

Technical details

In 2015, the sample for the British Social Attitudes survey was split into four equally-sized portions. Each portion was asked a different version of the questionnaire (versions A, B, C and D). Depending on the number of versions in which it was included, questions were thus asked either of the full sample (4,328 respondents) or of a random quarter, half or three-quarters of the sample.

Sample design

The British Social Attitudes survey is designed to yield a representative sample of adults aged 18 or over. Since 1993, the sampling frame for the survey has been the Postcode Address File (PAF), a list of addresses (or postal delivery points) compiled by the Post Office.¹

For practical reasons, the sample is confined to those living in private households. People living in institutions (though not in private households at such institutions) are excluded, as are households whose addresses were not on the PAF.

The sampling method involved a multi-stage design, with three separate stages of selection.

Selection of sectors

At the first stage, postcode sectors were selected systematically from a list of all postal sectors in Britain. Before selection, any sectors with fewer than 500 addresses were identified and grouped together with an adjacent sector; in Scotland all sectors north of the Caledonian Canal were excluded (because of the prohibitive costs of interviewing there). Sectors were then stratified on the basis of: 37 sub-regions; population density, (population in private households/ area of the postal sector in hectares), with variable banding used in order to create three equal-sized strata per sub-region; and ranking by percentage of homes that were owner-occupied.

This resulted in the selection of 362 postcode sectors, with probability proportional to the number of addresses in each sector.

Selection of addresses

Twenty-six addresses were selected in each of the 362 sectors or groups of sectors. The issued sample was therefore $362 \times 26 = 9,412$ addresses, selected by starting from a random point on the list of addresses for each sector, and choosing each address at a fixed interval. The fixed interval was calculated for each sector in order to generate the correct number of addresses.

The Multiple-Occupancy Indicator (MOI) available through PAF was used when selecting addresses in Scotland. The MOI shows the

number of accommodation spaces sharing one address. Thus, if the MOI indicated more than one accommodation space at a given address, the chances of the given address being selected from the list of addresses would increase so that it matched the total number of accommodation spaces. The MOI is largely irrelevant in England and Wales, as separate dwelling units (DUs) generally appear as separate entries on the PAF. In Scotland, tenements with many flats tend to appear as one entry on the PAF. However, even in Scotland, the vast majority (99.8%) of MOIs in the sample had a value of one. The remainder had MOIs greater than one. The MOI affects the selection probability of the address, so it was necessary to incorporate an adjustment for this into the weighting procedures (described below).

Selection of individuals

Interviewers called at each address selected from the PAF and listed all those eligible for inclusion in the British Social Attitudes sample – that is, all persons currently aged 18 or over and resident at the selected address. The interviewer then selected one respondent using a computer-generated random selection procedure. Where there were two or more DUs at the selected address, interviewers first had to select one DU using the same random procedure. They then followed the same procedure to select a person for interview within the selected DU.

Weighting

The weights for the British Social Attitudes survey correct for the unequal selection of addresses, DUs and individuals, and for biases caused by differential non-response. The different stages of the weighting scheme are outlined in detail below.

Selection weights

Selection weights are required because not all the units covered in the survey had the same probability of selection. The weighting reflects the relative selection probabilities of the individual at the three main stages of selection: address, DU and individual. First, because addresses in Scotland were selected using the MOI, weights were needed to compensate for the greater probability of an address with an MOI of more than one being selected, compared with an address with an MOI of one. (This stage was omitted for the English and Welsh data.) Secondly, data were weighted to compensate for the fact that a DU at an address that contained a large number of DUs was less likely to be selected for inclusion in the survey than a DU at an address that contained fewer DUs. (We used this procedure because in most cases where the MOI is greater than one, the two stages will cancel each other out, resulting in more efficient weights.) Thirdly, data were weighted to compensate for the lower selection probabilities of adults living in large households, compared with those in small households.

At each stage the selection weights were trimmed to avoid a small number of very high or very low weights in the sample; such weights would inflate standard errors, reducing the precision of the survey estimates and causing the weighted sample to be less efficient. Less than 1% of the selection weights were trimmed at each stage.

Non-response model

It is known that certain subgroups in the population are more likely to respond to surveys than others. These groups can end up over-represented in the sample, which can bias the survey estimates. Where information is available about non-responding households, the response behaviour of the sample members can be modelled and the results used to generate a non-response weight. This non-response weight is intended to reduce bias in the sample resulting from differential response to the survey.

The data was modelled using logistic regression, with the dependent variable indicating whether or not the selected individual responded to the survey. Ineligible households² were not included in the non-response modelling. A number of area-level and interviewer observation variables were used to model response. Not all the variables examined were retained for the final model: variables not strongly related to a household's propensity to respond were dropped from the analysis.

The variables found to be related to response were: region, the relative condition of the immediate local area, the relative condition of the address, whether there were entry barriers to the selected address, percentage of owner occupiers, and population density. The model shows that response increases if there are no barriers to entry (for instance, if there are no locked gates around the address and no entry phone) and if the general condition of the address is the same or better than other addresses in the area. Response decreases if the relative condition of the immediate surrounding area is mainly good or fair, and decreases as population density increases. Response is also higher for addresses in the North, the Midlands and the South East. The full model is given in Table A.1.

