EUROPEAN CENTRAL BANK

EUROSYSTEM

HOUSEHOLD FINANCE AND CONSUMPTION NETWORK

Oversampling the wealthy: eye for an eye, euro for a euro

October 2008

1. When a particular part of the population is especially important for a survey, oversampling that group can help the survey to provide better estimates. However, oversampling is not always easy or inexpensive. The ideal situation is one where information exists for the population that can help to discriminate the interesting sub-group and that information is available for sampling. Sometimes the information is weak for the intended purpose and sometimes there are restrictions on the use of information that make sampling difficult or impossible.[1]

2. This note tries to make explicit why, in the case of wealth surveys, oversampling of the relatively wealthy could be desirable and how some countries, with different procedures, have tried to achieve such oversampling. It draws on information available in papers by A. Kennickell, O. Bover and others (see the reference list). Technical questions of bias correction through unit non-response correction and variance reduction are not addressed directly here.

3. The first section of this note describes the distribution of wealth and differential response rates, and how oversampling can increase the efficiency of the survey and decrease non-response bias. The second section outlines possible features of oversampling as used in existing and prospective wealth surveys. The annex details some country procedures for oversampling, in the US, Spain, France and Cyprus

- **In countries where issues related to wealth and specific financial instruments are priorities of the survey, the possibility of oversampling should be carefully studied, taking into account institutional circumstances and related costs.**
- **Given that in some countries, good oversampling strategies may require cooperation with other institutions and, hence, solving complex questions of governance and confidentiality, related initiatives should start as early as possible.**
- **The methodology and practices applied to oversample the wealthy may be further developed and improved in successive waves of the survey.**

[1] Where such population information is not available, screening is a possibility. But effective screening requires that the identification of the interesting subgroup can be made cleanly for virtually all cases. Because screening requires some sort of contact with the selected cases, it tends to be quite expensive, and will not be mentioned further.

1 Wealth is unequally distributed and the wealthy participate differently in surveys

4. Wealth surveys are usually seen to face two conflicting constraints that affect the sample design. For the study of the broad financial behaviour of households, the sample should represent the population as a whole, and households should be selected such that the proportion of various "types" approximately mirrors the overall population. For the study of wealth, the sample should represent wealth as a whole, and each "euro" of wealth, which is highly skewed in its distribution, should contribute more or less in the same proportion; obviously, such a sample would require either oversampling of the wealthy or an enormous proportional sample. Practical considerations drive the need to consider a compromise to accommodate the two needs. The remainder of this section considers the trade-offs in greater detail.

1.1 Wealth is unequally distributed

5. Wealth inequality is thoroughly documented, and is usually much larger than income or consumption expenditure inequalities (Davies and Shorrocks [2000]). As an example, the two curves in Figure 1 show the Lorenz curves of real and financial assets of Italian households. The Lorenz curve, which shows the cumulative distribution of assets against the cumulative distribution of households ordered by wealth, can be used to assess the concentration of wealth. For example, based on the curve, the wealthiest 20% of the households in Italy hold approximately 75% of the financial wealth.

Figure 1: Lorenz curve for the distribution of real (left) and financial (right) assets (SHIW 2004)

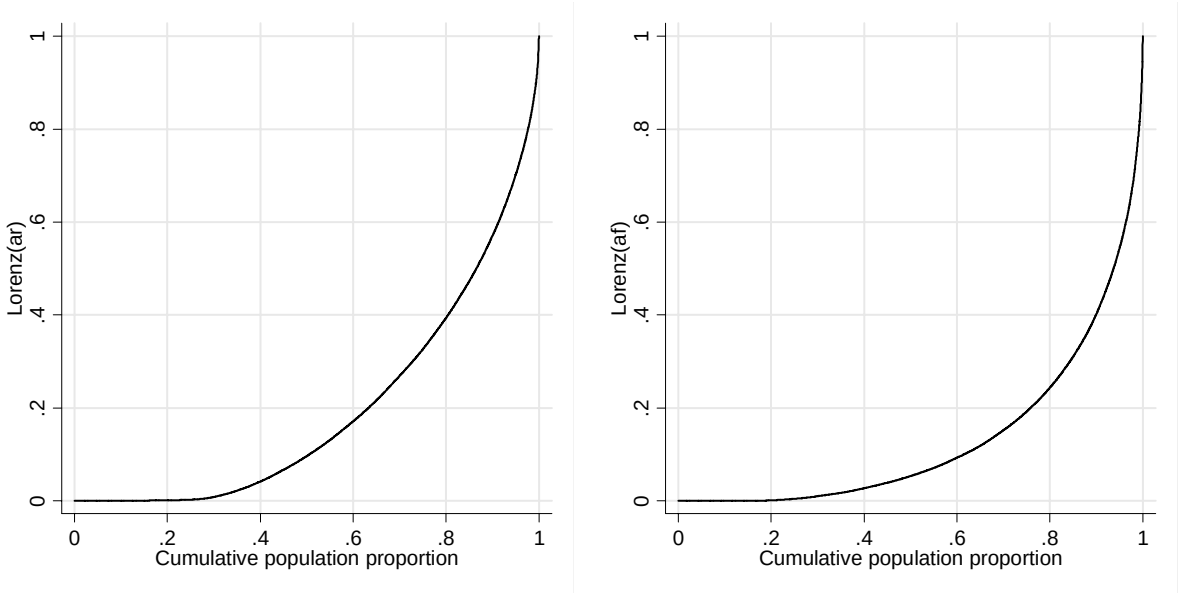

Source: SHIW 2004 – ECB calculations. The Lorenz curves use the sample weights of the households (Beach and Kaliski [1986]). One point on the curve reads as follows: the poorest x% of the population of households hold y% of the total real/financial assets.

6. The problem of dealing with the skewness of wealth is compounded by the fact that the variety of assets possessed also tends to increase with the level of wealth. Some financial products are held by only a relatively small fraction of the population. In other words, other variables of interest may be even more

unequally distributed over the population than aggregate net wealth: government securities for example are held by 8% of the population in Italy according to SHIW. This implies that analysis of the behaviour of bond holders out of a total sample of 8,000 households will have to be done only on 640 households in the absence of oversampling.

