Economic and Social Data Service

Weighting the Social Surveys

ESDS Government

Author: Alasdair Crockett
UK Data Archive and Institute for Social and Economic Research

Updated by: Reza Afkhami
Anthony Rafferty
Vanessa Higgins
Alan Marshall

Version: 1.9
Date: May 2011

University of Essex

MANCHESTER 1824 The University of Manchester

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL

This document is based upon the proceedings of the Weighting the Social Surveys event held by the Cathie Marsh Centre for Census and Survey Research (CCSR) at the Royal Statistical Society on Friday 12th March 2004.

The event was chaired by Jo Wathan (CCSR, University of Manchester) and the speakers were:

Ian Plewis, Institute of Education, University of London.
Jeremy Barton, Office for National Statistics.
Susan Purdon, National Centre for Social Research.
Peter Lynn, Institute for Social and Economic Research, University of Essex.
Nick Buck, Institute for Social and Economic Research, University of Essex.

The author would like to thank all the speakers. Their comments inform all sections of this report. To see the individual presentations see: http://www.ccsr.ac.uk/esds/events/2004-03-12/slides.shtml

Contents

1. Introduction

This report aims to provide a simple guide to weighting, for users of the major government social surveys supported by the Economic and Social Data Service Government function, known as ESDS Government.

The issue of weighting, and allowing for survey design more generally, remain poorly understood by parts of the UK social science research community. The important point to realise is that unless social survey data arise from an 'equal probability of selection method' (referred to as EPSEM) and almost everyone selected agrees to be interviewed, then the sample will provide a biased representation of the total population unless adequate correction is made in subsequent analysis. This correction is usually done by weighting, to correct for the non-equal probability of selection of respondents, and differential response rates within the group of selected individuals/households. Both effects combine to mean that some types or classes of individuals are more likely to be in the achieved sample than others, and this means the achieved sample will only be representative of the population the survey aimed to reflect once the data are weighted.

In practice, few social surveys use an EPSEM design. Typically, some sample members are given a higher selection probability than others – either through a desire to over-represent important small groups in the population, or because the available sampling frame gives the researcher no choice. An example of the first type of design is when certain groups (e.g. pensioners or ethnic minorities) or regions (e.g. Northern Ireland) are deliberately over-sampled to provide more precise estimates for that group or region. An example of the second type of design is when addresses are selected using EPSEM from the Postcode Address File and then one person is selected for interview at each address (it is not possible to select persons using EPSEM as the PAF does not indicate the number of persons resident at each address). In addition, response rates are seldom close to 100% (normally less than 80%), and usually vary by the type of respondent, such that non-response is unlikely to be at random with respect to the social, demographic or economic characteristics in which the analyst is likely to be interested. The figures below shows recent trends in response rates for the Office for National Statistics major social surveys.