Table A.1 The final non-response model

Variable	B	S.E.	Wald	Df	Sig.	Odds
Region			80.171	11	.000	
Inner London (baseline)						
North East	.553	.160	11.955	1	.001	1.738
North West	.627	.135	21.433	1	.000	1.871
Yorkshire and The Humber	.273	.143	3.640	1	.056	1.314
East Midlands	.786	.145	29.344	1	.000	2.195
West Midlands	.500	.141	12.495	1	.000	1.648
East of England	.394	.146	7.313	1	.007	1.482
Outer London	.309	.140	4.894	1	.027	1.362
South East	.600	.134	19.955	1	.000	1.822
South West	.443	.147	9.141	1	.002	1.558
Wales	.236	.159	2.211	1	.137	1.266
Scotland	.119	.142	0.706	1	.401	1.127
Barriers to address						
No barriers (baseline)						
One or more	-.483	.088	30.085	1	.000	.617
Relative condition of the local area			18.430	2	.000	
Mainly good (baseline)						
Mainly fair	-.160	.050	10.359	1	.001	.852
Mainly bad or very bad	-.462	.128	13.021	1	.000	.630
Relative condition of the address			38.411	2	0.000	
Better (baseline)						
About the same	-.488	.085	32.695	1	.000	.614
Worse	-.660	.118	31.207	1	.000	.517
Percentage owner-occupied (in quintiles)			12.350	4	0.015	
1 lowest (constant)						
2	-.137	.076	3.233	1	.072	.872
3	-.023	.080	.082	1	.774	.977
4	-.215	.086	6.323	1	.012	.807
5 highest	-.210	.087	5.778	1	.016	.811
Population density (population in private households/ area of postcode sector in hectares)						
	-.086	.024	13.141	1	.000	.917
Constant	.473	.179	6.968	1	.008	1.604

The response is 1 = individual responding to the survey, 0 = non-response

For consistency with previous years, region was also included in the non-response model. As region is used in the calibration weighting, including it in the non-response model can be beneficial in reducing the variance of the final weights

The model R^2 is 0.024 (Cox and Snell)

B is the estimate coefficient with standard error **S.E.**

The **Wald**-test measures the impact of the categorical variable on the model with the appropriate number of degrees of freedom (**df**). If the test is significant (**sig.** < 0.05), then the categorical variable is considered to be 'significantly associated' with the response variable and therefore included in the model

The non-response weight was calculated as the inverse of the predicted response probabilities saved from the logistic regression model. The non-response weight was then combined with the selection weights to create the final non-response weight. The top 1% of the weight were trimmed before the weight was scaled to the achieved sample size (resulting in the weight being standardised around an average of one).

Calibration weighting

The final stage of weighting was to adjust the final non-response weight so that the weighted sample matched the population in terms of age, sex and region.

Table A.2 Weighted and unweighted sample distribution, by region, age and sex

	Population	Unweighted respondents	Respondent weighted by selection weight only	Respondent un-calibrated non-response and selection weight	Respondent weighted by final weight
	%	%	%	%	%
Region					
North East	4.2	4.6	4.4	4.3	4.2
North West	11.3	12.8	12.7	11.7	11.3
Yorkshire and The Humber	8.5	7.9	7.4	8.1	8.5
East Midlands	7.4	10.1	10.0	8.6	7.4
West Midlands	9.0	9.0	9.1	9.0	9.0
East of England	9.6	8.2	8.6	8.6	9.6
London	13.4	9.9	10.3	12.2	13.4
South East	14.1	16.3	16.6	15.4	14.1
South West	8.8	8.7	8.4	8.2	8.8
Wales	5.0	5.0	4.8	5.2	5.0
Scotland	8.7	7.7	7.5	8.7	8.7
Age & sex					
M 18–24	5.9	3.1	4.4	4.4	5.9
M 25–34	8.6	5.9	6.3	6.5	8.6
M 35–44	8.2	7.1	7.3	7.2	8.2
M 45–54	8.8	8.0	8.4	8.4	8.8
M 55–59	3.7	3.0	3.3	3.3	3.7
M 60–64	3.4	4.0	4.1	4.0	3.4
M 65+	10.1	12.8	11.7	11.5	10.1
F 18–24	5.7	3.6	4.3	4.4	5.7
F 25–34	8.6	8.8	8.5	8.8	8.6
F 35–44	8.3	9.7	10.0	10.0	8.3
F 45–54	9.1	9.6	10.2	10.2	9.1
F 55–59	3.8	4.9	4.9	4.9	3.8
F 60–64	3.5	4.5	4.4	4.2	3.5
F 65+	12.3	14.9	12.2	12.1	12.3
<i>Base</i>	<i>49,501,761</i>	<i>4328</i>	<i>4328</i>	<i>4328</i>	<i>4328</i>

Only adults aged 18 or over are eligible to take part in the survey, therefore the data have been weighted to the British population aged 18+ based on 2011 Census data from the Office for National Statistics/General Register Office for Scotland.