7. Using data from a purely random selection of units, for example, would at best yield a statistically very inefficient estimate of the distribution of wealth. As Kennickell (2005) points out, where there are groups in the population that either posses a rare trait of interest or that exhibit relatively high variability of the variable of interest, there may be gains from sampling disproportionately larger fraction of observations from those groups. In the SCF, of 400 observations with direct holding of bonds, only 10% were from the area probability sample (the type of proportional sample used as one component of that survey).

8. To give an example taken from the Spanish survey on wealth (EFF), it is estimated from tax records, as shown in Bover (2004), that 0.4% of the population of households holds 40% of taxable wealth. In a random survey of 5000 households we would expect only 20 such households (this is the maximum given that the response rate for rich households is lower, see below). In contrast, the EFF sample contains 500 of them. In the SCF the effect is similar: without oversampling there would only be 41 households in the top 1% of the wealth distribution; with the oversampling there are 715.

1.2 Response rates are lower for wealthy households

9. A second issue is that response rates are most of the time not uniform across the sample, but tend to have a clear non-random component. Although non-response is a complex phenomenon, it is clear in at least three countries (ES, IT and the US) that wealthier households tend to be less likely to respond, either through outright refusal, more non-contact or even interviewer decisions (different conversion techniques, etc.). Figures 2 and 3 document the decrease in response rate as the wealth increases, in the SCF and in the EFF. Kennickell points out that this decrease could be the combination of several factors, both on the respondent and on the interviewer side, and has to be borne in mind in any case, whether oversampling is implemented or not.

10. Although no oversampling takes place in SHIW, a specific study in 1998 (d'Alessio and Faiella [2002]) with clients of a commercial bank showed that non-response in the experiment was correlated with wealth; the net financial wealth of respondents is only 58% of the wealth of non-respondents, resulting in an important bias. (As mentioned in D'Aurizio et al. [2007], there is another problem, which cannot be addressed with oversampling: the under-reporting of assets).

11. The differential non-response rate, if it is not compensated by post-survey adjustments, will cause measurement bias. However, if the sample is selected so that some factors correlated with aspects of wealth can be observed for all sample elements (register based or collected through the survey), this information may be used to guide post-survey adjustments to compensate for non-response and possibly to reduce sampling error.

Figure 2: Response rate in the SCF 1992-2001 by wealth stratum in the list sample (7: highest wealth index)

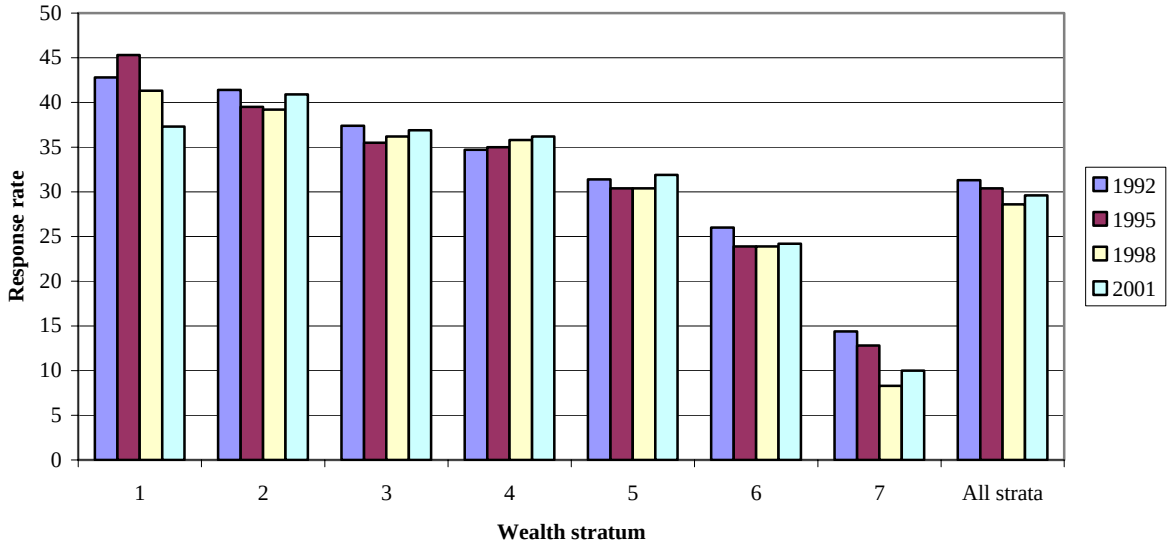

Figure 3: Never at home rate and cooperation rate, EFF 2002 (Bover 2004)

