

NatCen

Social Research that works for society

The effects of selecting multiple respondents per household for a survey of people in paid work

A statistical and cost assessment

Acknowledgements: This research was by funded by the Advisory, Conciliation and Arbitration Service (ACAS) and undertaken as part of the Skills and Employment Survey research programme based at WISERD, Cardiff University.

Authors: Susan Purdon (Bryson Purdon Social Research) and Martin Wood (National Centre for Social Research)

Date: May 2022

At NatCen Social Research we believe that social research has the power to make life better. By really understanding the complexity of people's lives and what they think about the issues that affect them, we give the public a powerful and influential role in shaping decisions and services that can make a difference to everyone. And as an independent, not for profit organisation we're able to put all our time and energy into delivering social research that works for society.

NatCen Social Research
35 Northampton Square
London EC1V 0AX
T 020 7250 1866
www.natcen.ac.uk

A Company Limited by Guarantee
Registered in England No.4392418.
A Charity registered in England and Wales (1091768) and Scotland (SC038454)
This project was carried out in compliance with ISO20252

Contents

	The effects of selecting multiple respondents per household for a survey of people in paid work	1
1	Background.....	1
	1.1 Selecting all adults per household	1
	1.2 Selecting one adult per household.....	2
	1.3 Selecting up to two adults per households.....	2
2	Effect of the different designs on standard errors	3
3	Potential change in non-response bias	5
4	Costs 6	
5	Conclusions	6

The effects of selecting multiple respondents per household for a survey of people in paid work

For PAF based surveys such as the Skills and Employment Survey a decision is needed on the number of working adults to select per household. In the past the SES has always selected just one, with the selection being at random from all working adults in the household. But other options are to select all working adults, or to apply a cap and select up to, say, two. The decision is not straightforward, but we have conducted analysis, described in this note, that suggests that selecting more than one adult per household would be a sensible change for the SES.

In addition to the statistical benefits of such a change, selecting multiple adults per household should provide fieldwork efficiencies that will enable an increase in the number of interviews that can be achieved within a fixed budget. We provide an estimate of the scale of the increase possible.

A further consideration for a long-running study with a valuable timeseries is that of comparability over time. There remains a concern that the nature of non-response bias is affected by a change in the selection approach. This is difficult to assess ahead of fieldwork. However, we argue this uncertainty is set against the more significant risks to changing response bias presented by a context of falling response across surveys in general which the post-pandemic situation may have exacerbated.

1 Background

There are advantages and disadvantages to all of the different within-household sampling approaches.

1.1 Selecting all adults per household

The disadvantages to selecting all adults per household are clustering effects and burden.

The usual reason given for avoiding selecting all adults per household is that it leads to within-household clustering effects in the data. These clustering effects occur when the responses given by adults from the same household are correlated. For instance, if adults in the same household share similar attitudes and opinions, or do similar jobs, their responses will be correlated. The clustering effects have the effect of increasing standard errors and confidence intervals for the survey statistics. The increases are greatest for survey variables that are the most highly correlated within households.

In addition, selecting all adults in a household will lead to an increased burden on larger households. This can lead to concerns that the response rate will fall as a consequence, although the evidence on this is hard to come by. Analysis of the Labour Force Survey provides the distribution of the number of working adults per household – Table 1 below. This suggests that for around 12% of households where there is a working adult, three or more adults would be asked to take part.

Table 1: Number of working adults per household, LFS Q4 2019

Number of working adults	% of households
1	29.5
2	58.5
3	9.4
4	2.4
5	0.3
6	0.0
Total	100.0

However, a very key advantage of selecting all eligible adults within a household is that it significantly reduces the number of households that need to be selected for a survey to achieve a target number of adult interviews. This reduces survey costs.

Another marked advantage of selecting all adults relates to the fact that, although clustering effects for 'all adult' statistics are a potential problem, as soon as the data is split by gender, most households will have just one responding adult and the clustering effects largely disappear.

1.2 Selecting one adult per household

To avoid within-household clustering effects, the standard approach is to select just one adult per household. However, although this does avoid the clustering effects, this approach still leads to increases in standard errors and confidence intervals because the data now has to be weighted to account for the within-household selection. (Adults from households with just one eligible adult get a weight of one; the selected adults from households with two eligible adults get a weight of two; and so on.) Furthermore, these weights apply across all analyses, including by gender, so the effect on standard errors persist across all statistics (whereas, as noted above, within-household clustering effects tend to largely disappear for within-gender analysis).