The survey data were weighted to the marginal age/sex and GOR distributions using raking-ratio (or rim) weighting. As a result, the weighted data should exactly match the population across these three dimensions. This is shown in Table A.2.

The calibration weight is the final non-response weight to be used in the analysis of the 2015 survey; this weight has been scaled to the responding sample size. The range of the weights is given in Table A.3.

	N	Minimum	Mean	Maximum
DU and person selection weight	4328	.56	1.00	2.23
Un-calibrated non-response and selection weight	4328	.38	1.00	3.47
Final calibrated non-response weight	4328	.27	1.00	4.91

Effective sample size

The effect of the sample design on the precision of survey estimates is indicated by the effective sample size (neff). The effective sample size measures the size of an (unweighted) simple random sample that would achieve the same precision (standard error) as the design being implemented. If the effective sample size is close to the actual sample size, then we have an efficient design with a good level of precision. The lower the effective sample size is, the lower the level of precision. The efficiency of a sample is given by the ratio of the effective sample size to the actual sample size. Samples that select one person per household tend to have lower efficiency than samples that select all household members. The final calibrated non-response weights have an effective sample size (neff) of 3,486 and efficiency of 81%.

Weighted bases

All the percentages presented in this report are based on weighted data. Only unweighted bases are presented in the tables. Details of weighted bases for standard demographics are shown in Table A.4.

Table A.4 Weighted bases for standard demographics, 2015

	Weighted base	Unweighted base
Sex		
Male	2105	1904
Female	2223	2424
Age		
18-24	500	290
25-34	743	635
35-44	712	728
45-54	776	764
55-59	326	344
60-64	299	367
65+	966	1193
Ethnicity		
White	3746	3861
Non-White	529	424
Religion		
No religion	2097	2067
Church of England/Anglican	736	850
Other Christian	740	758
Non-Christian	364	275
Roman Catholic	375	364
Class group (NSSEC)		
Managerial & professional occupations	1617	1619
Intermediate occupations	495	532
Employers in small org; own account workers	383	380
Lower supervisory & technical occupations	404	392
Semi-routine & routine occupations	1213	1241
Highest educational qualification		
Degree	1033	980
Higher education below degree	474	487
A level or equivalent	832	743
O level or equivalent	804	794
CSE or equivalent	307	318
Foreign or other	101	94
No qualification	766	903
Marital status		
Married	2120	1905
Living as married	500	401
Separated or divorced after marrying	402	623
Widowed	276	440
Not married	1027	957

Questionnaire versions

Each address in each sector (sampling point) was allocated to one of the portions of the sample: A, B, C or D. As mentioned earlier, a different version of the questionnaire was used with each of the four sample portions. If one serial number was version A, the next was version B, the next version C, and the fourth version D. There were 2,353 issued addresses for each of the four versions of the sample.

Fieldwork

Interviewing was mainly carried out between August and October 2015, with a small number of interviews taking place in November 2015.

Fieldwork was conducted by interviewers drawn from NatCen Social Research's regular panel and conducted using face-to-face computer-assisted interviewing.³ Interviewers attended a one-day briefing conference to familiarise them with the selection procedures and questionnaires, with the exception of experienced interviewers who completed a self-briefing containing updates to the questionnaire and procedures.

The mean interview length was 67 minutes for version A of the questionnaire, 62 minutes for version B, 65 minutes for version C, and 64 minutes for version D.⁴ Interviewers achieved an overall response rate of between 51.0 and 51.4%. Details are shown in Table A.5.

Table A.5 Response rate¹ on British Social Attitudes, 2015

	Number	Lower limit of response (%)	Upper limit of response (%)
Addresses issued	9412		
Out of scope	923		
Upper limit of eligible cases	8489	100	
Uncertain eligibility	61	0.7	
Lower limit of eligible cases	8428		100
Interview achieved	4328	51.0	51.4
With self-completion	3670	43.2	43.5
Interview not achieved	4100	48.3	48.6
Refused ²	3148	37.1	37.4
Non-contacted ³	448	5.3	5.3
Other non-response	503	5.9	6.0

1 Response is calculated as a range from a lower limit where all unknown eligibility cases (for example, address inaccessible, or unknown whether address is residential) are assumed to be eligible and therefore included in the unproductive outcomes, to an upper limit where all these cases are assumed to be ineligible and therefore excluded from the response calculation

2 'Refused' comprises refusals before selection of an individual at the address, refusals to the office, refusal by the selected person, 'proxy' refusals (on behalf of the selected respondent) and broken appointments after which the selected person could not be recontacted

3 'Non-contacted' comprises households where no one was contacted and those where the selected person could not be contacted

As in earlier rounds of the series, the respondent was asked to fill in a self-completion questionnaire which, whenever possible, was collected by the interviewer. Otherwise, the respondent was asked to post it to NatCen Social Research.

