The Advantages of Random Sampling versus Cutting-of-the-Tail: the Application of a Stratified Sample Design for the Collection of Data on Special Financial Institutions in the Netherlands

Raymond Chaudron and Krit Carlier[1]

15 August 2014

1. Introduction

The Netherlands is home to a large number of Special Financial Institutions (SFIs). SFIs are subsidiaries of foreign companies with almost exclusively financial relations with other foreign entities. The main motives behind the presence of these companies in the Netherlands are fiscal and regulatory. Interest in the sector has heightened since the 2008 financial crisis, especially from those working on macro-prudential issues relating to shadow banking. As mentioned in the *Global Shadow Banking Monitoring Report 2013* of the Financial Stability Board, the Netherlands has one of the largest non-bank financial intermediaries sectors in the world, relative to GDP[2]. SFIs make up the majority of this. It is therefore very important to have statistical data about SFIs that is accurate and reliable.

In the Netherlands SFIs were traditionally surveyed using a two-tiered collection system. The largest SFIs were surveyed on a monthly basis, while the remaining SFIs reported only limited information in a census every other year. The selection of SFIs for monthly reporting was determined on the basis of a cut-off sample of the largest companies aimed at a predetermined coverage of total assets of roughly 95%. In between the two-yearly censuses, population totals were calculated by simple grossing up. This system was quite reliable at first, since the population was extremely skewed, not too large and fairly stable. However, the growth of SFIs since the middle 2000's in both number and size (see figure 1) led to an ever increasing number of monthly reporters, as well as an upward shifting cut-off threshold.

The increasing number of monthly reporters placed an increasing workload on the organisation to check the monthly data, whereas the upward shift in the threshold increased the non-observed part of the population, both in absolute and relative terms. At first, compilers at the central bank increased the number of monthly reporters (from 600 in 2004 to 1200 in 2008) in order to maintain the required coverage of 95%. Thereafter it became necessary to lower the coverage from 95% to 89% in 2012, in order to prevent the number of reporters of growly to unsustainable levels. By 2012 a coverage of 95% would have required more than 1800 monthly reporters. The increase in the number of reporters led to a deterioration in the quality of the reports due to insufficient capacity to check them, provide support to reporting institutions and monitor population developments. The latter being of utmost

[1] Both authors work as economists with the Statistics and Information Division of the Nederlandsche Bank. Corresponding author: R.F.D.D. Chaudron, De Nederlandsche Bank, Statistics and Information Division, P.O. Box 98, 1000 AB Amsterdam, The Netherlands. E-mail: r.f.d.d.chaudron@dnb.nl.

[2] See FSB (2013), Exhibit 3-1, page 11.

importance for this very volatile population. Ex-post revisions of the data on 2010 exceeded EUR 100 billion due to newly detected SFIs from the census. The high level of dynamism in the population of SFIs also meant that the grossing-up ratio became outdated more quickly. It was thus deemed necessary to increase the frequency of the two-yearly census to an annual census, putting additional strains on the organisation.

