observations. It is unclear how the pseudo maximum likelihood properties are in this case. Also the threshold of 500 thousand is clearly too low for the SCF. In that case, the pseudo maximum likelihood estimate of the

---

[15]Ideally one would also want to increase the threshold further to 5 million, 10 million and so on. In practice, for most countries this would lead to very few observations, creating severe small sample problems. The Netherlands already has clear small sample problems starting at 1 million euro.

[16]Standard errors are based on a bootstrap procedure using replicate weights provided in the SCF and HFCS for that purpose and are calculated using Rubin's rule. Standard errors are based on the first 50 replicate weights.

tail index is much too low (1.02). More importantly, including the Forbes billionaires in the regression method leads, in most cases, to a reduction of the estimated tail index. This result confirms the main finding of the Monte Carlo study. Adding data from a rich list should lower the estimated tail index, at least in the presence of differential non-response. For instance, compare the regression estimates with and without the Forbes data, when the threshold is 1 million euro. The reduction is the largest for the countries with no oversampling of the wealthy. For instance, for the Netherlands the estimate drops from 2.97 to 1.53. The estimate from Italy drops from 1.86 to 1.59. Both estimates including the Forbes data are much closer to the focal estimate of 1.5 mentioned by Gabaix (2009). For the SCF, with heavy oversampling, including the Forbes billionaires changes the estimates of the tail index from 1.46 to 1.47, almost no change at all.

The effect of this is that estimates of the percentage wealth share of the top 1 percent of households is relatively affected the most for the countries with no or low oversampling, the Netherlands and Italy (results in Table 9). Indeed, using only the sample information 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 12 and 17 percent, close to the estimates of Spain and Belgium. Such increase in the estimated percentage suggest 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 estimated from the survey only is 14 percent. From the power law excluding the Forbes data, it is between 15 and 16 percent. Including the Forbes data it is between 20 and 21 percent. For the SCF, the wealth share calculated from the survey is 34 percent, while it is estimated to be between 35 and 37 percent including the Forbes data, so only a marginal increase. Note that relative to the survey estimate, also in France and Spain, with heavy oversampling, the estimate using the Forbes data is very close (for France from 18 percent using the survey to 19-20 using the Forbes data, for Spain from 15 percent to 15-16 percent). For the other countries without strong oversampling, the survey estimate is also much below the regression estimate using the Forbes data.

Table 10 shows the 99th percentile of the wealth distribution, constructed using the survey or using the estimated power law. In contrast to the wealth share of the top 1 percent, the 99th percentile is relatively less affected by estimation method and the inclusion of the Forbes billionaires.

Table 11 shows the percentage wealth share of the top 5 percent of households. The estimates using the HFCS and SCF data samples directly results in estimates ranging from 61 percent in the US (comparable to the earlier mentioned number of 59 percent for historical SCF) to 26 percent in the Netherlands. Again for the US the pseudo maximum likelihood estimates are higher, but the estimates using the regression method are quite close to the sample estimate and range from 58 percent to 61 percent depending on the threshold and the in-or-exclusion of the Forbes data. For France and Spain, with heavy

oversampling, all estimates, irrespective of the method, are very similar. For France they range from 36 to 38 percent, for Spain from 29 to 32 percent. Again, the Netherlands and Italy, without oversampling, show an increase in the wealth share when using the Forbes data. For the Netherlands the sample estimate is 26 percent; using the regression method including Forbes data, the estimate ranges from 28 to 35, depending on the threshold. For Italy, the sample estimate is 32 percent; using the regression method including Forbes data, the estimate ranges from 37 to 38, depending on the threshold. Also the estimates for Belgium, Austria and Germany increase a few percentage points when using pseudo maximum likelihood or regression methods relative to the sample estimate only. Table 12 shows the 95th percentile of the wealth distribution. In line with the earlier results for the 99th percentile, the different estimates of the percentiles are less affected by the estimation method.

