

Aspects of Sample Allocation in Business Surveys

Gareth James, Mark Pont and Markus Sova

Office for National Statistics, Government Buildings, Cardiff Road, NEWPORT, NP10 8XG, UK.

Gareth.James@ons.gov.uk, Mark.Pont@ons.gov.uk, Markus.Sova@ons.gov.uk

Abstract

The UK's Office for National Statistics runs business surveys using stratified designs; Neyman optimal allocation is usually used to allocate the sample. However, the designs are often constrained for various reasons with the consequence that the resulting allocations are sub-optimal.

This paper draws on some allocation issues that arose during two recent survey redesigns in the ONS - the Retail Sales Inquiry (RSI) and the Vacancy Survey. In particular, we consider the effect on precision, respondent load for small businesses and bias for ratio estimators of:

- imposing minimum and maximum stratum sample sizes;
- allocations to minimise the variances of estimated population totals or month-to-month changes; and
- the use of expansion (number-raised) or ratio estimation.

The paper also discusses the conflicts that can occur when more than one constraint is applied. The empirical findings from these surveys have implications for future redesigns or re-allocations. The chosen allocations are explained, and the process of checking and comparing allocations individually for overall plausibility shows that deciding on a sample allocation is much more involved than simply plugging numbers into a formula.

1 Introduction

1.1 Office for National Statistics business surveys

The Office for National Statistics (ONS) conducts many surveys of businesses in the UK. The surveys periodically come up for review, and these give the chance to examine the sample design, allocation, estimation methodology and so on. Though the idea of allocating a sample may seem trivial, to the point of saying that allocation using the standard Neyman formula is all that is needed, there are many more practical considerations and constraints on the design which need to be addressed.

Some background to business survey sampling at ONS is given here, as it helps to put the rest of the paper into context. Most surveys have a stratified design, and the strata are usually defined by two variables: employment (in size bands) and industrial classification. One size band is usually reserved for those businesses with employment less than ten, as we have special rules for keeping the burden on these businesses low. At the other end, surveys often have a cut-off employment size, above which all businesses are sampled - complete enumeration (CE); the cut-off varies from survey to survey. Standard Industrial Classification, 2003 - previously 1992 (SIC(92)) - is used to classify businesses in terms of their main economic activity. SIC is a hierarchical coding system, and one level of the hierarchy (or sometimes an amalgamation of industries) is chosen for stratification of business surveys.

The most-often used sampling frame at ONS for business surveys is the Inter-Departmental Business Register (IDBR). ONS is responsible for maintenance of the frame, which contains details of about 1.9 million businesses, accounting for about 99 per cent of economic activity in the UK. Rotational sampling is used for businesses not in the CE bands, using a permanent random number system. (See, for example, Co-ordination of Samples Using Permanent Random Numbers, by E. Ohlsson in Business Survey Methods.)

Neyman allocation is usually used to allocate the sample in a business survey; the strata for which have been previously defined. However, with desires to limit the burden on small businesses, bias in estimates if using a ratio estimator with small samples, potentially low response, several outputs from the one survey, and domain estimation, there are many other considerations in choosing an allocation than just the overall standard error of one estimate.

1.2 The RSI and Vacancy Survey - some introductory notes

Sample allocations are discussed here with reference to two monthly ONS surveys, which have recently been redesigned. The first is the monthly Retail Sales Inquiry (RSI), from which the Retail Sales Index is constructed. The sample covers 27 different industries (retail sectors), and each is split into four employment size bands (0-9, 10-19, 20-99 and 100+), making 108 strata in total. In each industry the top size band is completely enumerated. The total sample size is 5000.

A review of the RSI took place during 2002, and was part of a regular review of ONS surveys. As one part of this, the sample allocation was examined, as the then current one was believed to have become non-optimal. Other parts of the review examined the estimation procedure (including a comparison of using the ratio approach instead of matched pairs to measure change), imputation procedures and so on. The remit of the redesign did not include redefinition of stratum boundaries, and the CE strata were to remain. In investigation of the allocation, the standard deviation of the change in monthly retail sales was examined, with averaged, smoothed and modelled estimates produced for use in optimal allocation. Other constraints were imposed, and these are described in detail in the section 2.