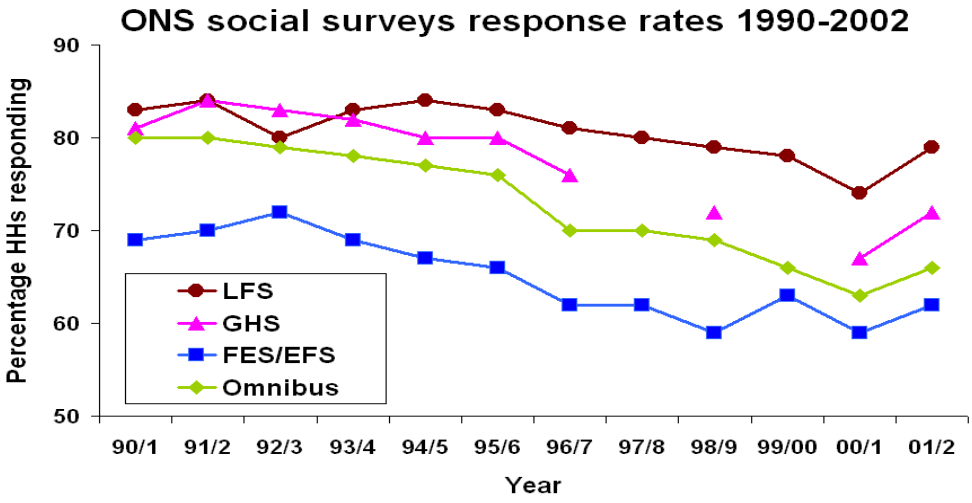

LFS = Labour Force Survey
GLF = General Lifestyle Survey (formerly the General Household Survey - GHS)
FES/LCF = Family Expenditure Survey/ Living Costs and Food Survey (formerly the Expenditure and Food Survey – EFS)
Opinions = ONS Opinions Survey (formerly the ONS Omnibus Survey)
Source: Weighting on National Statistics Household Surveys, Jeremy Barton, Office for National Statistics, presentation given at ESDS Weighting Meeting 12th March 2004

Another common reason for weighting is to use the data for a different unit of analysis than that for which the sample was primarily designed. For example we may wish make the data representative of households as opposed to individuals, to reflect our research interests.

This report covers the issues of weighting in the government surveys supported by ESDS Government, which are generally repeated cross-sectional surveys (i.e. a different sample of people are interviewed each time the survey is conducted). Weighting is also an important, and more complex, issue in the major longitudinal surveys, in which the same group or panel of respondents are repeatedly interviewed for several years or decades (but attrition means that fewer take part each time). This report covers only surveys within the remit of the ESDS Government service. There are several important longitudinal surveys, which are covered by the ESDS Longitudinal service. Peter Lynn and Nick Buck of the Institute for Social and Economic Research were present at the meeting on which this report is based, and the slides which were produced for their excellent talks are available on the ESDS website.

1.1. What is weighting and why is it important?

Almost all the major British social surveys require weighting. If data requiring weighting are not weighted the resulting estimates will be biased if they are interpreted as estimates for the wider population (as opposed to estimates relating to the achieved sample). In almost all social science analysis, one is interested in the characteristics of the wider population (typically, this being the population of Britain, the United Kingdom or one or more of its constituent countries) rather than the achieved sample. For example, the British Social Attitudes Survey (BSAS) is designed to provide estimates of attitudinal data for the adult British population, but due to both differential selection probabilities (interviewing one dwelling per address, one household per dwelling[1], and one adult per household), one cannot interpret the achieved sample of the BSA as providing unbiased estimates of the social attitudes of the adult British population. To generate estimates that are unbiased estimates of the British adult population, one has to weight the BSAS data.

It is important to note that the issue of bias does not just relate to complex multivariate methods such as regression, it also relates to simple descriptive statistics such as mean income, or the proportion that say they will vote at the next general election. Weighted analysis is for all social scientists, not just specialist statisticians. Indeed, for simple descriptive statistics weighting is invariably the correct thing to do, whereas for multivariate modelling there may be alternative methods that generate more precise estimates than can be achieved via weighting (see section 1.3).

A second important point to note is that weighting also involves adjustment to the precision of one's estimates. The standard measure of precision is the standard error, which tells you how close to the real value (i.e. the actual value among the population) the point or parameter estimate (e.g. means, proportions, regression coefficients) is likely to be. When data are weighted, the precision of point and parameter estimates will tend to decline.[2] The precision of weighted parameter estimates will typically be lower than the corresponding precision of the unweighted estimates, though this is not always the case.

The third and final important point is that the effects of weighting are specific to each and every variable in the dataset. Some characteristics might be more common among the under-represented class(es) of respondent, others might be less common. The former characteristics will appear more common when the data are weighted, while the latter will appear less common. If the characteristics vary at random with respect to the weighting variable, the weighted and unweighted parameter estimates will be the same. Similarly, when examining relationships via regression or other techniques, the relationship between two variables might be stronger or weaker among under-

[1] Usually, an address is a small-user address point from the Postcode Address File. A dwelling unit is a self-contained unit of accommodation. A household is usually defined as a group of individuals who either share living acocomodation or a meal per day. Definitions can vary by survey.

[2] The exception is post-stratification weights considered at section 2.3, which may increase precision. In practice, the net effect of weighting in the major social surveys is to reduce precision.

represented groups. If the former, the resulting parameter estimates (e.g. regression coefficients) will be larger once the data are weighted, if the latter they will be smaller..