Figure 1: Size of the special financial institutions sector

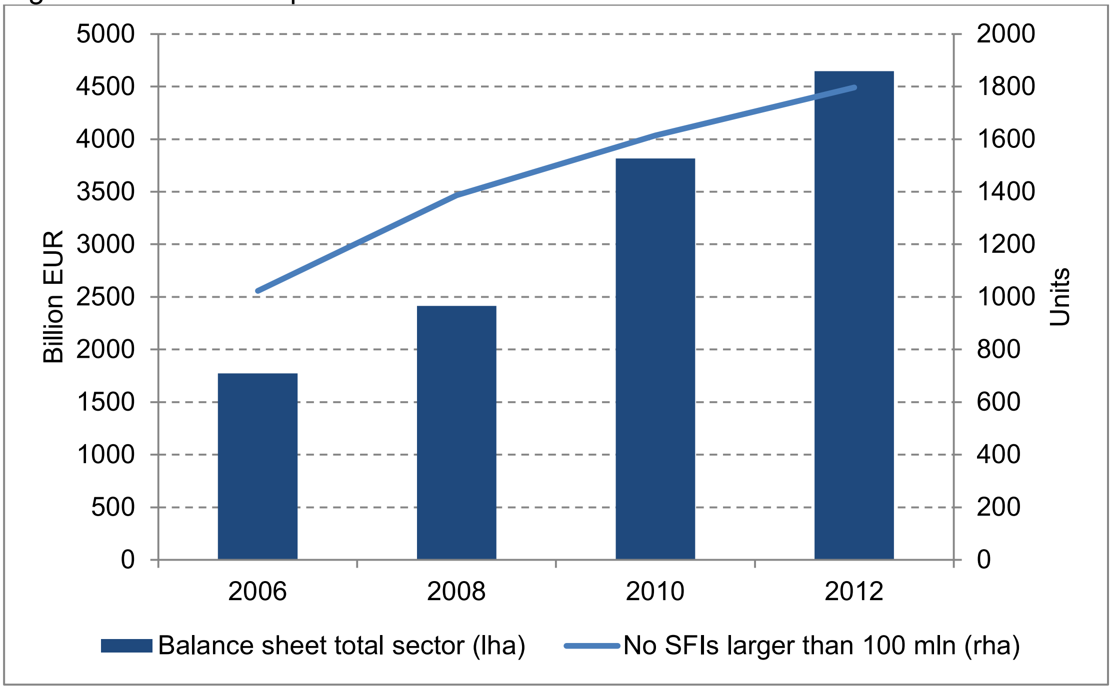

The above-mentioned problems prompted the need for a more efficient data collection system. In 2011 an investigation was initiated to explore the possibilities to move to a more sophisticated sample design incorporating random sampling in order to reduce the number of monthly reporters and increase quality. These analyses showed that the cut-off sample was introducing an increasingly large bias into the population estimates, particularly regarding individual items of the balance sheet. It also showed that a stratified sample design incorporating an include-all top stratum and an exclude-all tail, with a random sample for the mid-sized SFIs, could significantly lower the number of monthly reporters while improving the quality of the population estimates. By 2013 a sample design was chosen and implemented for the SFI sector. This paper elaborates on the rationale to move to this new sample design and its features.

The rest of the paper is structured as follows. The following section provides a brief description of the SFIs in the Netherlands as well as an indication of their size. Section 3 explains the limitations to cut-off sampling, while section 4 presents an outline of the sample design incorporating random sampling. Section 5 illustrates the results of the Monte Carlo simulations which were used to fine-tune the sample design. Section 6 ends with conclusions.

2. A brief description of SFIs in the Netherlands

Broadly speaking, SFIs can be classified into four types: holding company SFIs, financing company SFIs, Financial Vehicle Corporations (FVCs) for the securitisation of assets and other SFIs. These types of SFIs are distinguished according to the structure of their balance sheets[3]. Holding company SFIs primarily have equity participations in or by group companies on both sides of their balance sheet. They are usually part of a multinational's group structure in order to optimise the tax payments of the whole group. Financing company SFIs are set up to concentrate a conglomerate's external (wholesale) funding into one entity. Financing company SFIs obtain funding through either the issue of debt securities or through loans which is then distributed throughout the conglomerate. Both the assets and liabilities of financing SFIs are therefore dominated by borrowings. The balance sheet of SPVs contain the assets which are securitised (usually mortgages or commercial loans) on the asset side, financed by the issue of asset backed securities on the liabilities side. The group of other SFIs is fairly heterogeneous. Some are hybrids, mixtures of the first three, while some are set up for a specific purpose, such as the management of royalties or intra-company factoring.