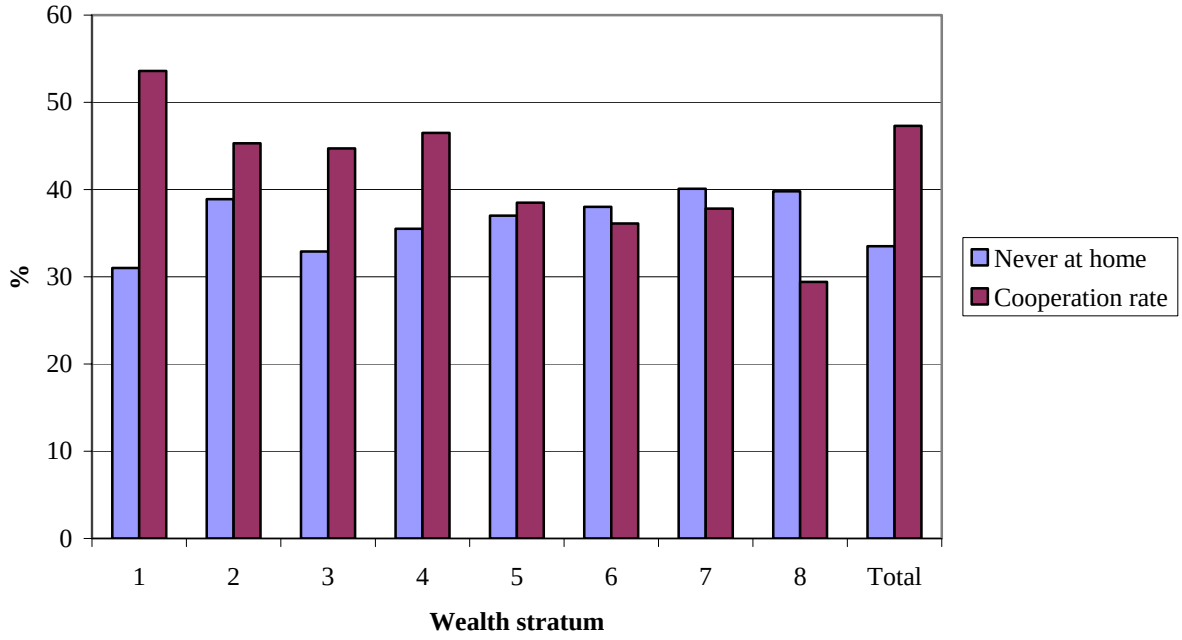

12. The lower response rate for wealthy households may hold in all countries. In the French Enquête de Patrimoine 2004, no direct wealth measure or estimate was used for the oversampling. Instead, some occupations and geographical areas were oversampled: white-collar workers, areas in and around Paris, etc. The response rate is not specifically lower in those areas. In the Cyprus SCF, the response rate is only marginally lower in the "wealthy" stratum: 69% vs. 75% in the area probability sample. However, in these examples it is not clear whether the stratification is sufficient to detect a difference by wealth.

1.3 The effect of oversampling on sample estimates

13. Where there is differential nonresponse over wealth groups, a well structured design for oversampling allows efficient correction through post-survey weight adjustments. Kennickell (2007) estimates the effect of oversampling in the US by comparing estimates only from the area probability sample of the SCF to estimates using the full area probability plus list samples.

Figure 4: Distribution of net worth in 2004 based on full sample minus distribution based on area probability sample only, 2004 dollars, by percentile of net worth (from Kennickell [2007])

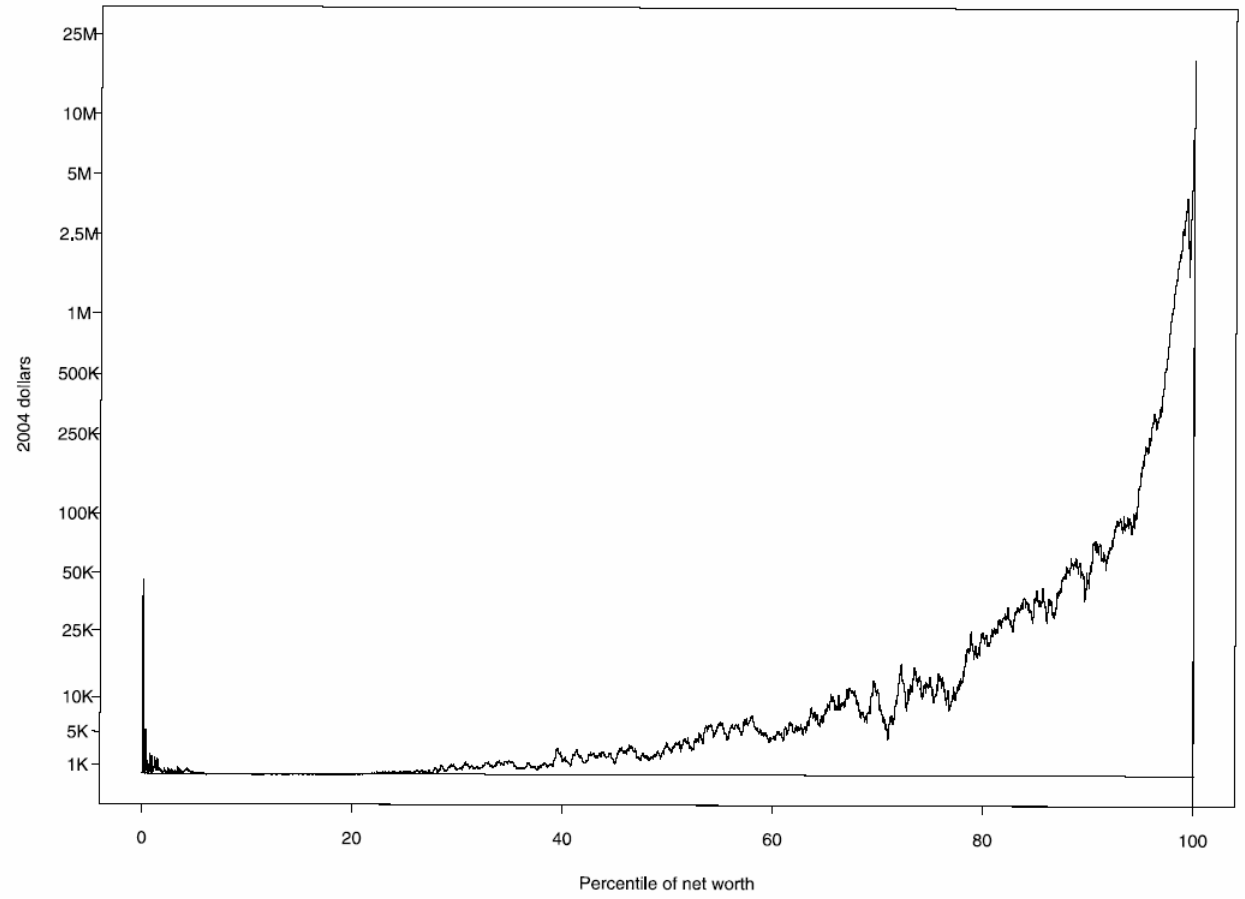

Figure 4 shows graphically the effect on the distribution of wealth. Kennickell (2007) comments:

14. From about the 20th percentile and higher, the level of net worth at each percentile under the combined AP and list samples is higher than under the AP sample alone. The approximate increase in the level of the wealth distribution from integrating the list sample is $250 (1.9%) at the 25th percentile, $3,500 (3.9%) at the median, $13,600 (4.3%) at the 75th percentile, $43,600 (5.5%) at the 90th percentile and $2,661,000 (74.0%) at the 99th percentile.