A total of 658 respondents (15% of those interviewed) did not return their self-completion questionnaire. Version A of the self-completion questionnaire was returned by 86% of respondents to the face-to-face interview, version B of the questionnaire was returned by 84%, version C by 86%, and version D by 84%. As in previous rounds, we judged that it was not necessary to apply additional weights to correct for non-response to the self-completion questionnaire.

Advance letter

Advance letters describing the purpose of the survey and the coverage of the questionnaire, were sent to sampled addresses before the interviewer made their first call.⁵

Analysis variables

A number of standard analyses have been used in the tables that appear in this report. The analysis groups requiring further definition are set out below. For further details see Stafford and Thomson (2006). Where relevant the name given to the relevant analysis variable is shown in square brackets – for example [REarn].

Region

The dataset is classified by 12 regions, formerly the Government Office Regions.

Standard Occupational Classification

Respondents are classified according to their own occupation, not that of the 'head of household'. Each respondent was asked about their current or last job, so that all respondents except those who had never worked were coded. Additionally, all job details were collected for all spouses and partners in work.

Since the 2011 survey, we have coded occupation to the Standard Occupational Classification 2010 (SOC 2010) instead of the Standard Occupational Classification 2000 (SOC 2000). The main socio-economic grouping based on SOC 2010 is the National Statistics Socio-Economic Classification (NS-SEC). However, to maintain time-series, some analysis has continued to use the older schemes based on SOC 90 – Registrar General's Social Class and Socio-Economic Group – though these are now derived from SOC 2000 (which is derived from SOC 2010).

National Statistics Socio-Economic Classification (NS-SEC)

The combination of SOC 2010 and employment status for current or last job generates the following NS-SEC analytic classes:

- Employers in large organisations, higher managerial and professional
- Lower professional and managerial; higher technical and supervisory
- Intermediate occupations
- Small employers and own account workers
- Lower supervisory and technical occupations
- Semi-routine occupations
- Routine occupations

The remaining respondents are grouped as “never had a job” or “not classifiable”. For some analyses, it may be more appropriate to classify respondents according to their current socio-economic status, which takes into account only their present economic position. In this case, in addition to the seven classes listed above, the remaining respondents not currently in paid work fall into one of the following categories: “not classifiable”, “retired”, “looking after the home”, “unemployed” or “others not in paid occupations”.

Registrar General’s Social Class

As with NS-SEC, each respondent’s social class is based on his or her current or last occupation. The combination of SOC 90 with employment status for current or last job generates the following six social classes:

I		Professional etc. occupations	}	‘Non-manual’
II		Managerial and technical occupations		
III	(Non-manual)	Skilled occupations		
III	(Manual)	Skilled occupations	}	‘Manual’
IV		Partly skilled occupations		
V		Unskilled occupations		

They are usually collapsed into four groups: I & II, III Non-manual, III Manual, and IV & V.

Socio-Economic Group

As with NS-SEC, each respondent’s Socio-Economic Group (SEG) is based on his or her current or last occupation. SEG aims to bring together people with jobs of similar social and economic status, and is derived from a combination of employment status and occupation. The full SEG classification identifies 18 categories, but these are usually condensed into six groups:

- Professionals, employers and managers
- Intermediate non-manual workers
- Junior non-manual workers
- Skilled manual workers
- Semi-skilled manual workers
- Unskilled manual workers

As with NS-SEC, the remaining respondents are grouped as “never had a job” or “not classifiable”.

Industry

All respondents whose occupation could be coded were allocated a Standard Industrial Classification 2007 (SIC 07). Two-digit class codes are used. As with social class, SIC may be generated on the basis of the respondent’s current occupation only, or on his or her most recently classifiable occupation.

Party identification

Respondents can be classified as identifying with a particular political party on one of three counts: if they consider themselves supporters of that party, closer to it than to others, or more likely to support it in the event of a general election. The three groups are generally described respectively as ‘partisans’, ‘sympathisers’ and ‘residual identifiers’. In combination, the three groups are referred to as ‘identifiers’. Responses are derived from the following questions:

Generally speaking, do you think of yourself as a supporter of any one political party? [Yes/No]

[If “No”/“Don’t know”]

Do you think of yourself as a little closer to one political party than to the others? [Yes/No]

[If “Yes” at either question or “No”/“Don’t know” at 2nd question]

Which one?/If there were a general election tomorrow, which political party do you think you would be most likely to support?

[Conservative; Labour; Liberal Democrat; Scottish National Party; Plaid Cymru; Green Party; UK Independence Party (UKIP)/Veritas; British National Party (BNP)/National Front; RESPECT/Scottish Socialist Party (SSP)/Socialist Party; Other party; Other answer; None; Refused to say]

Income

Two variables classify the respondent’s earnings [REarn] and household income [HHInc] The bandings used are designed to

be representative of those that exist in Britain and are taken from the Family Resources Survey (see <http://research.dwp.gov.uk/asd/frs/>). Four derived variables give income deciles/quartiles: [REarnD], [REarnQ], [HHIncD] and [HHIncQ]. Deciles and quartiles are calculated based on individual earnings and household incomes in Britain as a whole.