The other survey to be discussed here is the Vacancy Survey, and a relatively brief review of its sample allocation, which was conducted in late 2001. The survey provides an estimate of the number of job vacancies in the economy and is now published as a monthly series as part of National Statistics. When the review was carried out, however, the survey wasn't being published as it was still at the trial stage.

The survey had, initially, been stratified by 29 industries and ten employment size bands, though a few industries had had some size bands collapsed due to small numbers. The top size band in each industry is completely enumerated, whereas the others are sampled. Each sampled business is included in the sample every third month until being rotated out of the sample. There are, thus, three waves of sampled businesses in the survey. Businesses that are in the completely enumerated band are included in the sample every month. In total, the sample size at the time of the review was a little less than the original size of 6000 due to decay. Estimation of the total number of vacancies in the economy was being made using a ratio estimator with registered employment (from the IDBR) as the auxiliary variable.

The review was carried out after about five months' data had been collected. This gave the opportunity to examine the allocation in light of survey responses. Other parts of the review concentrated on the form of the stratification, and the type of estimator used.

The main part of the review was concerned with the sample allocation. The remit included proposing a re-allocation of the full sample of size 6000, and investigation into reducing the number of size bands by merging (collapsing) some together (though not redefining the boundaries of the current bands).

2 Sample Allocations

2.1 Constraints on the allocation

Naturally, when making an allocation, if no constraints are imposed at all, an allocation can be found that is optimal. In other words, one allocation can be found that will minimise the standard error of a chosen estimator. However, at ONS, a number of constraints are imposed on allocations. Firstly, most surveys have CE strata. These are usually the strata containing businesses with greatest employment, as these tend to contribute the most to the overall variance of the survey. Once the number of businesses in each CE stratum has been deducted from the overall sample size, the remaining sample can be allocated optimally over the sampled strata.

Some other aspects of allocation constraints are now explored, and the effects they have had on the allocations of the Vacancy Survey and the RSI explored.

2.2 Setting limits on the allocation by stratum

The allocation in a particular stratum is sometimes controlled. Both minimum and maximum sample size restrictions can be imposed, together with extra rules for those strata where the optimum allocation would not meet these conditions. On the one hand, a maximum sampling fraction is sometimes imposed. In many surveys this is set at 0.5, i.e. half the population. The reason for doing so lies in the rotation of contributors from the survey. Sampling no more than half the contributors each month ensures that each contributor is out of the sample for at least as long a time as it is in the sample.

On the other hand, a minimum stratum size may be introduced. The reasons for doing this are essentially two fold. The first is an attempt to overcome the issue of non-response, by trying to ensure that some contributors will be present in the sample from each domain for which an estimate is derived. The second relates to estimation: often a ratio estimator is used, or an

estimate of a ratio may be calculated. It is known that the estimator is biased in small samples (see, for example, Sampling Techniques by Cochran), and ensuring a minimum sample size helps to reduce the potential bias.

2.3 The RSI reallocation

Provisional results from the RSI are first published about 15 days after the end of the reporting period (month). Revised figures are published one month later, with the final figures one month after that. It is the first release, that of provisional results, that makes the headlines, and could be considered the most important. An allocation was made to try to ensure a minimum numbers of responses in each stratum at the time of the provisional results; this varied from stratum to stratum depending on both the size of the stratum, and the response rate (an average figure was used for this). This is a new approach for ONS surveys.

The effects of imposing different minimum return numbers, by stratum, were investigated for sizes between zero (i.e. allocation without restriction other than CE strata) to 30, which virtually defined the whole allocation. For each minimum size investigated, a stratum could have a free allocation, or be restricted as


- being completely enumerated, either because it is the top size band, or the minimum number of returns is greater than the population size;
- taking the minimum size necessary to achieve the minimum number of returns; or
- is bounded by half the population size.