15. Another effect also mentioned by Kennickell is the increase in precision (i.e. the reduction in standard errors), for the high net worth groups but also for the other groups. The estimates of wealth are thus also more precise, as illustrated with the confidence intervals in Figure 5.[2]

Figure 5: Net worth shares in 2004 for various percentiles of the net worth distribution (Kennickell [2007])

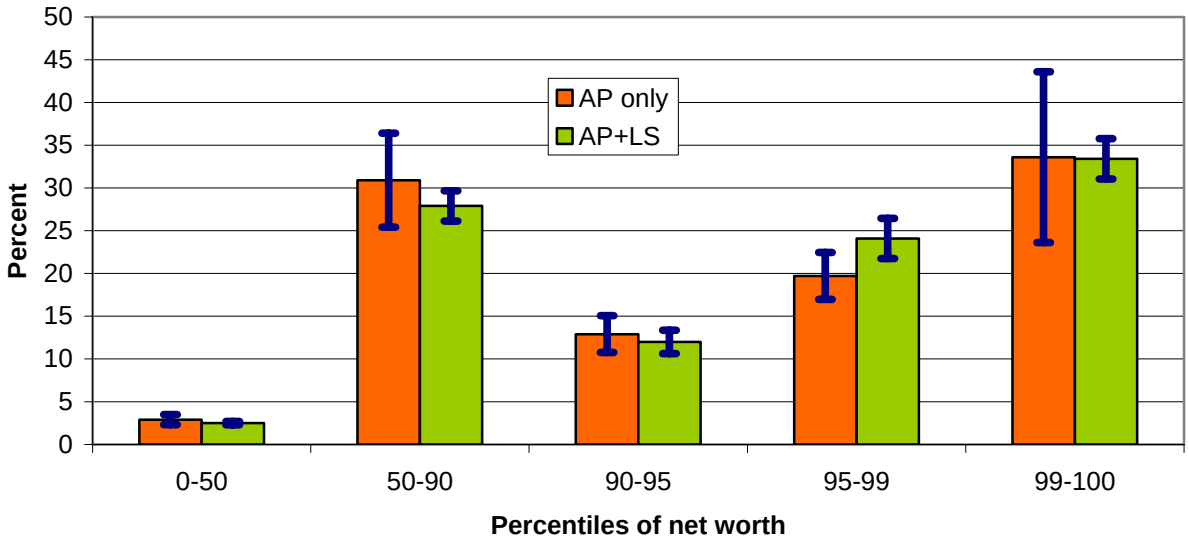

2 Overview of possible oversampling procedures for the wealthy

16. There is no ready-made solution for the oversampling of the wealthy that would work effectively in every country. Each country has different frame registers available, different institutional factors and sometimes different rules that affect what is possible in designing an oversample of the wealthy. However some general features are common to most, if not all, countries.

2.1 Oversampling can be achieved through the use of different sources

17. The main difficulty in oversampling the wealthy is the requirement to have information available that exists for the complete population (or at least a well controlled part of it) and that is a good indication for the variable for which we wish to oversample—in this case wealth (usually gross assets but this could be used as well for net wealth).

18. The accuracy of the oversampling will depend on the richness of the sample frame. The more information that is available and the higher the correlation between this information and wealth, the larger the effect of oversampling on the accuracy and precision of the survey estimates. External data can come from very different sources: electricity consumption (as in Cyprus), income tax returns (US), wealth tax (Spain), or census/dwelling based (France, see the Annex below for all examples).

[2] The underlying figures include an adjustment to the estimates for the area-probability sample to compensate for the fact that fewer observations were used in computing the standard error.

2.1.1 Wealth tax

19. In principle, the most appropriate source for the oversampling of wealth seems to be the wealth tax registers, as the wealth tax usually covers most of the aspects of wealth that we are concerned with in the HFCS: properties, financial assets, other valuables.

20. However among euro area countries only a few countries have a wealth tax, namely France, Spain and Greece, while Austria, Finland, Germany, Luxembourg, and the Netherlands have abandoned this tax in recent years.

2.1.2 Income tax

21. By contrast, all euro area countries have some form of income tax, though the income that is included varies (e.g., in DE most capital income is taxed at the source, not on the individual tax return). Although there is no perfect match between income and wealth, income tax information has been successfully used to oversample the wealthy.

22. The crucial step is to construct an estimate of wealth based on the information available in the income tax form. This can be done with a simple accounting model, as in the case of the 1989 SCF. In that survey, wealth was approximated for sampling primarily as the sum of each capital income divided by the average market return in that category. But key problems with this approach are that rates of return do in fact vary considerably across individuals, capital gains and negative incomes are difficult to incorporate, and housing wealth can only be poorly estimated.

23. The estimate of wealth can also be based on an estimated model, if a matched dataset with observations on wealth and income tax information is available. This model has however other drawbacks—principally that the model implicitly embeds rates of return at a given period. In the case of the SCF, it has been fruitful to employ a mixture of an accounting model and an estimated model to based non-devise the variable used for sample stratification.

2.1.3 Property tax

24. A potential source of information that could be used is the system of property taxes that exist in most countries. In principle, property taxes are a function of the value of the real estate or other property being taxed. Land and business properties may be taxed as well. Since real estate is the largest store of wealth for most households in many European countries, it may be possible to oversample the wealthy on the basis of this tax.

25. However, problems may arise if the value of the property is not reappraised often enough, or if the information is only available per property instead of per owner. Furthermore, property taxes are sometimes levied at the local government level, making variations from county to county potentially problematic and the implementation of a national sample difficult.

26. The information available in the property tax records might also complement well the income tax records on the real estate wealth.

2.1.4 Municipality level information

27. Information at the municipality level is often more readily available. Information from tax records or census estimates is aggregated by small geographical area and gives the distribution of some characteristics of the population. Of interest for the oversampling of the wealthy are reported income, but also socio-economic or occupation classifications. This allows the determination of affluent neighbourhoods.

28. Nevertheless intra-municipality variation is often high and an oversampling based uniquely on these criteria may only be marginally useful. There have not been, to the best of our knowledge, wealth studies based on such oversampling. The main advantage of this source is its relative ease of access.

2.1.5 Dwelling level information

29. Other sources may be available, at the dwelling level. For example, detailed census records can show the characteristics of the residents of a dwelling at the last census. This information can be used to oversample characteristics (or combinations thereof) positively correlated with wealth. Electricity bills have been successfully used by Cyprus (see Annex) to oversample the wealthy, though the correlation between electricity consumption and wealth could be low in other countries.