Table 1: International investment position (IIP) of the Netherlands, 2013 (billion euro)

	Assets			Liabilities		
	SFIs	Other sectors	Total	SFIs	Other sectors	Total
Direct investments	2,814	757	3,571	2,438	498	2,936
Equity capital	2,014	524	2,565	2,000	294	2,294
Other capital	772	233	1,005	439	203	642
Portfolio investment	102	1,223	1,325	526	1,253	1,779
Equity securities	47	567	614	21	413	434
Debt securities	55	656	711	505	840	1.345
Financial derivatives	58	225	283	19	229	248
Other investment	139	726	865	128	706	834
Official reserves	0	34	34	-	-	-
Total	3,112	2,965	6,077	3,111	2,686	5,797

All SFIs have in common that they have no or only a few employees and that they have no relationship with the real economy of the Netherlands. Most small SFIs are managed by external management companies called trust offices. These companies should not be confused with the Anglo-Saxon trust company – which is a legal construct for the administration of financial assets for a group of beneficiaries. Dutch trust companies can be best described as corporate service providers, providing administrative, legal and compliance services. They are subject to supervision of the Dutch central bank under the Trust Offices Supervision Act ('Wet toezicht

[3] SFIs may also be classified according to the activity of the group to which they belong. For macro-prudential purposes, the Nederlandsche Bank uses a distinction between financial groups and other groups.

trustkantoren'). SFIs themselves are not supervised. SFIs which are part of financial conglomerates together with SPVs are considered to be an important part of the shadow banking sector in the Netherlands. Statistical data on SFIs was originally collected only for the Dutch balance of payments, but is used increasingly as a stand-alone dataset since the interest in shadow-banking activities has surged after the financial crisis, for instance in the annual FSB *Global Shadow Banking Monitoring Report*. Table 1 provides an overview of the size of the SFIs' financial position in relation to the international investment position of the Netherlands. As the data in the table clearly shows, SFIs make up roughly half of all cross-border assets and liabilities of the Netherlands.

3. Limitations to cut-off sampling

A cut-off sample design works well if the population is both stable and homogenous. A stable population does not show many new institutions which appear (births) or disappear (deaths) in between the updates, nor shows many fundamental changes of entities themselves. In a stable population the risk of missing out on important developments concerning out-of-sample entities is therefore minimal. In turn, in a homogenous population institutions vary by size, but not by other characteristics. In other words, a cut-off sample design assumes the largest institutions to be representative for the small institutions in every respect. Their balance sheet structure is comparable and they grow or contract in the same way as the small institutions. Obviously, once these assumptions are less realistic the cut-off sample design will lead to biased estimates of the population as a whole. As such biases most severely affect the breakdown of the estimates, whether these matter or not depends on user demands. The more statistical demands of users concern certain breakdowns, the more important the aspect of homogeneity becomes.

Surveys of financial institutions carried out by central banks have traditionally used a cut-off sample design. In such a design, institutions are ordered by size (for instance total assets) and the largest institutions, representing a certain percentage of the population total (usually 90 or 95%) for the size measure are selected for regular (e.g. monthly or quarterly) reporting. The whole population is enumerated using a less frequent census, on the basis of which the cut-off sample and the grossing-up ratios are updated. This sample design is efficient in terms of sample size and reporting burden, as most financial sectors exhibit a strongly skewed population distribution. In financial sectors such as banks, pension funds and insurance companies, a small proportion of the largest institutions represent a disproportionally large fraction of the total assets of the whole sector. Furthermore, in most countries, the financial sectors consist of a few dozen or a few hundred institutions at the most and are fairly stable over time. Advantages of the cut-off sample design in these cases are the simplicity of selecting reporters, the concentration of reporting burden with the largest institutions and the procedure to gross-up the observed figures to population totals. The data from the monthly reporting institutions is simply multiplied with the inverse of the proportion of the size measure represented by the largest institutions in the sample.

Like other financial sectors, the population SFIs is highly skewed. Figure 2 on the next page shows that in 2008 the largest 1000 SFIs covered just over 90% of the population total assets. After the 1000 largest however, the additional coverage of each company quickly declines. However, unlike other financial sectors the number of SFIs involved is quite large: by the latest count there were 14.366 SFIs in the Netherlands (2013) and the population has also become very volatile. For instance, between 2008 and 2010 about 2,500 new entities were established of which more than 100 with total assets in excess of EUR 1 billion. During the same period 1,700 entities disappeared, of which 40 were larger than EUR 1 billion. Next to the high incidence of 'births' and 'deaths', entities can also mushroom in size or reversely shrink to limited proportions from one period to the other.