The number of such "free" strata (out of 108) is an indicator of the extent of constraint on the allocation. The results, below, show how the freedom in the allocation has been virtually removed when at least 30 returns are required in each stratum:

Minimum return size	0	5	10	15	20	25	30
No. of free strata (of 108). NB: 27 are CE anyway.	81	45	36	33	20	12	3

Naturally, as the minimum stratum size increases, the allocation in the "free" strata - which tend to be the smallest size band - falls. For example, the size when the minimum return size is 30, the allocations in the "free" strata are only about 30 per cent of the size when there is no minimum.


Of course, the effects on bias and standard error are important. The graph below shows the overall standard error of the estimate of monthly change in total retail sales, based upon optimum allocation with the minimum number of returns at the provisional results stage imposed on the allocation. Note that the units of the standard error is %, as the change is measured in percentage points. The disadvantages of imposing larger minimum stratum sizes can clearly be seen. Comparison should be made to the standard error with no minimum stratum size, as this is optimal.



The expression for the bias was not formulated in full, due to time constraints. In its place, the leading term (except for a constant) was investigated, as giving an indication of the likely bias. That term is

$$\sqrt{\frac{1-f}{n}}$$

where f is the sampling fraction and n the sample size. This value of this term is plotted below against minimum stratum size; again comparison should be made to the value when the minimum stratum size is zero. This plot shows the benefit of increasing the minimum stratum size.



The decision on what to recommend as a minimum stratum size came down to a compromise between standard error and bias. One of the most reassuring outcomes from the work was that there were minimum sizes (any in the range 5 to 25, in fact) which would result in gains in both standard error and bias on the current allocation. However, the final decision also involved consideration of what the allocation actually looked like, and what it would mean in practice for each of the strata and the businesses that comprise them. For example, greatly increasing the sample size (to meet minimum stratum requirements) in the fish sales industries when their overall contribution to retail sales is proportionately extremely small seems intuitively wrong, especially when considering the increased burden that would be placed upon these businesses. After much discussion, 15 was recommended, and the sample has now been reallocated.

2.4 The Vacancy Survey

2.4.1 Allocation The simplest way to re-allocate the sample was within the ten size bands already defined. Neyman allocation was used with the constraints of completely enumerated bands and a minimum stratum size (allocation) of five. The standard deviations, by stratum were estimated as the square root of the average of the five monthly variances then available. The resulting allocation was generally similar to that which had previously been defined, though there were some strata which changed notably.

The suggested reallocation resulted in standard errors of the total that were about, on average, 8 per cent lower than before. Further investigation showed that the overall gain in precision could be attributed to two sources of improvement: roughly one-third came from the overall increase in sample size (to 6000, again) and two-thirds from the re-allocation. The standard errors in the tables below were calculated with the expected response rates in mind.

	Estimate of total, 000s	Standard error, 000s (c.v., %)	
		Current allocation	Re-allocation
April 2001	659	17.3 (2.6)	16.4 (2.5)
May 2001	682	22.5 (3.3)	21.7 (3.2)
June 2001	689	17.8 (2.6)	16.7 (2.4)
July 2001	667	18.7 (2.8)	17.0 (2.6)
August 2001	647	19.7 (3.0)	16.0 (2.5)
Average	669	19.2 (2.9)	17.5 (2.6)

A rolling quarterly average is also constructed. The expected gains in standard error of these estimates were also calculated:

	Estimate of total, 000s	Standard error, 000s (c.v., %)	
		Current allocation	Re-allocation
April-June 2001	677	11.0 (1.6)	10.5 (1.5)
May-July 2001	679	11.2 (1.6)	10.5 (1.6)
June-August 2001	668	10.6 (1.6)	9.4 (1.4)