2.2 Oversampling can be a complement to another household frame or the only used frame

30. In some cases the frame available for the oversampling of the wealthy will not be appropriate for the poorer households. For example, wealth tax information may only be available for households close to the threshold of such a tax. Since the HFC survey needs to ensure the representativity of the average household, this implies that two frames may have to be used in such instances. Alternatively, the practice of the Spanish EFF is to merge (wealth and income) information from the wealth tax (which is levied on a selected group of households) to a frame that is representative of the non-institutionalised population (the Census of Municipal Population), thus resulting in a single frame.

31. Oversampling can be done together with the general proportional sample if the frame with wealth information is also representative of the whole population of households, or can be done in two independent sub-samples (area probability and list sample in the case of the US and Cyprus SCF for example). In that case the question of weighting of these two subsamples has to be addressed, as it may be in some cases more difficult to establish the relationship between the two samples.

2.3 Cooperation between survey and owner of frame data can be rewarding to both

32. Since oversampling of the wealthy may require the exchange between tax authorities and the survey institution of information that is usually highly confidential, there has to be mutual trust and interest of all partners in the process. The owner of the frame data (we will use tax authority for short here) can be involved in the procedure and will usually gain information on general structural characteristics of the nature on the taxpayers (of course no identifiable data should ever be transmitted

back). In countries where such transfer of confidential data has been done, the involvement of both parts is viewed as positive and successful.

33. Furthermore, oversampling is not a one-step affair that should be done once or only every two or three years. The example of the SCF is illustrative of what can be achieved through the collaboration of the institution in charge of the survey and the tax authority: the sample selection of one year can be used to increase the effectiveness and the appropriateness of the sample selection for the following wave. This can include for example matching the micro data from the survey to the information from the tax register, to model in more detail the forecasting of wealth. For this to be possible the tax authority must be willing to (and has to be allowed to) transmit these data.

34. In existing surveys with oversampling, the confidentiality of the information has been protected in such a way that nobody has had access to the full set of survey and register information. In the case of the EFF, a complex agreement was negotiated among the BdE, the tax authority, the national statistical agency and the field work company to design and implement their sample. In the case of the SCF, the Federal Reserve accesses tax-based information with the permission of the tax authority under a specific provision of the tax law that seriously limits the degree to which the data may be used and that imposes a number of stringent requirements. Privacy laws and administrative rules may impede direct access to the most useful data and may inhibit the use of such information for post-survey weighting adjustments as well. Thus, it is important for countries planning for oversampling to begin early to explore the availability of data and the possible constraints on their use.

2.4 The extent of oversampling can vary, depending on the frame information and purpose of the survey

35. As mentioned earlier in this note, a survey devoted entirely to wealth estimation is most efficient if the probability of inclusion is proportional to wealth. However the HFCS will have many more uses than this. The sample has to be a compromise between conflicting objectives. In all countries implementing oversampling of the wealthy, there has been a compromise composed of a general purpose sample (where each unit has a minimum probability of inclusion) and some type of stratification with increasing probabilities of inclusion in strata deemed likely to consist of progressively wealthier households.

36. There are no hard guidelines as to how much oversampling should be done, although in principle some form of Neyman allocation could be used.[3]

37. In practice, the number of wealth strata varies between countries: from one (in Cyprus) to eight (in Spain). The sampling rate increases with the wealth strata (the highest stratum in the EFF seems to have been oversampled by at least more than 25 times; in the SCF the rate of oversampling in the top stratum cannot be revealed exactly, but it is far higher. In contrast, if the strata are only marginally related to

1. [3] The Neyman allocation indicates the allocation of a sample across strata such that the sample minimizes the variance for a particular variable. If the variable is some form of wealth, and the strata are wealth related, in the Neyman allocation the size of the strata will depend on the variance of wealth and the number of households inside each strata.

wealth, the effective oversampling for a given allocation of cases has to be lower (only 4 times for the highest strata in the Enquête de Patrimoine).

38. The amount of oversampling will therefore also vary between surveys. In the SCF around 35% of the final sample comes from the list (wealthy) sample; in the EFF over 40% of the completed interviews correspond to wealth tax filers.

2.5 Oversampling may have some disadvantages but improves the quality of many estimates in the survey

39. Giving the outright cost of oversampling is difficult. Oversampling can have several disadvantages, some monetary, others in the added complexity of running the survey. As mentioned earlier, the initial difficulty is in obtaining the sampling frame appropriate for the oversampling of wealth. Another disadvantage is that administrative data are usually available only with a delay: for example tax files are only available long after the fiscal year (2 years in the US). Therefore the survey has to be run with older data – which can complicate somewhat the field procedure, weighting and analysis and can weaken the effect of the oversampling.

40. Since wealthier households are in general more difficult to contact, oversampling the wealthy will necessitate more attempted contacts; because wealthier households tend to have more complicated finances, they also tend to have longer interviews. These factors will likely imply higher costs per completed interview. Interviewers will probably need to be trained specifically for high net worth households. Additionally, if oversampling is done with two separate sub-samples, this will be an additional complexity in the field and in the post-field phases of the survey. Finally, oversampling can make the problem of anonymisation of the data more acute.

41. On the positive side, if done correctly, the improvement in wealth estimation achieved by the oversampling may partly compensate the increased cost, and the total sample size could be reduced. Taking the above-mentioned example from Kennickell (2007), adding 1500 high net worth households decreases the standard error of the wealth share of the 50-90% wealth group by a factor of 3. In terms of achieving a similar effect on the level of precision, a very substantial increase in the sample size would be necessary to obtain this increase in precision without oversampling.

2.6 Oversampling in a cross-country comparative survey

42. When designing a cross-country survey, it is important that every step is taken to ensure the comparability of results. The sample is no exception. The comparability of the wealth distributions across countries, especially for the top of the distribution, might suffer if only some countries implemented oversampling of the wealthy. This could also be the case if the procedures for oversampling differed significantly over countries. However it might be argued that in the case of oversampling it would be counterproductive to opt for the lowest common denominator to be applied in all countries. On the contrary, best practices (oversampling based on wealth tax/income tax/other proxies) should ideally be the target for all countries and resolving systematic differences due to sample design effects might require

additional research. Detailed metadata on sample features and design will be important for combining different designs with and without oversampling.