The second major disadvantage of selecting just one adult is that it leads to larger survey costs. To meet a fixed target sample of adults the exact same number of households have to be recruited. That is, the number of households recruited to the survey has to be much larger than with within-household clustering. This increases costs.

The main advantage cited for selecting one adult is generally that it reduces the survey burden on households. There is also an assumption made quite frequently that the effect of the weighting to account for the selection of one adult on standard errors is less damaging than the effect of within-household clustering. The analysis we set out below challenges that assumption.

1.3 Selecting up to two adults per households

A compromise between these two approaches is to select more than one eligible adult per household, but to cap the number selected at two (or perhaps three). This leads to some clustering effects, and the selection within the larger households means that some weighting is still needed, so there is still inflation of standard errors.

This approach has similar advantages to the ‘select all adult’ design, but ensures that the burden on larger households is reduced. In terms of cost it is between the ‘select all adult’ and ‘select one adult’ designs, but closer to the former than the latter.

2 Effect of the different designs on standard errors

To reach a decision on the best approach, the natural first step is to establish which of the designs leads to the least inflation of standard errors (this being measured via ‘design factors’ – see below). If within-household clustering leads to smaller design factors than weighting then, because within-household clustering is less costly, this would be a strong argument for moving to a model where more than one adult is selected per household.

To test this, we have mimicked the different approaches using Labour Force Survey data (restricted to working adults). That survey collects data on all adults within a household so is naturally clustered. To simulate the effect of selecting just one adult, one adult has been selected at random for each household with at least two eligible adults, and the data weighted accordingly. Likewise, to simulate the effect of selecting two adults, two adults have been selected at random within household with three or more eligible adults and the data weighted. This gives three different datasets to compare.

We have then estimated the ‘design factor’ for a range of survey variables from the LFS that are similar to the variables that will be collected as part of the SES. The ‘design factor’ represents the degree to which the standard error, and hence confidence interval, is inflated, relative to a simple random sample, because of the sample design. For example, if the achieved sample size was 1,000, for a survey percentage of 50%, a simple random sample would give a standard error of 1.58 (and a 95% confidence interval of (47%, 53%)). If the design factor was 1.2 for a particular design, then the standard error for a sample of 1,000 would be $1.58 \times 1.2 = 1.90$, with a 95% confidence interval of (46.3%, 53.7%).

Table 2 below sets out the estimated design factors for the three designs for a broad range of statistics. In each case the smallest design factor is shown in green, the middle one in amber, the largest one in red.

All else being equal, the design with the most green cells is to be preferred. Looking down the table it is clear that this suggests that selecting all eligible adults is the optimal of the three designs, and that selecting just one adult is the least effective. Selecting up to two eligible adults falls between these two.

The natural conclusion to draw from this is that the SES should move to a design where either all eligible adults are selected, and accept that this might be quite burdensome for the largest households. Otherwise, up to two adults should be selected so as to reduce this burden, but accepting that this will lead to slightly larger design factors.

Table 2: Design factors estimated from the LFS for the three sample designs

	Estimate	Select all eligible adults	Select up to two eligible adults	Select one eligible adult
		DEFT	DEFT	DEFT
NS-SEC				
Higher managerial and professional	18.2%	1.172	1.169	1.173
Lower managerial and professional	28.9%	1.107	1.117	1.155
Intermediate occupations	12.8%	1.058	1.080	1.143
Small employers and own account workers	10.4%	1.129	1.132	1.122
Lower supervisory and technical	7.3%	1.097	1.109	1.194
Semi-routine occupations	11.7%	1.097	1.129	1.145
Routine occupations	8.4%	1.125	1.160	1.170
Never worked, unemployed, and nec	2.3%	1.229	1.308	1.322
Industry section in main job				
A Agriculture, forestry and fishing	0.9%	1.149	1.145	1.111
B Mining and quarrying	0.4%	1.083	1.094	1.160
C Manufacturing	9.4%	1.095	1.108	1.138
D Electricity, gas, air cond supply	0.6%	1.101	1.087	1.067
E Water supply, sewerage, waste	0.7%	1.070	1.082	1.107
F Construction	7.1%	1.117	1.142	1.193
G Wholesale, retail, repair of vehicles	11.9%	1.131	1.148	1.167
H Transport and storage	4.8%	1.113	1.130	1.140
I Accommodation and food services	4.8%	1.207	1.247	1.224
J Information and communication	4.8%	1.259	1.258	1.228
K Financial and insurance activities	4.0%	1.158	1.183	1.224
L Real estate activities	1.1%	1.102	1.089	1.075
M Prof, scientific, technical activ.	7.9%	1.145	1.138	1.154
N Admin and support services	4.6%	1.140	1.143	1.170
O Public admin and defence	6.7%	1.117	1.139	1.144
P Education	10.5%	1.098	1.101	1.120
Q Health and social work	14.0%	1.117	1.123	1.134
R Arts, entertainment and recreation	2.5%	1.194	1.197	1.241
S Other service activities	2.8%	1.079	1.099	1.153
T Households as employers	0.2%	1.226	1.268	1.192