TABLE 9

Percentage wealth share of top 1 percent of households

when tail is replaced by estimated Pareto distribution

| | | Pseudo max.likelihood | | | Regression method | | | | | |
| | | | | | excluding Forbes | | | including Forbes | | |
	data	$\geq$ 2M	$\geq$ 1M	$\geq$ 500T	$\geq$ 2M	$\geq$ 1M	$\geq$ 500T	$\geq$ 2M	$\geq$ 1M	$\geq$ 500T
USA	34	50	55	93	34	36	41	35	35	37
	1	4	5	3	1	1	1	0	0	0
Germany	24	30	31	26	26	28	27	32	33	33
	3	18	9	4	14	6	3	1	1	1
France	18	18	17	17	18	18	18	20	19	19
	2	2	1	1	2	1	1	1	1	1
Italy	14	15	16	16	15	16	16	21	21	20
	1	1	2	1	2	2	1	1	0	0
Spain	15	15	13	14	15	15	15	16	16	15
	1	3	1	1	1	1	1	1	1	1
Netherlands	9	9	8	9	7	9	10	12	17	17
	1	1	1	1	122	1	1	4	2	1
Belgium	12	14	17	16	14	16	17	17	17	17
	1	1	2	2	2	2	1	1	1	1
Austria	23	30	41	35	31	39	40	33	34	36
	7	19	35	31	23	36	21	4	6	8
Finland	12	13	12	13	13	13	13	15	14	13
	1	1	1	1	1	1	1	1	1	1
Portugal	21	26	19	22	42	24	24	26	26	25
	3	55	2	3	158	5	3	2	2	2

Notes: Mean estimate using all five implicates. Standard errors below mean estimate.

TABLE 10
Estimates of 99th percentiles of the wealth distribution
when tail is replaced by estimated Pareto distribution
million euro

| | | Pseudo max.likelihood | | | Regression method | | | | | |
| | | | | | excluding Forbes | | | including Forbes | | |
	data	≥ 2M	≥ 1M	≥ 500T	≥ 2M	≥ 1M	≥ 500T	≥ 2M	≥ 1M	≥ 500T
USA	4.91	5.07	5.07	5.04	4.17	3.85	3.45	4.24	3.81	3.16
	0.26	0.22	0.29	0.36	0.14	0.09	0.08	0.16	0.08	0.05
Germany	1.93	1.97	1.93	1.93	1.93	1.91	1.93	1.93	2.02	2.20
	0.22	0.15	0.19	0.22	0.15	0.18	0.18	0.16	0.15	0.11
France	1.78	1.81	1.78	1.78	1.78	1.83	1.83	1.78	1.90	1.90
	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.06
Italy	2.11	2.05	2.14	2.14	2.11	2.04	2.07	2.15	2.31	2.36
	0.18	0.13	0.12	0.18	0.13	0.14	0.13	0.17	0.10	0.07
Spain	1.86	1.88	1.86	1.86	1.86	2.01	2.00	1.86	2.06	2.07
	0.12	0.12	0.11	0.12	0.12	0.11	0.10	0.12	0.11	0.09
Netherlands	1.04	1.04	1.04	1.04	1.04	1.04	1.03	1.04	1.08	1.46
	0.09	0.09	0.07	0.09	0.09	0.08	0.09	0.09	0.15	0.09
Belgium	2.62	2.63	2.86	2.86	2.60	2.52	2.58	2.75	2.61	2.62
	0.26	0.18	0.20	0.26	0.17	0.16	0.17	0.22	0.13	0.15
Austria	3.06	3.13	3.25	3.25	3.09	3.05	3.12	3.09	2.97	2.93
	1.05	1.07	1.16	0.94	1.12	1.14	1.08	0.74	0.43	0.49
Finland	1.09	1.13	1.09	1.09	1.09	1.15	1.13	1.09	1.17	1.16
	0.04	0.04	0.03	0.04	0.04	0.03	0.03	0.04	0.04	0.03
Portugal	1.24	1.31	1.24	1.24	1.24	1.37	1.34	1.24	1.40	1.39
	0.11	0.11	0.09	0.11	0.11	0.11	0.10	0.11	0.11	0.08

Notes: Mean estimate using all five implicates. Standard errors below mean estimate.