Annex: Sampling procedures and oversampling of the wealthy in selected surveys

Sampling procedure in the Survey of Consumer Finances

This section is largely taken from a collection of articles by Frankel and Kennickell (1995) and Kennickell (1999, 2005, and 2007) and does not do justice to the original articles.

Area probability sample

A national area-probability (AP) sample provides good coverage of widely spread characteristics. The sample selects household units with equal probability through a multistage selection procedure. In the first stage, a group of major metropolitan areas (24 such areas in the sample used for the 2007 SCF) is selected as high-level sampling units ("national frame areas" or NFAs) with probability one; for the remainder of the country, counties and metropolitan statistical areas are stratified by a variety of characteristics, and NFAs (56 such areas for the 2007 SCF) are selected proportional to their population. In the second stage, smaller areas ("segments") are selected within NFAs to serve as the basis for the ultimate samples. Finally, at the time of the execution of the SCF, housing units within selected areas are enumerated (either by using postal address sequences for urban areas or by using "listers" to record every address on a map of the area), and a set of ultimate observations is drawn.

43. Although an AP sample is an efficient means of generating a nationally-representative sample, by itself it has two important shortcomings for the SCF. First, given the concentration of ownership of many assets, it is very unlikely without an enormous sample size, that an AP sample would yield sufficient observations for the analysis of many types of financial behaviour. Second, there is ample evidence (discussed below) that the nonresponse in the SCF depends on wealth (or variables correlated with it).

List sample selection

44. The second part is a list sample, which is selected by FRB staff from statistical records derived from tax returns. These records are provided by Statistics of Income (SOI), a statistical arm of the tax authority, to the FRB under strict controls on the use of the data. This file is a sample from the full set of tax returns filed in a given year.

45. The list sample for the 1983 survey, the first of the series, was stratified by income categories created using a set of rules that have not been cleared for public release. The 1989 SCF was the first to use an explicit wealth index approach. This general type of index might be expressed as

$$WINDEX\,0 = \sum_{j} \frac{1}{r_{ij}} Y_{ij}$$

where "*i*" indexes individuals, "*j*" indexes components of capital income, and *rij* is the rate of return associated with capital income component *Yij*. Thus, if a person had interest income of $100 and the rate of return were 5 percent, then the contribution of this item to the wealth index would be $2,000.

46. Following a very long negotiation, permission was gained during the preparation for the 1995 SCF to perform a match of the 1992 survey wealth data with the single year of SOI data that had been used in selecting the 1992 sample. A variety of specifications were tried to find the best fit of the wealth data, consistent with the inclusion of a core set of variables and a few general theoretical principles. The data strongly rejected the WINDEX0 income-return model as the optimal specification, and coefficients estimated for the narrow set of variables in that model implied unbelievable rates of return.

47. Beginning with the 1995 SCF, the wealth index used for stratification was altered to include a component calculated from a model estimated from a regression of actual wealth observed in the previous survey on its corresponding original frame data. This model, referred to here as “WINDEX1", offers a flexible way of accounting for systematic patterns in the structure of various types of income–including non-capital income–on wealth.

48. To hedge against the possibility of misclassification from the WINDEX1 model, the 1995 and later surveys have pooled estimates of WINDEX0 and WINDEX1; that pooled estimate is referred to here as “WINDEXM”.

49. After defining the stratifying variable in terms of the whole population, the data file was reduced for the actual selection to include only cases that filed returns from a county included in the 100 PSUs underlying the AP sample. An implicit substratification was imposed by sorting the resulting file by age and a measure of financial income. Within each stratum cases were selected using systematic selection from a random starting point. Cases were oversampled by a progressively larger proportion with increasing strata.

50. Analysis of the performance of this sampling method strongly suggested that there would be an additional gain from using multiple years of tax-based data in order to smooth transitory variations in income components. The 2001 and 2004 SCFs used two years of data and the 2007 used three years of data.

Figure 4: Distribution in the rank of net worth in the 2001 SCF list sample; 10th, 25th, 50th, 75th and 90th percentiles of the distribution; by list sample stratum.

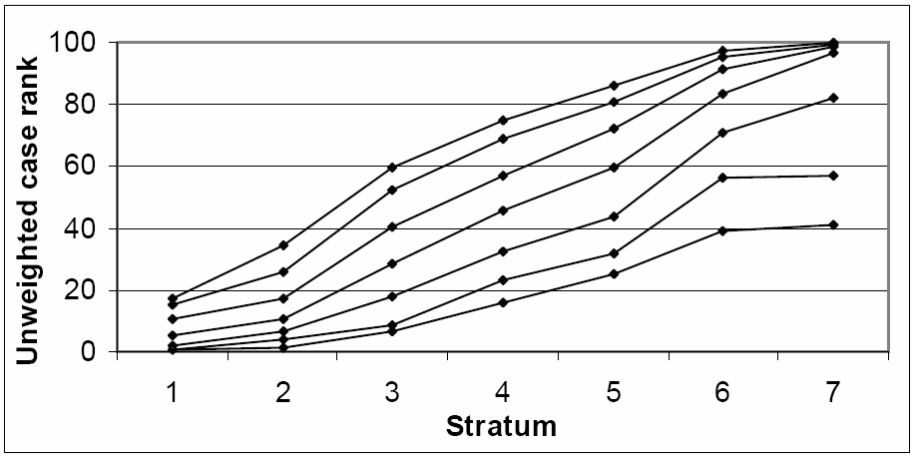

Composition of the final sample

51. About 98 percent of SCF cases with at least $5 million of net worth in 2004 derived from the list sample; more than 85 percent of cases with at least $1 million of net worth and about 75 percent of the cases with at least $500 thousand dollars came from this sample. Thus, it is clear in an informal sense that the list sample adds substantially to analysis that depends on good representation of the upper tail of the wealth distribution.