U Extraterritorial organisations	0.2%	1.238	1.218	1.211
Managerial status				
Manager	26.2%	1.133	1.135	1.143
Foreman or supervisor	11.5%	1.082	1.091	1.135
Not manager or supervisor	62.3%	1.117	1.128	1.150
Number of employees at workplace				
1-10	18.2%	1.125	1.134	1.142
11-19	7.7%	1.091	1.100	1.126
20-24	4.3%	1.110	1.132	1.172
Don't know but under 25	2.2%	1.171	1.221	1.336
25-49	12.8%	1.094	1.125	1.168
50-249	23.6%	1.122	1.116	1.144
250-499	6.9%	1.137	1.158	1.206
Don't know but between 50 and 499	3.9%	1.128	1.153	1.183
500 or more	20.3%	1.174	1.175	1.168
Respondent works flexi-time				
Yes	12.8%	1.154	1.129	1.156
Respondent has zero hours contract				
Yes	2.5%	1.130	1.174	1.230
Job-related training or education in the last three months				
Yes	25.3%	1.151	1.159	1.152
Length of time with current employer				
Less than 3 months	3.5%	1.154	1.198	1.249
3 months, less than 6	4.2%	1.143	1.172	1.273
6 months, less than 12	6.6%	1.127	1.156	1.217
1 year, less than 2	11.5%	1.114	1.147	1.218
2 years, less than 5	23.6%	1.122	1.144	1.189
5 years, less than 10	18.8%	1.097	1.110	1.152
10 years, less than 20	20.7%	1.085	1.087	1.092
20 years or more	11.1%	1.057	1.065	1.071

3 Potential change in non-response bias

There remains a concern that the nature of non-response bias is affected by a change in the selection approach. The perceived higher burden to the household where two people are to be interviewed may lead to higher household non-response on the one hand and additional non-response within households where the second adult decides not to participate having observed the interview with the first (although it is also possible that the second adult could be encouraged to participate by seeing the first take part). The possible implication were there to be substantial differences in non-response patterns is different biases that are not observable/correctable could be introduced, leading to a reduction of comparability on some estimates with previous surveys.

This is difficult to assess ahead of fieldwork – there are few surveys that have been through this specific change in survey design and in any case it is likely that the context of the specific survey would be important in non-response patterns (interview topic saliency, interview length).

However, our contention is that the impact on response of the change in approach does not seem likely to be substantial, and in any case is just one area where comparability over time is at risk. Many National Statistics household surveys, including the LFS, have seen declines in response rates over the last two decades and the public's response to the pandemic is likely to have changed the nature of non-response to in-home surveys still further.

The potential gains to statistical efficiency, particularly where it facilitates a higher number of interviews overall, seems worth pursuing. The worst of the impact of a multi-person design would be avoided by selecting a maximum of two people per household.

4 Costs

The process for fieldwork in previous designs of the SES has been for field interviewers to screen for adults aged 20 to 65 in paid employment at selected addresses, and then conduct a random selection of one individual.

LFS data for Q4 2019 suggests that 67% of households will have at least one eligible adult for the study. In practice, the proportion will be lower due to non-working households being more likely to agree to the screening process (driven by retired households). This is therefore a relatively costly design already compared with surveys where all households are eligible.

An estimated 32% of households have just a single eligible person in the household, with the remaining 35% containing two or more eligible people. This means that for 100 selected households, a design where a single eligible adult is selected would result in 67 eligible adults from whom to try to obtain an interview. In a design where 2 eligible people can be selected per household, the number of eligible adults is 102.

Consideration however needs to be given to how within-household dynamics might affect response rates at the individual level where a 2-person design is implemented: due to perceived burden on the household, it is possible that the second individual will be less likely to take part than the first.

5 Conclusions

Our analysis of LFS data suggests that a switch to selecting multiple eligible adults per household would lead to improved statistical efficiency for the SES, with design factors reduced on relevant variables of interest. The benefits remain when selecting a maximum of two adults, and this would to some extent deal with concerns about changes in non-response patterns. Such a change in the design could facilitate significant gains in the number of interviews that could be achieved within a given budget.