TABLE 11
Percentage wealth share of top 5 percent of households
when tail is replaced by estimated Pareto distribution

| | | Pseudo max.likelihood | | | Regression method | | | | | |
| | | | | | excluding Forbes | | | including Forbes | | |
	data	≥ 2M	≥ 1M	≥ 500T	≥ 2M	≥ 1M	≥ 500T	≥ 2M	≥ 1M	≥ 500T
USA	61	70	73	96	59	59	61	60	59	58
	1	3	3	2	1	1	1	1	1	1
Germany	46	50	50	47	47	48	47	51	52	53
	3	13	7	3	10	5	3	2	1	1
France	37	37	36	36	37	36	36	38	38	37
	1	1	1	1	2	1	1	1	1	1
Italy	32	33	33	34	33	33	34	37	38	38
	1	1	2	1	2	2	1	1	1	1
Spain	31	31	29	30	31	31	31	32	32	32
	1	2	1	1	1	1	1	1	1	1
Netherlands	26	26	26	26	25	26	26	28	32	35
	1	1	1	1	99	1	2	3	2	1
Belgium	31	32	34	34	32	33	34	34	34	34
	1	1	2	2	1	2	2	1	1	1
Austria	48	52	59	55	53	57	58	55	54	55
	8	14	25	21	17	25	16	5	6	7
Finland	31	31	30	31	31	31	31	32	32	31
	1	1	1	1	1	1	1	1	1	1
Portugal	41	45	39	41	56	43	42	44	45	44
	2	41	2	2	118	4	3	2	1	2

Notes: Mean estimate using all five implicates. Standard errors below mean estimate.

TABLE 12

Estimates of 95th percentiles of the wealth distribution

when tail is replaced by estimated Pareto distribution

million euro

| | | Pseudo max.likelihood | | | Regression method | | | | | |
| | | | | | excluding Forbes | | | including Forbes | | |
	data	≥ 2M	≥ 1M	≥ 500T	≥ 2M	≥ 1M	≥ 500T	≥ 2M	≥ 1M	≥ 500T
USA	1.37	1.34	1.37	1.37	1.37	1.28	1.03	1.37	1.28	0.99
	0.05	0.05	0.04	0.05	0.05	0.03	0.02	0.05	0.03	0.02
Germany	0.66	0.66	0.66	0.66	0.66	0.66	0.69	0.66	0.66	0.71
	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
France	0.78	0.78	0.78	0.78	0.78	0.78	0.75	0.78	0.78	0.76
	0.02	0.02	0.02	0.02	0.02	0.02	0.01	0.02	0.02	0.01
Italy	0.86	0.86	0.86	0.86	0.86	0.86	0.85	0.86	0.86	0.89
	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.02	0.02	0.02
Spain	0.88	0.88	0.88	0.88	0.88	0.88	0.85	0.88	0.88	0.86
	0.04	0.04	0.04	0.04	0.04	0.04	0.03	0.04	0.04	0.03
Netherlands	0.58	0.58	0.58	0.58	0.58	0.58	0.56	0.58	0.58	0.60
	0.03	0.03	0.03	0.03	0.03	0.03	0.02	0.03	0.03	0.04
Belgium	1.07	1.07	1.07	1.07	1.07	1.07	1.05	1.07	1.07	1.06
	0.07	0.07	0.06	0.07	0.07	0.05	0.04	0.07	0.06	0.04
Austria	0.93	0.93	0.93	0.93	0.93	0.93	0.92	0.93	0.93	0.91
	0.12	0.12	0.11	0.12	0.12	0.10	0.13	0.12	0.10	0.08
Finland	0.55	0.55	0.55	0.55	0.55	0.55	0.55	0.55	0.55	0.55
	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Portugal	0.48	0.48	0.48	0.48	0.48	0.48	0.48	0.48	0.48	0.48
	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03

Notes: Mean estimate using all five implicates. Standard errors below mean estimate.

5 Conclusion

The wealth distribution is an important variable for both economic research and policymakers. Not only do policymakers care about wealth for fiscal policy purposes, the share of wealth held at the very top has become an important parameter used to calibrate macroeconomic models. Yet, our knowledge of the wealth distribution is less than perfect. This paper has investigated how differential non-response in household wealth surveys affects tail wealth estimates.

The results clearly indicate that survey wealth estimates are very likely to underestimate wealth at the top. This is caused by differential non-response that cannot be remedied by appropriate reweighting of sample observations. Journalist list such as the Forbes billionaires can help in improving top wealth estimates. This is not so much so because of the wealth numbers of these billionaires itself, rather, the combination of survey data and rich list leads to improved estimates of the pareto tail index. 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 non-response*. In that case, journalist 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, if there is evidence that differential 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

[1] Benhabib Jess, Bisin Alberto and Shenghao Zhu, 2011, "The distribution of wealth and fiscal policy in economies with finitely lived agents, *Econometrica* Vol 79, No 11, 123-157.