Figure 6: List sample as a proportion of the combined samples (SCF 2004)

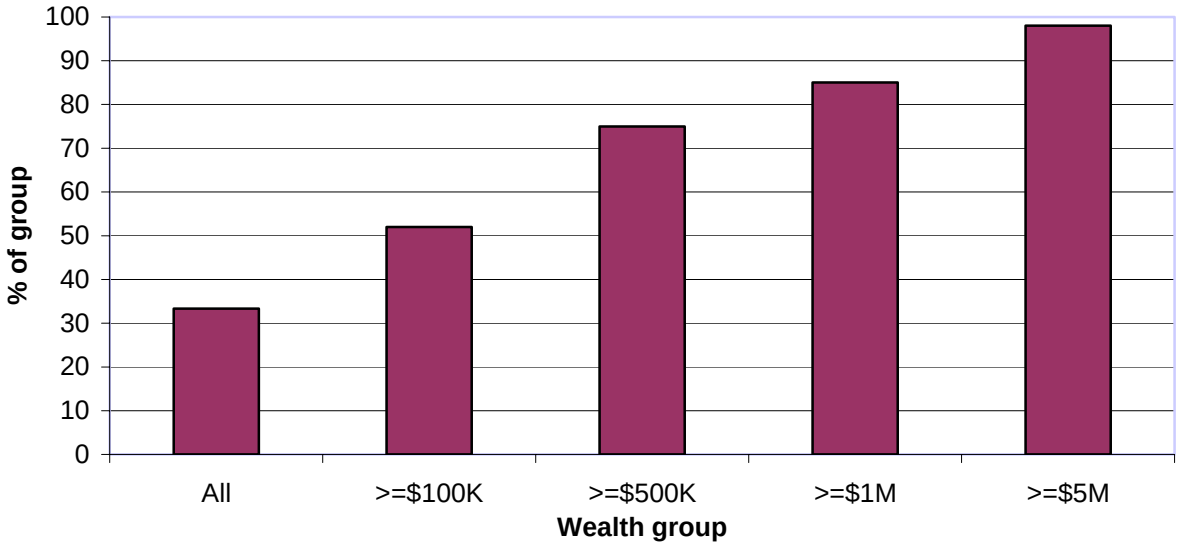

Confidentiality and protection of privacy

52. Proper treatment of data to minimize the intrusion on the privacy of respondents is a very high priority. This treatment is regulated formally by contracts between the Federal Reserve Board (FRB) and SOI, FRB and National Opinion Research Center (NORC) at the University of Chicago, and NORC and SOI, and by the provisions of the Internal Revenue Code and the Confidential Information Protection and Statistical Efficiency Act (CIPSEA). FRB staff select the list sample from a file containing no names, and they pass the Social Security Numbers (SSN) of the cases selected to SOI along with a temporary identification number. SOI retrieves the names and addresses corresponding to the SSNs, and passes that information to NORC. NORC provides the FRB with the answers to the survey questions along with the original temporary ID number for each case. In no case is the FRB allowed to know the name of any survey participants. SOI is never given the link between the survey answers and the final ID numbers; indeed, SOI encourages FRB staff to take rigorous measures to minimize the likelihood that anyone at SOI or elsewhere might be able to use the survey answers to guess the identity of the participants. NORC never receives any financial data from tax returns. Disclosure of identifying information on any SCF respondent is a serious criminal offence under the law.

53. By agreement with SOI, the members of the list sample are given an opportunity to refuse participation in the survey before they are contacted by an interviewer.

Sampling in the Spanish EFF survey

54. This section is based on Bover (2004).

55. The oversampling in the EFF was achieved thanks to the collaboration of the INE (Spain's statistical institute) and the Tax Authorities (TA), through a complex co-ordination mechanism that enabled the TA's strict confidentiality requirements to be observed at all times. Specifically, the TA devised a wealth strata-based random sample drawing on the Padrón Continuo (a continuously updated municipal population census) provided by INE, following the guidelines of the sample design prepared by INE. This provides the EFF with a unique population frame for its sample, thereby ensuring the representativeness of the information obtained while attaining accurate information on the behaviour of the richest household segment. Finally, a complex procedure for replacing non-respondent households was incorporated into the sample design, thus ensuring the maintenance of the sample's desirable characteristics.

56. In Spain there is a wealth tax ('Impuesto sobre el Patrimonio') and it is on the individual wealth tax files information that the EFF oversampling is based. This is in contrast with the SCF where a wealth index is constructed by integrating the information about asset income coming from the individual income tax files since there is no wealth tax in the US. People liable to the wealth tax in Spain were, in 1999 (which was the tax year used in selecting the sample), those with taxable wealth over 104,000 €. In 1999 around 980,000 individuals (corresponding to approximately 700,000 households) filed a wealth tax return.

57. The choice of defining the wealth strata was based on the SCF intervals and on the households' percentile distribution of the wealth tax for Spain. Banco de España defined eight strata which were oversampled progressively at higher rates. Strata 2 and 3 capture slightly less than half of the distribution of taxable wealth. Strata 4, 5 and 6 capture the third quartile except for the last percentile and a half approximately which is represented by the last two strata.

58. Finally, in Navarre and the Basque Country there was no oversampling of the wealthy because the national Tax Office does not hold the personal wealth tax file information for those regions.

Confidentiality guarantees

59. The Tax Office is subject to very stringent confidentiality requirements and cannot release, even to the Statistics Office, any personal tax information (not even in the form of intervals). To overcome the problem and enable wealth tax oversampling while preserving confidentiality, the National Tax Office volunteered to do the random sample selection itself following the sample design requirements, as instructed by the Bank of Spain and the National Statistics Office.

60. Thanks to the collaboration of both the Statistics Office and the Tax Office there is a unique population frame for the sampling. The population frame for the sample was the Continuous Municipal Census dated mid-2001, where the units are the households as defined by their address. With this information sent by the Statistics Office to the Tax Office, the Tax Office constructed for each address

three variables based on information from both the wealth and the income tax. These data were the starting point for the sampling.

61. The first variable, the wealth stratum indicator, is based on total declared taxable wealth for the household, which was obtained by adding up the returns of all its members when applicable. The second one, for those filing income tax but not wealth tax, is a variable indicating to which quartile in the national taxable income distribution the household belongs. Finally, information on the per capita income of the household was also added. The income variables were helpful in the selection of sample replacements (as we shall see below), and to ensure that households from all income levels were selected into the sample. The latter was obtained by using systematic sampling with random start in a properly ordered data frame.