In fact, SFIs do not form a homogenous population, even when the most basic characteristic is considered, the composition of the balance sheet. As explained in the introduction, there are four basic types of SFIs based on the structure of the balance sheet. This structure largely depends on the size of the SFI. As illustrated by figure 3, large SFIs are financed with securities and equity from group companies, whereas the smaller SFIs are predominantly financed through intercompany loans. Since the large SFIs are not representative for the smaller ones, accurate estimates for the whole population can not be provided through a cut-off sample but require sampling of all size groups. As statistical demands of users of SFI data, for instance for the analysis of financial stability, are increasingly interested in certain breakdowns, this lack of homogeneity clearly necessitates an unbiased sample design.

Figure 2: Cumulative coverage of total assets of companies ordered by size (2008)

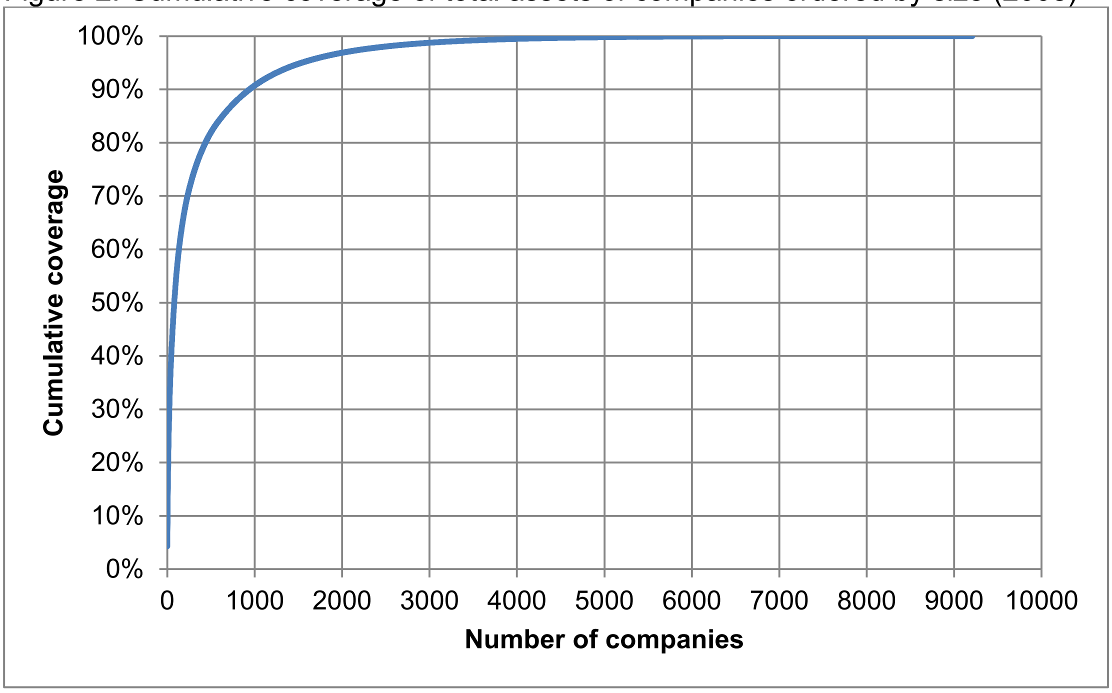

Figure 3: Composition of the liabilities of SFIs by size group

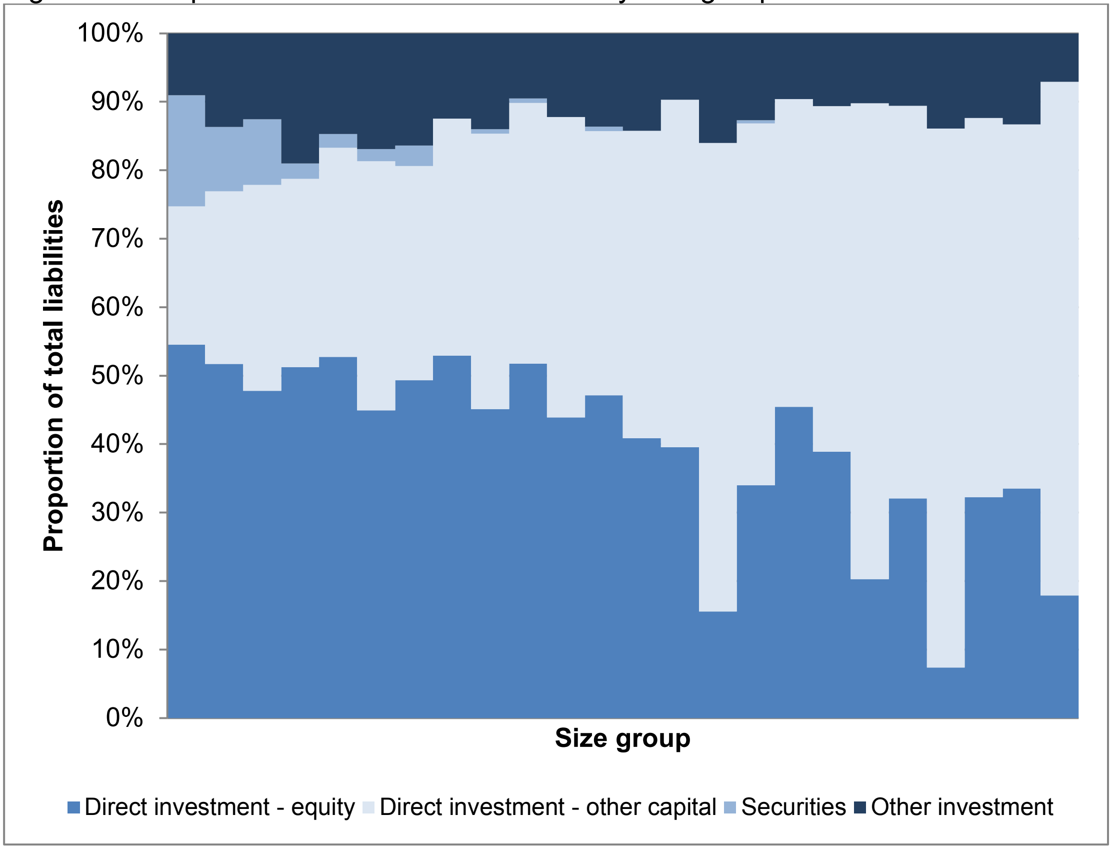

Finally, another drawback regarding the cut-off design is that it is very difficult to provide a measure of accuracy (such as a confidence interval) for the estimates of the population totals. As the selection of the reporting population is non-random (but systematic) it is impossible to estimate sampling errors. By contrast, a sample design which incorporates a random selection of specific groups of reporters does allow for the calculation of statistical confidence measures. Furthermore, as we show later, random selection enables to fine-tune the sample design using stochastic techniques such as Monte-Carlo simulations.

4. Outline of the new sample design

The aim of choosing the optimal sample design was to combine the efficiency in terms of sample size and reporting burden of a cut-off sample with the advantages of a stochastic sample design. Through this strategy one can still benefit from the skewness of the population, which is further documented in tables 2 and 3.

Combining the efficiency of a cut-off sample with the superior estimation of a random sample is possible by using a stratified sample design in which a top stratum is defined from which all reporters are selected and a second stratum from which reporters are selected through a random process. This sample design follows the basic set-up presented in Benedetti et al. (2010). Assuming a given number of reporters in a cut-off design and in a stratified design, the number of reporters in the top stratum will of course be smaller in this new sample design than in a cut-off sample to accommodate sampling in the sample stratum. In addition to a top stratum and a sample stratum, one can also define a tail stratum from which the reporters – whose contribution to population totals is very small - are completely exempt from reporting. This is also a cut-off in a sense and can be employed to relieve the very smallest of institutions from reporting, at the same time limiting the sample size. This cut-off stratum can be kept small enough to be estimated using a relatively naive estimator, such as a fixed amount or a proportion of the total, analogous to the cut-off sample design.