Attitude scales

Since 1986, the British Social Attitudes surveys have included two attitude scales which aim to measure where respondents stand on certain underlying value dimensions – left–right and libertarian–authoritarian.⁶ Since 1987 (except in 1990), a similar scale on ‘welfarism’ has also been included. Some of the items in the welfarism scale were changed in 2000–2001. The current version of the scale is shown below.

A useful way of summarising the information from a number of questions of this sort is to construct an additive index (Spector, 1992; DeVellis, 2003). This approach rests on the assumption that there is an underlying – ‘latent’ – attitudinal dimension which characterises the answers to all the questions within each scale. If so, scores on the index are likely to be a more reliable indication of the underlying attitude than the answers to any one question.

Each of these scales consists of a number of statements to which the respondent is invited to “agree strongly”, “agree”, “neither agree nor disagree”, “disagree” or “disagree strongly”.

The items are:

Left–right scale

Government should redistribute income from the better off to those who are less well off [Redistrb]

Big business benefits owners at the expense of workers [BigBusnN]

Ordinary working people do not get their fair share of the nation’s wealth [Wealth]7

There is one law for the rich and one for the poor [RichLaw]

Management will always try to get the better of employees if it gets the chance [Indust4]

Libertarian–authoritarian scale

Young people today don’t have enough respect for traditional British values. [TradVals]

People who break the law should be given stiffer sentences. [StifSent]

For some crimes, the death penalty is the most appropriate sentence. [DeathApp]

Schools should teach children to obey authority. [Obey]

The law should always be obeyed, even if a particular law is wrong. [WrongLaw]

Censorship of films and magazines is necessary to uphold moral standards. [Censor]

Welfarism scale

The welfare state encourages people to stop helping each other. [WelfHelp]

The government should spend more money on welfare benefits for the poor, even if it leads to higher taxes. [MoreWelf]

Around here, most unemployed people could find a job if they really wanted one. [UnempJob]

Many people who get social security don't really deserve any help. [SocHelp]

Most people on the dole are fiddling in one way or another. [DoleFidl]

If welfare benefits weren't so generous, people would learn to stand on their own two feet. [WelfFeet]

Cutting welfare benefits would damage too many people's lives. [DamLives]

The creation of the welfare state is one of Britain's proudest achievements. [ProudWif]

The indices for the three scales are formed by scoring the leftmost, most libertarian or most pro-welfare position, as 1 and the rightmost, most authoritarian or most anti-welfarist position, as 5. The “neither agree nor disagree” option is scored as 3. The scores to all the questions in each scale are added and then divided by the number of items in the scale, giving indices ranging from 1 (leftmost, most libertarian, most pro-welfare) to 5 (rightmost, most authoritarian, most anti-welfare). The scores on the three indices have been placed on the dataset.⁸

The scales have been tested for reliability (as measured by Cronbach's alpha). The Cronbach's alpha (unstandardised items) for the scales in 2014 are 0.82 for the left–right scale, 0.82 for the welfarism scale and 0.74 for the libertarian–authoritarian scale. This level of reliability can be considered “good” for the left–right and libertarian–authoritarian scales and “respectable” for the welfarism scale (DeVellis, 2003: 95–96).

Other analysis variables

These are taken directly from the questionnaire and to that extent are self-explanatory. The principal ones are:

- Sex
- Age
- Household income
- Economic position
- Religion
- Highest educational qualification obtained
- Marital status
- Benefits received

Sampling errors

No sample precisely reflects the characteristics of the population it represents, because of both sampling and non-sampling errors. If a sample were designed as a random sample (if every adult had an equal and independent chance of inclusion in the sample), then we could calculate the sampling error of any percentage, p , using the formula:

$$\text{s.e. } (p) = \sqrt{\frac{p(100 - p)}{n}}$$

where n is the number of respondents on which the percentage is based. Once the sampling error had been calculated, it would be a straightforward exercise to calculate a confidence interval for the true population percentage. For example, a 95% confidence interval would be given by the formula:

$$p \pm 1.96 \times \text{s.e. } (p)$$

Clearly, for a simple random sample (srs), the sampling error depends only on the values of p and n . However, simple random sampling is almost never used in practice, because of its inefficiency in terms of time and cost.