[2] Cagetti, M. and DeNardi, M., 2008, "Wealth Inequality: Data and Models," *Macroeconomic Dynamics*, 12, S2, 285313.

[3] Castañeda Ana, Díaz-Giménez Javier and José-Víctor Ríos-Rull, 2003,"Accounting for the U.S. Earnings and Wealth Inequality,'*Journal of Political Economy*, Vol 111, N0 4, 818-857.

[4] Clauset Aaron, Shalizi Cxosma Rohilla, and M.E.J. Newman, 2009," Power-law distributions in empirical data", *SIAM Review* 51(4), 661-703

[5] Davies James B., and Shorrocks Anthony ,1999,"The distribution of wealth," Chapter 11 In *Handbook of Income Distribution: Volume 1* Edited by A.B. Atkinson and F.Bourguignon 121: 223-254.

[6] Davies James B., Sandstrom Susanna, Shorrocks Anthony and Edward N. Wolff, 2010,"The level and distribution of global household wealth," *The Economic Journal* 121: 223-254.

[7] HFCN, 2013, "The Eurosystem Household Finance and Consumption Survey: Methodological report for the first wave, *ECB Statistical Paper Series*, No. 1.

[8] Gabaix, Xavier, 2009,"Power laws in economics and finance",*Annual Review of Economics* 1: 255-293

[9] Kennickell, Arthur, 1997, "Consistent weight design for the 1989,1992, and 1995 SCFs, and the distribution of wealth, Appendix A" *Federal Reserve Board, Survey of Consumer Finances Working Papers.*.

[10] Kennickell, Arthur, 2000, "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, Oregon.*.

[11] Kennickell, Arthur, 2007, "The Role of Over-sampling of the Wealthy in the Survey of Consumer Finances," *Federal Reserve Board, Survey of Consumer Finances Working Papers.*.

[12] Kennickell, Arthur, 2009, "Getting to the Top: Reaching Wealthy Respondents in the SCF," *Paper prepared for the 2009 Joint Statistical Meetings, Washington, DC.*

[13] Kennickell, Arthur and D. McManus, 1993,"Sampling for Household Financial Characteristics Using Frame Information on Past Income," *Paper presented at the 1993 Joint Statistical Meetings, Atlanta, GA*

[14] Kennickell, Arthur and R.L. Woodburn, 1997, "Consistent Weight Design for the 1989, 1992, and 1995 SCFs, and the Distribution of Wealth", *Federal Reserve Board, Survey of Consumer Finances Working Papers.*

[15] Klass, Oren S., Ofer Biham, Moshe Levy, Ofer Malcai and Sorin Solomon, 2006, "The Forbes 400 and the Pareto wealth distribution" *Economics Letters* (90): 290-295.

[16] Levy, Moshe and Sorin Solomon, 1997, "New evidence for the power-law distribution of wealth" *Physica A* (242): 90-94.

[17] Ogwang, Thomson, 2011, "Power laws in top wealth distributions: evidence from Canada," *Empirical Economics* 41: 473-486.

[18] Wolff, Edward N., 1996, "International comparisons of wealth inequality, "*Review of Income and Wealth* 42: 433-451.

[19] Wolff, Edward N., 1990, "Methodological issues in the estimation of the size distribution of household wealth," *Journal of Econometrics* 43: 179-195.

[20] Wolff, Edward N., 2006, "Chapter 6: Changes in Household wealth in the 1980s and 1990s in the US", in Edward N. Wolff, Editor, *International Perspectives on Household Wealth* , Edgar Publishing Ltd.

A APPENDIX

The Appendix shows the tail of the wealth distribution (starting at 1 million euro) together with the estimated relationship on a log-log scale.

Figure 3: *Tail of the wealth distribution: USA*

Figure 4: *Tail of the wealth distribution: Germany*

Figure 5: *Tail of the wealth distribution: France*

Figure 6: *Tail of the wealth distribution: Italy*

Figure 7: *Tail of the wealth distribution: Spain*

Figure 8: *Tail of the wealth distribution: Netherlands*

Figure 9: *Tail of the wealth distribution: Belgium*

Figure 10: *Tail of the wealth distribution: Portugal*

Figure 11: *Tail of the wealth distribution: Austria*

Figure 12: *Tail of the wealth distribution: Finland*