Table 2: Size distribution of the SFI population over time, number of entities

Size class	2006		2008		2010		2012	
(million EUR)	Number	Share (%)	Number	Share (%)	Number	Share (%)	Number	Share (%)
> 10.000	31	0.4%	39	0.3%	75	0.6%	97	0.8%
1.000 – 10.000	265	3.4%	345	3.1%	504	4.0%	560	4.4%
100 – 1.000	1,023	13.2%	1,386	12.4%	1,614	12.9%	1,796	14.0%
10 – 100	1,914	24.7%	2,407	21.6%	2,761	22.0%	2,915	22.7%
< 10	4,519	58.3%	5,070	45.4%	5,291	42.2%	5,102	39.8%
Liquidated			1,909	17.1%	2,283	18.2%	2,344	18.3%
Total	7,752	100.0%	11,156	100.0%	12,528	100.0%	12,814	100.0%

Table 3: Size distribution of the SFI population over time, total aggregate assets

Size class	2006		2008		2010		2012	
(million EUR)	Aggregate assets	Share (%)	Aggregate assets	Share (%)	Aggregate assets	Share (%)	Aggregate assets	Share (%)
> 10.000	638,035	34.5%	943,210	37.6%	1,746,970	44.5%	2,354,968	49.4%
1.000 – 10.000	800,500	43.2%	1,019,109	40.6%	1,498,698	38.2%	1,669,676	35.1%
100 – 1.000	335,269	18.1%	452,657	18.0%	569,336	14.5%	621,677	13.1%
10 – 100	68,878	3.7%	87,458	3.5%	100,181	2.6%	107,655	2.3%
< 10	9,003	0.5%	9,373	0.4%	11,016	0.3%	9,650	0.2%
Liquidated			-	0.0%	-	0.0%	-	0.0%
Total	1,851,685	100.0%	2,511,807	100.0%	3,926,200	100.0%	4,763,626	100.0%

The three strata are often denoted by the take-all or the certainly include stratum (CI), the take-some or random sample stratum (S) and the take-none or certainly exclude stratum (CE), where the sample stratum can be further subdivided into two or more substrata with different sampling proportions. Stratification can be based on a single criterion (variable), but also on combinations of two or more criteria. The main objective in stratification is to divide the population in homogenous groups. A common practice is to use the target variable(s) itself as auxiliary variable, provided that data on this variable is available for the entire population, for instance from the most recent census. In fact, it is best to choose stratification variables that are highly correlated with the target variables for the most accurate estimates. Since total assets have high serial correlation, lagged total assets (or other balance sheet items) make good stratification variables.

In the simplest sample design each entity in the sample stratum has an equal chance of being selected. Such a design can be described as follows. If *Y* denotes the target variable to be estimated for the whole of the population, N_a is the total number of SFIs in stratum a and n_a is the number of SFIs sampled from stratum a (where CI, S and CE indicate the certainly include, sample and certainly exclude strata), then the estimate of the population sum of *Y* is as follows:

$$[1]\quad \hat{Y} = \left(\sum_{i=1}^{N_{CI}} Y_i\right) + \left(\sum_{j=1}^{n_S} Y_j \frac{N_s}{n_s}\right) + f(N_{CE})$$

The first term of the right hand side of this equation is simply the sum of the observations from the companies in the certainly-include stratum. This sum is of course completely non-stochastic since all entities in this stratum are sampled. The second part of the equation is the sum of observations from the companies in the sample stratum multiplied with the inverse of the sample proportion. This is the only stochastic part of the estimate. The error of the population estimate will almost exclusively depend on the variance of this estimator. The final part, which is not further specified here, is a term to denote the estimate for the companies in the certainly-exclude stratum.

It is of course possible to employ a different sampling procedure than simple random sampling for the sample stratum, in which the chance of being selected is not identical for all entities within a stratum but for instance depends on the size of the entity (sampling with a probability proportional to size, PPS).

5. Choosing the optimal sampling parameters

A major advantage of stochastic sample designs over cut-off sampling is that the quality of the sample design can be objectively compared based on quantitative measures, notably by comparing the mean squared error of the estimates (MSE) to the actual realisation. Using Monte-Carlo simulations for different values of the sample design parameters, including the choice of stratification variables and stratum boundaries, we were able to fine-tune the parameters in order to produce the best population estimates. In order to find the optimal boundaries between the different strata (the certainly include stratum, the take-some or random sample stratum and the take-none or certainly exclude stratum) Monte Carlo simulations were performed for many combinations of these two boundaries, while keeping the total number of reporters fixed.