2.4.2 Collapsing size bands One of the main aims of the review was to reduce the number of size bands. Merging sizebands creates larger strata. A major benefit to come from an increase in stratum sample size is the corresponding decrease in weights, for the lower/lowest of the bands in each collapsed stratum. It is the smallest businesses which tend to have fewest vacancies, and for any such business a typical monthly series of the number of vacancies might look like: ..., 0, 0, 0, 0, 2, 0, 0, 1, 1, 0, 0, 0, 1, The larger the weights applied to this series, the more erratic the series. This is the main justification for reducing the number employment bands, though the simplifications in running the survey, and compiling the results are also beneficial. The table below shows average sample sizes in each employment band, over all industry classes, for various different band combinations:

Employment sizeband	4	5	9	10	10 (Current)
1: 0-4	61	70	77	48	47
2: 5-9				22	22
3: 10-19	39	31	17	19	17
4: 20-49			14	15	16
5: 50-99		21	9	10	9
6: 100-249	51		11	12	12
7: 250-499		29	7	7	8
8: 500-999			9	10	8
9: 1000-2499			8	9	6
10: 2500+	CE	CE	CE	CE	CE

The CE strata have an average size of 56.

Options for collapsing size bands were investigated by removing some of the stratum boundaries and pooling standard deviations. Naturally, there are various ways in which ten sizebands can be collapsed to form five bands, say, and alternatives were investigated. As expected, the actual way in which the collapsing was done had little effect on the standard errors, but rather it was the number of sizebands which had a far greater impact. The table below shows the effects on standard errors of collapsing sizebands and then reallocating the sample:

Standard error of the total	Re-allocation of 6000								Current allocation
	3	4	5	6	7	8	9	10	10
Number of bands									
April 2001	64.4	22.1	20.5	20.0	19.6	18.9	18.5	16.4	17.3
May 2001	83.0	26.0	22.3	21.2	21.1	20.8	20.6	21.7	22.5
June 2001	72.7	23.3	19.7	19.1	18.9	18.5	19.0	16.7	17.8
July 2001	74.2	23.0	20.0	19.9	19.4	18.9	18.1	17.0	18.7
August 2001	104.4	24.9	19.6	19.2	18.2	17.9	17.6	16.0	19.7
<i>Average</i>	<i>79.8</i>	<i>23.9</i>	<i>20.4</i>	<i>19.9</i>	<i>19.4</i>	<i>19.0</i>	<i>18.6</i>	<i>17.5</i>	<i>19.2</i>

The average sample sizes in each new stratum (over all industries) were also investigated, these are shown below, for a number of possible sizebands and ways of collapsing strata.

The results suggest that having only three bands is much worse than having four, itself a little worse than five, and after that there is little to be gained by having more. Thus, our recommendation was to collapse sizebands to form just five. The exact way in which the five bands would be created was chosen to minimise the respondent burden on the smallest businesses, though there was little to choose between them.

2.4.3 Method of estimation The final part of the review was to examine how much better the use of ratio estimation (with employment as the auxiliary variable) is when compared to number-raised estimation. The results suggested that ratio estimation should continue to be used, and switching to number-raised estimation would incur an increase in the standard error of the quarterly estimate of about 11 per cent.

Since the review, the survey has developed further and is now being published. The recommendations to collapse the number of size bands to five and reallocate the sample have been accepted and will be implemented as soon as resources permit.

3 Concluding remarks

Having to reallocate samples has shown how the process is not as simple as it might, at first, appear. There are many ways to go, practical issues to consider, and a blanket approach may not be the best one. The need for someone to check the outcome each time, and not to just use the output from a program has been shown. It is essential that the allocation really does make sense, and is the suitable for the job.

4 References

- Ohlsson, E. 1995. Co-ordination of Samples Using Permanent Random Numbers, in *Business Survey Methods & Sampling* (Cox, *et al*), Wiley.
- Cochran, W.G. 1977. *Sampling Techniques*, Wiley.