As noted above, the British Social Attitudes sample, like that drawn for most large-scale surveys, was clustered according to a stratified multi-stage design into 261 postcode sectors (or combinations of sectors). With a complex design like this, the sampling error of a percentage giving a particular response is not simply a function of the number of respondents in the sample and the size of the percentage; it also depends on how that percentage response is spread within and between sample points.

$$\text{DEFT} = \sqrt{\frac{\text{Variance of estimator with complex design, sample size } n}{\text{Variance of estimator with srs design, sample size } n}}$$

The complex design may be assessed relative to simple random sampling by calculating a range of design factors (DEFTs) associated with it, where:

and represents the multiplying factor to be applied to the simple random sampling error to produce its complex equivalent. A design factor of one means that the complex sample has achieved the same precision as a simple random sample of the same size. A design factor greater than one means the complex sample is less precise than its simple random sample equivalent. If the DEFT for a particular characteristic is known, a 95% confidence interval for a percentage may be calculated using the formula:

$$p \pm 1.96 \times \text{complex sampling error (p)}$$

$$= p \pm 1.96 \times \text{DEFT} \times \sqrt{\frac{p(100 - p)}{n}}$$

Table A.6 gives examples of the confidence intervals and DEFTs calculated for a range of different questions. Most background questions were asked of the whole sample, whereas many attitudinal questions were asked only of a third or two-thirds of the sample; some were asked on the interview questionnaire and some on the self-completion supplement.

Table A.6 Complex standard errors and confidence intervals of selected variables

	% (p)	Complex standard error of p	95% confidence interval	DEFT	Base
Classification variables					
Party identification (full sample)					
Conservative	31.5	.9	29.8-33.4	1.288	4326
Labour	29.8	.9	28.0-31.7	1.322	4326
UKIP	5.5	.4	4.7-6.3	1.163	4326
Liberal Democrat	4.5	.4	3.8-5.4	1.323	4326
Scottish National Party	3.3	.3	2.8-4.0	1.082	4326
UK Independence Party	5.5	.4	4.7-6.3	1.163	4326
None	14.6	.8	13.1-16.2	1.467	4326
Housing tenure (full sample)					
Owns	63.9	1.1	61.7-66.1	1.520	4328
Rents from local authority	9.9	.7	8.6-11.3	1.463	4328
Rents privately/HA	24.9	1.0	23.0-26.8	1.466	4328
Religion (full sample)					
No religion	48.5	1.0	46.5-50.5	1.343	4328
Church of England	17.0	.7	15.7-18.4	1.231	4328
Roman Catholic	8.7	.5	7.7-9.7	1.154	4328
Age of completing continuous full-time education (full sample)					
16 or under	44.9	1.1	42.7-47.0	1.422	4328
17 or 18	21.5	.8	20.0-23.1	1.283	4328
19 or over	28.9	1.0	27.1-30.9	1.392	4328

	% (p)	Complex standard error of p	95% confidence interval	DEFT	Base
Can I just check, would you describe the place where you live as... (full sample)					
...a big city	11.5	1.0	9.6-13.7	2.120	4328
the suburbs or outskirts of a big city	23.8	1.5	21.1-26.8	2.254	4328
a small city or town	44.7	1.7	41.3-48.1	2.282	4328
a country village	17.2	1.3	14.8-19.9	2.262	4328
or, a farm or home in the country?	2.1	.3	1.6-2.7	1.279	4328
Attitudinal variables (face-to-face interview)					
Benefits for the unemployed are ... (three-quarters of sample)					
... too low	23.5	.8	21.9-25.2	1.133	3266
... too high	58.8	1.0	56.7-60.8	1.195	3266
How serious a problem is traffic congestion in towns, cities (three-quarters of sample)					
A very serious problem	16.8	.7	15.4-18.3	1.133	3223
A serious problem	37.7	1.0	35.7-39.8	1.205	3223
Not a very serious problem	32.5	1.0	30.6-34.5	1.204	3223
Not a problem at all	12.8	.7	11.5-14.3	1.231	3223
How much interest do you generally have in what is going on in politics... (full sample)					
... a great deal	11.7	.6	10.6-12.8	1.159	4328
quite a lot	24.1	.8	22.6-25.7	1.219	4328
some	32.6	.9	30.9-34.3	1.198	4328
not very much	19.9	.7	18.5-21.4	1.225	4328
or, none at all?	11.6	.7	10.3-13.1	1.473	4328
Please say how much you agree or disagree with the following statements. 'There is no reason to worry about being a bit overweight'? (one half of sample)					
Agree strongly	4.4	.5	3.4-5.6	1.240	2188
Agree	41.9	1.3	39.3-44.6	1.266	2188
Neither agree nor disagree	10.6	.7	9.3-12.0	1.069	2188
Disagree	34.4	1.2	32.1-36.9	1.204	2188
Disagree strongly	8.6	.7	7.3-10.1	1.183	2188
Attitudinal variables (self-completion)					
Government should redistribute income from the better off to those who are less well off (full sample)					
Agree strongly	11.8	.6	10.6-13.2	1.201	3670
Agree	32.0	.8	30.3-33.7	1.089	3670
Neither agree nor disagree	28.2	.9	26.5-30.0	1.203	3670
Disagree	21.2	.8	19.7-22.9	1.213	3670
Disagree strongly	4.7	.4	4.0-5.5	1.124	3670
Some working couples with children find it hard to make ends meet on low wages. In these circumstances, do you think ... (three-quarters of sample)					
...the government should top-up their wages	55.3	1.1	53.2-57.4	1.132	2781
...or, is it up to the couple to look after themselves and their children as best they can?	29.4	1.0	27.4-31.4	1.178	2781
Can't choose	14.1	.7	12.8-15.6	1.089	2781