1.2 What software to use

The functionality offered by statistical software is constantly increasing. However prior to version 12, SPSS was not capable of correct weighted data analysis because it did not estimate the precision of parameter estimates correctly. This means the standard errors generated by SPSS are too small, which can lead to spurious statistical significance (as illustrated in section 1.3). The Complex sample module of version 12 of SPSS does conduct weighted analysis correctly (and also allows for design effects due to clustering and stratification), but this only covers descriptive statistics. The complex samples modules in versions beyond V.12 of SPSS contain an increasing number of multivariate commands that allow for correct survey weighting. See http://www.spss.com/complex_samples/ for further details.

The other major multi-purpose statistical packages, Stata, SAS and R, all conduct weighted analysis correctly. Users of ESDS government data are advised to use Stata, SAS or R for their analyses for this reason. All ESDS government data are made available for immediate download in Stata format (as well as SPSS and tab-delimited text), however, not all procedures are available in this module.

Though not covered by this report in any detail, there are other important aspects of survey design, in addition to weighting, that affect the standard errors of survey estimates. For most of the large social surveys, these should be incorporated into one's analysis in order for standard error estimates to be correct (see 1.4 below). If you wish to incorporate these other design features, Stata, R and the specialist packages SUDAAN and WESVAR offer the greatest functionality. Stata is the most easy to use of these options and offers easy to use functionality for conducting weighted analysis and including other design features via use of the 'svy' commands. The design need only be specified once, and all subsequent commands prefixed by 'svy' will calculate standard errors in an appropriate way. The additional menu support for version 8 of Stata makes setting weights and other design features even easier. R is open source (i.e. free) package and has been ported to run in Windows as well as LINUX/UNIX, and is the best choice if your institution has no license for Stata and financial restrictions prevent you from purchasing your own license.

1.3 Should one always weight one's analyses?

As a general principle, one should always carry out weighted analysis. If you weight by the appropriate weight variable, the point and parameter estimates you generate (e.g. means, proportions, and regression coefficients) will be unbiased population estimates. Information about weighting variables should be available in the appropriate documentation, which can be obtained from the survey pages on the ESDS webpages. The documentation should always be consulted before attempting any analysis.

Weighting to adjust for unequal sampling probabilities is therefore never a 'wrong' thing to do, but it can be sub-optimal for certain multivariate analyses in that it may reduce precision more than alternative ways of accounting for the same effects. In some models it may be possible to incorporate the information encapsulated in the weight variable (and other design features) as explicit variables on the right-hand side of one's equation. In so doing you can achieve estimates that are unbiased population estimates and may be of higher precision than would result from a weighted analysis.

If you do not weight you must incorporate all the effects encapsulated in the weighting variable, and this usually requires substantial statistical expertise. If in doubt, always weight your analyses. Incorporating design features in other ways requires specialist knowledge. If you do not know how and lack a source of expertise to ask at your institution, then you should weight your analysis; at worst this will be sub-optimal.

1.4 What other design effects are there?

The effects of unequal inclusion probabilities - controlled by applying weights - are usually the most important to incorporate in one's analysis as the weighting ensures unbiased population estimates (as well as reducing the precision of estimates). Other features of survey design affect only the precision of estimates; some act to reduce precision, some to increase it. So, for statistically rigorous hypothesis testing, these design features are important. The precise nature of design effects is specific to the design of each survey. Two additional effects commonly affect British and UK social surveys; these are known as clustering and stratification effects.

Clustering Effects
Many surveys have primary sampling units (PSUs), for example post code sectors if the sampling frame is the post code address file. This means that rather than selecting the same proportion of respondents from every PSU in the population - which is very expensive and time consuming (because of the travel involved) – sample designers select a sample of PSUs and then select sample elements (e.g. households) only from the sampled PSUs. The result is that respondents are clustered within certain geographical areas. To the extent that the characteristic of interest to the researcher (e.g. income) is homogeneous within a PSU but varies between PSUs, the effect of this clustering will be to reduce the precision of population estimates.