The boundary of the certainly-include stratum was varied between 40% and 80% of the population total for assets, with intervals of 5%. The certainly-exclude stratum was varied between 2% and 10% with 1% intervals. By performing 1000 simulations with sample size n divided between the certainly include and the sample strata for each of the combinations of the boundaries for the certainly-include and certainly-exclude stratum, we could obtain measures for the accuracy of each variation of the sample design.

Figure 3 presents the MSE for the population estimate of total assets. The MSE is shown to initially decline with an increase of the certainly-include boundary, but roughly from the 65% mark[4] it starts to increase with a further increase of this boundary. This pivot point results from the fact that more reporters in de CI stratum implies less reporters in the sample stratum (when keeping the total sample size constant).The figure shows a reverse pattern regarding the certainly-exclude boundary: the MSE decreases with an increase of the boundary which indicates that too low a boundary includes volatility introduced by a very limited number (low chance of selection) of small SFIs growing very fast. Similar analyses have been performed for individual balance sheet items and the geographical breakdown, showing comparable results. The shape of the figure also demonstrates that the optimal sample design includes all three strata. As the certainly-include stratum exceeds 65%, the MSE starts to increase again. This effect is caused by the fact that there are too few reporters from the sample stratum. In the end, the optimal stratification chosen was for a 65% certainly-include stratum and a 4% certainly-exclude stratum.