62. Furthermore, the income quartile indicator was used to correct for non-response in large cities. The tax information available at the time was dated 1999. This entailed some limited mismatch between the two sources.

Refusal

63. As we can see in Figure 3 (included in page 4 of this note), there is a clear non-random component in cooperation rates [defined as completed/(completed+refused)], decreasing as we move up the wealth strata, ranging from 53.6% to 29.4%. It is clear from this pattern that overall cooperation or response rates are not very informative in case of oversampling since they are dependent on the degree of such oversampling. For some meaningful comparison, we constructed cooperation rates by strata for the 1992 SCF. There cooperation rates for the list sample ranged from 52.6% for stratum 1 to 20.1% for stratum 7.

Sampling in the French Enquête Patrimoine

64. This section is based on the documentation of the Enquête Patrimoine 2004.

65. In the Enquête Patrimoine 2004, 15,025 dwellings were sampled from the population of all dwellings from the 1999 census and the population of new dwellings (completed after the census). 9,692 households answered, for a total of 22,821 persons.

66. The survey attempts to oversample wealthy households and those with business wealth, by increasing the sample of white-collar ("cadres"), independent professionals, retired persons and households living in rich neighbourhoods. Compared to the other households, 4 times more households whose reference person was self-employed were sampled; households with cadres were oversampled 3 times.

Table 1: Oversampling coefficient

	Active self-employed	Active cadre	Retired	Other
Rich neighbourhood	4	3	3	2
Other neighbourhood	2	1.5	1.5	1

67. The rich neighbourhoods are defined as municipalities (communes) belonging to urban areas of 100,000 inhabitants or more, whose coordinate on the first axis in the Tabard typology of neighbourhoods is over the 75th percentile.

68. However the information used to oversample is only available for the dwellings found in the 1999 census, not on the new dwellings, whose sampling rate was uniform.

69. The procedure used in the sampling for the 2009 Enquête patrimoine will differ from the one shown here, as it will probably use a combination of wealth and income tax returns.

Sampling in the Cyprus Survey of Consumer Finances

70. This section is based on Karagrigoriou (2005) and Karagrigoriou and Michael (2001).

71. The 1999 CySCF sample consists of a randomly selected, nationally representative, area probability sample of all Cyprus households supplemented by a sample of high-income households. The area probability design is a multi-stage design that samples successively smaller geographic areas (counties, regions, municipalities, etc.). The basic rule of the area probability design is that the sample units at each stage of the design are selected with probability proportional to their size/population.

72. The list of households used for the selection of the samples for the Cyprus project is the 1999 list of customers of the Electricity Authority of Cyprus (EAC) that consists of approximately 270,000 customers/households. The average electricity consumption for the year 1998 is 400KW bimonthly. Due to the fact that the list consists of all summer houses and secondary residences, all residences consuming on the average at most 100KW bimonthly, were removed from the list. Approximately 13% of the households fall into this category, 3% of which had zero electricity consumption. The final household population from which the samples are selected is 230,000 households.

73. For the CySCF, the wealthy sample consists of a sample from the administrative records maintained by the Electricity Authority of Cyprus (EAC) that contain data coded from household electricity consumption. To obtain good measures of highly concentrated assets one should draw a sample of households with either high net worth or some other characteristic correlated with wealth. Although extensive data are collected on income as a product of the administration of federal tax returns in Cyprus, the reliability of such data is questionable. As a result it was decided to fill the wealthy sample with households having high consumption of electricity, due to the fact that the EAC list is the only available list that is both accurate in terms of measurements and complete in terms of household population coverage.

74. Note that the high consumption of electricity has been used in various Cyprus surveys as a criterion of wealth. In order to ensure the highest possible accuracy of the results for the CySCF, the EAC files were preferred to any other Authority’s files. Note that although the high consumption of electricity is not regarded as the best indicator of wealth, the files of Electricity Authorities around the world are considered as providing accurate and reliable household information. Note that in the United States, the Federal Reserve Board is using for this purpose the highly reliable tax files of the Internal Revenue Service (IRS).

References

O. Bover, “The Spanish survey of household finances (EFF): Description and methods of the 2002 wave”, *Occasional paper n°0409*, Banco de España, 2004.

D’Alessio and I. Faiella, “Non-response behaviour in the Bank of Italy’s Survey of Household Income and Wealth”, *Temi di discussione*, Banca d’Italia n°462, 2002.

L. D'Aurizio, I. Faiella, S. Iezzi, and A. Neri, “L'under-reporting della ricchezza finanziaria nell'indagine sui bilanci delle famiglie”, *Temi di discussione*, Banca d’Italia n°610, 2006.

J. Davies and A. Shorrocks, “The distribution of wealth”, in *Handbook of Income Distribution*, Handbook of Income Distribution, A.B. Atkinson & F. Bourguignon (ed.), Elsevier, 2000.

M. Frankel and A. Kennickell, “Toward the development of an optimal stratification paradigm for the survey of consumer finances”, 1995.

INSEE, « Enquête Patrimoine 2004 », *Sources et Méthodes*, http://www.insee.fr/fr/nom_def_met/sources/pdf/Patrimoine.pdf, 2007.

A. Karagrigoriou, “The Survey of Consumer Finances: Sampling and surveying in Cyprus”, http://www.econ.ucy.ac.cy/~echalias/survey.html, 2005.

A. Karagrigoriou and M. C. Michael, The Sample Design of the 1999 Cyprus Survey of Consumer Finances, http://www.econ.ucy.ac.cy/~echalias/survey.html, 2001.

A. Kennickell, “List Sample Design for the 1998 Survey of Consumer Finances”, 1998.

A. Kennickell, “Using Income Data to Predict Wealth”, 1999.

A. Kennickell, “Modeling Wealth with Multiple Observations of Income: Redesign of the Sample for the 2001 Survey of Consumer Finances”, 2001.

A. Kennickell, “The Good Shepherd: Sample Design and Control for Wealth Measurement in the Survey of Consumer Finances”, 2005.

A. Kennickell, “The Role of Over-sampling of the Wealthy in the Survey of Consumer Finances”, ISI 56th conference, 2007.

All papers by A. Kennickell except [2007] are available from http://www.federalreserve.gov/Pubs/oss/oss2/method.html