	% (p)	Complex standard error of p	95% confidence interval	DEFT	Base
People should be able to travel by plane as much as they like (one quarter of sample)					
Agree	62.7	1.7	59.3-66.1	1.075	889
Neither agree nor disagree	23.3	1.4	20.5-26.2	1.011	889
Disagree	6.7	1.1	4.9-9.1	1.264	889

The table shows that most of the questions asked of all sample members have a confidence interval of around plus or minus two to four of the survey percentage. This means that we can be 95% certain that the true population percentage is within two to four per cent (in either direction) of the percentage we report.

Variables with much larger variation are, as might be expected, those closely related to the geographic location of the respondent (for example, whether they live in a big city, a small town or a village). Here, the variation may be as large as six or seven per cent either way around the percentage found on the survey. Consequently, the design effects calculated for these variables in a clustered sample will be greater than the design effects calculated for variables less strongly associated with area. Also, sampling errors for percentages based only on respondents to just one of the versions of the questionnaire, or on subgroups within the sample, are larger than they would have been had the questions been asked of everyone.

Analysis techniques

Regression

Regression analysis aims to summarise the relationship between a 'dependent' variable and one or more 'independent' variables. It shows how well we can estimate a respondent's score on the dependent variable from knowledge of their scores on the independent variables. It is often undertaken to support a claim that the phenomena measured by the independent variables *cause* the phenomenon measured by the dependent variable. However, the causal ordering, if any, between the variables cannot be verified or falsified by the technique. Causality can only be inferred through special experimental designs or through assumptions made by the analyst.

All regression analysis assumes that the relationship between the dependent and each of the independent variables takes a particular form. In *linear regression*, it is assumed that the relationship can be adequately summarised by a straight line. This means that a one percentage point increase in the value of an independent variable is assumed to have the same impact on the value of the dependent variable on average, irrespective of the previous values of those variables.

Strictly speaking the technique assumes that both the dependent and the independent variables are measured on an interval-level scale, although it may sometimes still be applied even where this is not the case. For example, one can use an ordinal variable (e.g. a Likert scale) as a *dependent* variable if one is willing to assume that there is an underlying interval-level scale and the difference between the observed ordinal scale and the underlying interval scale is due to random measurement error. Often the answers to a number of Likert-type questions are averaged to give a dependent variable that is more like a continuous variable. Categorical or nominal data can be used as *independent* variables by converting them into dummy or binary variables; these are variables where the only valid scores are 0 and 1, with 1 signifying membership of a particular category and 0 otherwise.

The assumptions of linear regression cause particular difficulties where the *dependent* variable is binary. The assumption that the relationship between the dependent and the independent variables is a straight line means that it can produce estimated values for the dependent variable of less than 0 or greater than 1. In this case it may be more appropriate to assume that the relationship between the dependent and the independent variables takes the form of an S-curve, where the impact on the dependent variable of a one-point increase in an independent variable becomes progressively less the closer the value of the dependent variable approaches 0 or 1. *Logistic regression* is an alternative form of regression which fits such an S-curve rather than a straight line. The technique can also be adapted to analyse multinomial non-interval-level dependent variables, that is, variables which classify respondents into more than two categories.

The two statistical scores most commonly reported from the results of regression analyses are:

A measure of variance explained: This summarises how well all the independent variables combined can account for the variation in respondents' scores in the dependent variable. The higher the measure, the more accurately we are able in general to estimate the correct value of each respondent's score on the dependent variable from knowledge of their scores on the independent variables.

A parameter estimate: This shows how much the dependent variable will change on average, given a one-unit change in the independent variable (while holding all other independent variables in the model constant). The parameter estimate has a positive sign if an increase in the value of the independent variable results in an increase in the value of the dependent variable. It has a negative sign if an increase in the value of the independent variable results in a decrease in the value of the dependent variable. If the parameter estimates are standardised, it is possible to compare the relative impact of different independent variables; those variables with the largest standardised estimates can be said to have the biggest impact on the value of the dependent variable.

Regression also tests for the statistical significance of parameter estimates. A parameter estimate is said to be significant at the 5% level if the range of the values encompassed by its 95% confidence interval (see also section on sampling errors) are either all positive or all negative. This means that there is less than a 5% chance that the association we have found between the dependent variable and the independent variable is simply the result of sampling error and does not reflect a relationship that actually exists in the general population.

Factor analysis

Factor analysis is a statistical technique which aims to identify whether there are one or more apparent sources of commonality to the answers given by respondents to a set of questions. It ascertains the smallest number of *factors* (or dimensions) which can most economically summarise all of the variation found in the set of questions being analysed. Factors are established where respondents who gave a particular answer to one question in the set tended to give the same answer as each other to one or more of the other questions in the set. The technique is most useful when a relatively small number of factors are able to account for a relatively large proportion of the variance in all of the questions in the set.