Stratification effects
By contrast, some sample designs include stratification. Strata are groupings defined by criteria that are likely to be important to subsequent analysis, such as geographical location, social, demographic and ethnic composition, and units are sampled within these. Stratification serves to ensure that the sample is distributed over the strata in the same way as the wider population. The sample therefore better reflects the population than it would have been likely to if it were selected entirely at random. For this reason, stratification effects act to increase the precision of population estimates. The effect is stronger the stronger the relationship between the characteristic of interest to the researcher and the characteristics used to define the strata.
It is common for sample designs to incorporate both clustering and stratification elements. Each has effects on the accuracy of your results. If you ignore clustering effects (where these exist), your estimates will appear too precise - i.e. the standard errors you obtain will be under-estimates, and apparent statistical significance may be spurious as a result. Stratification effects (where these exist) act in the opposite direction though are generally weaker than clustering effects, such that clustering and stratification in combination will generally cause a modest reduction in the precision of your estimates. Information about the sample design used in your survey of interest should be available in the documentation. However, it should be noted that it will only be possible for you to obtain unbiased estimates of standard errors, taking into account the clustering and stratification, if the data set includes variables indicating PSUs and strata. Not all data sets include this information.

1.5 Illustrative case study using the British Social Attitudes Survey data

Let us imagine that we are interested in changes in the rates of religious affiliation between 1994 and 2001. Let us further imagine, for the sake of simplicity, that the British Social Attitudes (BSAS) Survey was only conducted in 1994 and 2001. One then has a simple question to test from the BSAS data: was there a statistically significant difference in the proportion of British adults reporting a religious affiliation in 2001 compared with 1994..

If one examines the unweighted data the results look like this:

	Percent of BSAS respondents with a religious affiliation	Standard error
1994	62.0	0.83
2001	58.5	0.86

In terms of a formal test, the probability of this difference between 1994 and 2001 arising by chance (i.e. that both percentages arise from the same binomial distribution) is 0.004 (i.e. a 1 in 250 chance), so affiliation rates were significantly different between the two years. However, this result

relates to the achieved BSAS samples, not to the adult British population. To make the BSAS estimates unbiased estimates of the adult British population; we need to apply the weight variable 'wtfactor' (a sample design weight). If we do this we get the following:

Data weighted by 'wtfactor'

	Percent of adult British population with religious affiliation	Standard error	Standard error according to SPSS
1994	61.4	0.92	0.83
2001	58.5	0.94	0.86

Note how the percentages in the second column have changed very little (so the bias of the unweighted estimate was minimal), but that the weighted standard errors in the third column are substantially higher, indicating that precision has been reduced by weighting. Note also the standard errors generated by SPSS (not using the post version 12 Complex sample module) are too small, they have not altered from the unweighted analysis shown in the previous table[3].

In terms of a formal test, the probability of this difference between 1994 and 2001 arising by chance is 0.028 (i.e. the odds are 1 in 36 that the difference is due to chance).

Notice how the weighted analysis gives a much higher likelihood of a chance result (1 in 36) compared with the unweighted analysis (1 in 250). This is largely because weighting has reduced the precision of the results.

Including the other BSAS design effects

If one adds in the clustering and stratification effects in the BSAS,[4] the precision of the population estimates is reduced further, as the results below show.

Full design effect results (using Stata)

	Percent of adult British population with religious affiliation	Standard error (incorporating full design effect)
1994	61.4	1.01
2001	58.5	1.00

Note the modest additional increase in the standard errors in the third column. The probability of this difference between 1994 and 2001 arising by chance (i.e. that both percentages arise from the same binomial distribution) is now equal to 0.047, in other words there is a 1 in 21 chance that difference is due to chance.

In this case study, most of the reduction in precision arose from weighting and the additional reduction from specifying the full design effect was relatively small. This is because the BSAS design

[3] In BSAS the mean weight is 1. The effective sample size in SPSS will be the same as the real achieved sample size. However, many ESDS Government surveys use weights that have a much larger mean, as the weights are used to produce population estimates. In this case, the effective sample size will be the achieved sample size times the mean weight. This will appear to reduce the standard error hugely, but this is an artefact of how SPSS applies weights in standard comands and is incorrect.