Figure 4: Mean squared error of the population sum of assets

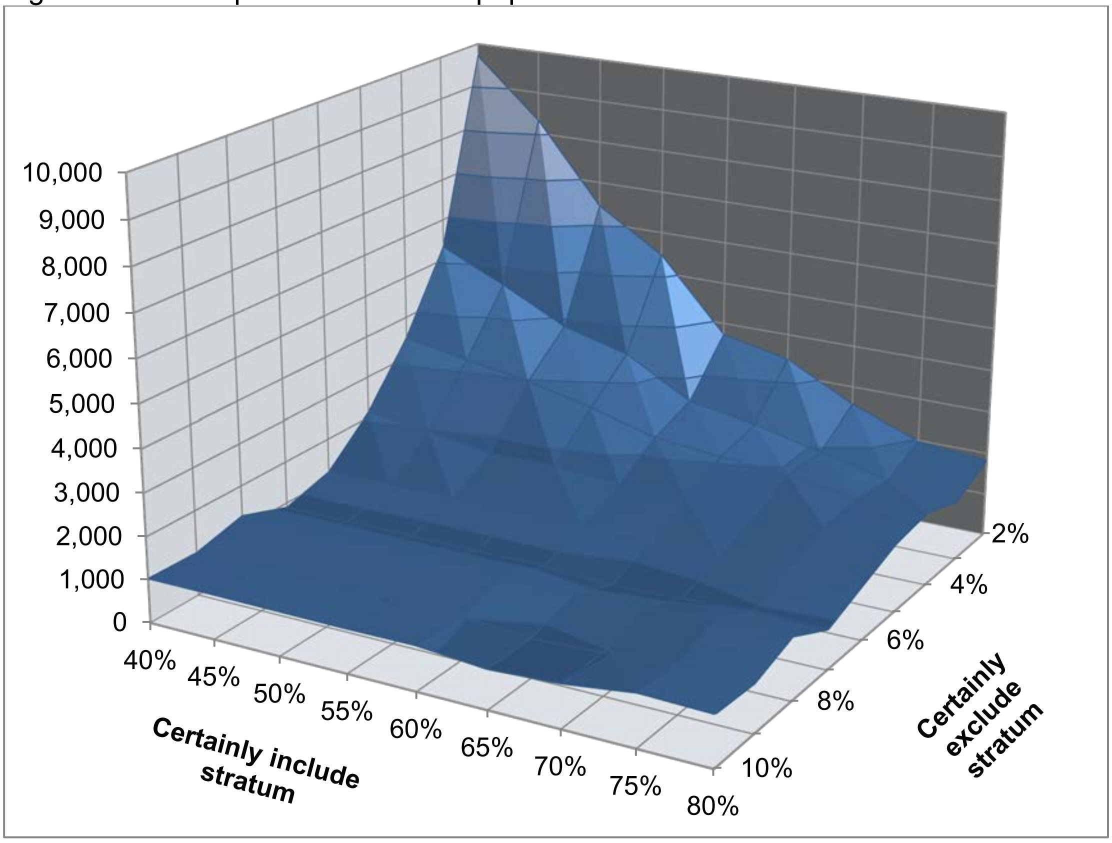

After the boundaries of the certainly-include and certainly-exclude had been roughly determined, the required size of the reporting population – certainly-include stratum

[4] At which 65% of the population's balance sheet total is covered by the CI stratum.

plus sample stratum – needed to be defined. For this purpose again Monte Carlo simulations were performed, the boundaries set as defined above (at 65% and 4% respectively), but for different sizes of the reporting population. From this a downward sloping curve results (more reporters always mean higher quality estimates) but the slope becomes increasingly flat with each additional reporter (implying that each additional reporter has a lower effect on the quality than the previous one). This curve provides information of the quality of the estimate with different numbers of reporters and thus allows setting the size of the reporting population providing the quality required. In the case of the SFIs, a total sample size of around 850 was considered sufficient. Additional reporters were shown not to increase the quality substantially. Compared to the pure cut-off sample design, this entailed reducing the number of reporters with approximately one-third compared to the cut-off sample used previously.
The new sample design thus offers two clear benefits. Firstly, the new approach has made the accuracy an explicit criterion in the sample design. This enables statisticians to engage in an informed discussion with users and management about the priorities and the resources devoted to the survey. It also allows statisticians to determine the limits for which the data can be used. Secondly, the new sample design is shown to provide better quality estimates compared to a cut-off sample (of equal size). In the case of SFIs it also allowed a reduction of the sample size significantly from around 1400 to 850. This in turn made it possible to devote more time and effort to every individual reporter and to the maintenance of the business register and the annual census, thus improving the overall quality of the statistics even more.

6. Conclusions

The main point we want to make here is that the use of stochastic sample designs – as opposed to the purely systematic selection of companies through for instance a cut-off sample of the largest units – can significantly increase the efficiency of a data collection system as well as the quality of estimates of both aggregates and breakdowns. Additionally, in contrast to systematic collection, stochastic sample designs allow quantifying the quality of the estimates with varying sample sizes, so that the sample size can be set in an informed manner. As a result, stochastic sample designs can be used to limit the sample size considerably as well as improving output quality (reporting burden, though less a concern since the financial crisis, still also plays an important role in the background), as we have illustrated for the case of Dutch SFI statistics.

We also want to mention that besides the choice of sample design, there are many other practical issues that have to be resolved in putting a new sample design into practice. One important issue is that of sample rotation. By renewing the random sample each year, companies risk having to report on and off. In order to limit this side-effect (which is difficult to understand from the perspective of the reporter), we have chosen to renew only one third of the sample each year using a system with so-called ‘permanent random numbers’ (see Ohlson, 1995). Monte Carlo simulations have shown that this did not significantly worsen the accuracy of the estimates. Another practical issue to be dealt with is that of maintaining the business register from which the sample is drawn and the treatment of 'births and deaths'. By limiting the number of monthly reporters, we were able to free up resources for the maintenance of the register, which is a vital prerequisite for reliable samples.

Acknowledgements

This research has benefitted greatly from the work of Frans Bookholt, who worked at the Nederlandsche Bank as an intern during the summer and fall of 2013 to perform the Monte-Carlo simulations presented in this article.

References

Benedetti, R., M. Bee, and G. Espa (2010), A Framework for Cut-off Sampling in Business Survey Design, *Journal of Official Statistics*, 26(4), pp. 651-671.

Financial Stability Board (2013), Global Shadow Banking Monitoring Report 2013, Basel.

Ohlsson, E. (1995). Coordination of Samples Using Permanent Random Numbers. In Cox, B. et al. (eds.): Business Survey Methods. New York: Wiley, 153-169.