The technique produces a *factor loading* for each question (or variable) on each factor. Where questions have a high loading on the same factor, then it will be the case that respondents who gave a particular answer to one of these questions tended to give a similar answer to each other at the other questions. The technique is most commonly used in attitudinal research to try to identify the underlying ideological dimensions which apparently structure attitudes towards the subject in question.

International Social Survey Programme

The International Social Survey Programme (ISSP) is run by a group of research organisations in different countries, each of which undertakes to field annually an agreed module of questions on a chosen topic area. Since 1985, an International Social Survey Programme module has been included in one of the British Social Attitudes self-completion questionnaires. Each module is chosen for repetition at intervals to allow comparisons both between countries (membership is currently standing at 48) and over time. In 2015, the chosen subject was Work orientations, and the module was carried on the B and D versions of the self-completion questionnaire (JbErnMny – AgeYngC).

Notes

1. Until 1991 all British Social Attitudes samples were drawn from the Electoral Register (ER). However, following concern that this sampling frame might be deficient in its coverage of certain population subgroups, a 'splicing' experiment was conducted in 1991. We are grateful to the Market Research Development Fund for contributing towards the costs of this experiment. Its purpose was to investigate whether a switch to PAF would disrupt the time-series – for instance, by lowering response rates or affecting the distribution of responses to particular questions. In the event, it was concluded that the change from ER to PAF was unlikely to affect time trends in any noticeable ways, and that no adjustment factors were necessary. Since significant differences in efficiency exist between PAF and ER, and because we considered it untenable to continue to use a frame that is known to be biased, we decided to adopt PAF as the sampling frame for future British Social Attitudes surveys. For details of the PAF/ER 'splicing' experiment, see Lynn and Taylor (1995).
2. This includes households not containing any adults aged 18 or over, vacant dwelling units, derelict dwelling units, non-resident addresses and other deadwood.
3. In 1993 it was decided to mount a split-sample experiment designed to test the applicability of Computer-Assisted Personal Interviewing (CAPI) to the British Social Attitudes survey series. CAPI has been used increasingly over the past decade as an alternative to traditional interviewing techniques. As the name implies, CAPI involves the use of a laptop computer during the interview, with the interviewer entering responses directly into the computer. One of the advantages of CAPI is that it significantly reduces both the amount of time spent on data processing and the number of coding and editing errors. There was, however, concern that a different interviewing technique might alter the distribution of responses and so affect the year-on-year consistency of British Social Attitudes data.

Following the experiment, it was decided to change over to CAPI completely in 1994 (the self-completion questionnaire still being administered in the conventional way). The results of the experiment are discussed in the British Social Attitudes 11th Report (Lynn and Purdon, 1994).

4. Interview times recorded as less than 20 minutes were excluded, as these timings were likely to be errors.
5. An experiment was conducted on the 1991 British Social Attitudes survey (Jowell *et al.*, 1992) which showed that sending advance letters to sampled addresses before fieldwork begins has very little impact on response rates. However, interviewers do find that an advance letter helps them to introduce the survey on the doorstep, and a majority of respondents have said that they preferred some advance notice. For these reasons, advance letters have been used on the British Social Attitudes surveys since 1991.

6. Because of methodological experiments on scale development, the exact items detailed in this section have not been asked on all versions of the questionnaire each year.
7. In 1994 only, this item was replaced by: Ordinary people get their fair share of the nation's wealth [*Wealth1*].
8. In constructing the scale, a decision had to be taken on how to treat missing values ("Don't know" and "Not answered"). Respondents who had more than two missing values on the left-right scale and more than three missing values on the libertarian-authoritarian and welfarism scales were excluded from that scale. For respondents with fewer missing values, "Don't know" was recoded to the midpoint of the scale and "Not answered" was recoded to the scale mean for that respondent on their valid items.

References

- DeVellis, R.F. (2003), *Scale Development: Theory and Applications*, 2nd edition, Applied Social Research Methods Series, 26, Thousand Oaks, California: Sage
- Jowell, R., Brook, L., Prior, G. and Taylor, B. (1992), *British Social Attitudes: the 9th Report*, Aldershot: Dartmouth
- Lynn, P. and Purdon, S. (1994), 'Time-series and lap-tops: the change to computer- assisted interviewing', in Jowell, R., Curtice, J., Brook, L. and Ahrendt, D. (eds.), *British Social Attitudes: the 11th Report*, Aldershot: Dartmouth
- Lynn, P. and Taylor, B. (1995), 'On the bias and variance of samples of individuals: a comparison of the Electoral Registers and Postcode Address File as sampling frames', *The Statistician*, 44: 173-194
- Spector, P. (1992), *Summated Rating Scale Construction: An Introduction, Quantitative Applications in the Social Sciences*, 82, Newbury Park, California: Sage
- Stafford, R. and Thomson, K. (2006), *British Social Attitudes and Young People's Social Attitudes surveys 2003: Technical Report*, London: NatCen Social Research