[4] These are a clustering effect - due to the non-equal probability of selection by PSU (post code sector), and weak stratification effects arising from the criteria used to select PSUs. In terms of using Stata, one specifies what variable corresponds to the PSU (post code sector) to account for the clustering effect, and to account for the modest stratification effect, one needs to create a new variable which is based on consecutive pairs of PSUs in terms of the order they were selected. This can be done as the data creators, the National Centre for Social Research (NATCEN), leave the variable called 'spoint' in the data supplied to ESDS government (and this gives PSU selection order by region).

involves considerable variation in selection probabilities, while attitudes tend not to cluster greatly within postcode sectors. But this will not always be the case. For example, on a survey such as the Health Survey for England, where all persons are sampled within each sampled household (so no variation in selection probabilities), there will be no reduction in precision due to weighting (of the adult sample), whereas health measures do tend to cluster within postal sectors, resulting in a reduction in precision due to clustering.

The BSAS case study illustrates how a naïve unweighted analysis of the BSAS data would lead one to reject without hesitation the null hypothesis that a different proportion of British adults reported a religious affiliation in 2001 than in 1994. When, however, the analysis was weighted using appropriate software (Stata), and when the full design effect was specified, the difference between 1994 and 2001 was at the margins of whether we would accept or reject the null hypothesis (in both instances it would be rejected at the 0.05 significance level but accepted at the 0.01 significance level).

2. Types of Weight

To understand more fully what weighted analysis entails, one needs to distinguish the three primary types of weight that can exist in a given social survey dataset. These are sample design weights, non-response weights, and post-stratification weights. These three types of weight are explained in the following sub-sections.

2.1 Sample design or probability weights

Sample design or probability weights correct for cases having unequal probabilities of selection that result from sample design. It is important to note that non-equal selection probabilities can also occur due to differentials in non-response, which is corrected by non-response weights described at 2.2. below. Minor discrepancies may also require adjustment if the sampling frame (e.g. the postcode address file) does not entirely reflect the population, and these would constitute a type of post-stratification weight outlined at 2.3 below.

To illustrate how a sample design weight is calculated, consider a survey design that interviews one dwelling per address, one household per dwelling and one adult per household. Provided information concerning dwellings per address, households per dwelling and adults per household is enumerated by the interviewer, one can subsequently calculate sample design weights that correct for the lower selection probabilities of adults in multi-adult (and household/dwelling) households. The general formula for a sample design weight is arithmetically very simple, it is 1 divided by the probability of selection due to the survey design. However, these are usually scaled, so we define the weight as proportional to this number. For example, if there are 3 adults in a given household the resulting sample design weight for the single interviewed adult will be proportional to 1/(1/3), i.e. proportional to 3. In a one adult household, the weight will be simple proportional to 1/1, i.e. proportional to one. In other words the influence of the former respondent is being increased threefold relative to the influence of the latter respondent to exactly compensate for the fact the former respondent was three times less likely to be included in the sample.

The weights are often scaled to have a mean of 1, which maintains an effective sample size when the data are weighted.

2.2 Non-response weights

Non-response weights compensate for differential response rates. Response rate in this sense refers to unit non-response, whereby someone refuses to take part in the survey at all, as opposed to item non-response, which relates to refusing to answer specific questions, which is addressed by missing data methods rather than weighting.

Non-response weights are typically obtained by defining weighting classes, which are based on information available for both responding and non-responding households. Such information typically relates to geographical location, primary sampling unit (PSU) characteristics (which are derived from other data sources, often the Census) and often household and dwelling type (which need to be recorded by the interviewer).

Respondents in each weighting class are weighted to compensate for the proportion of non-respondents in that class. More formally, the non-response rate weight is proportional to 1 divided by the response rate for the weighting class, i.e. directly analogous to sample design weights.

Sample design weights usually control exactly for differences in selection probability due to sample design, but non-response weights are seldom entirely accurate. The utility of the non-response weights is governed by the amount of information available to define the weighting classes. By definition, information about non-respondents is limited. The assumption of non-response weights is that the characteristics of respondents and non-respondents within each weighting class are the same; only if they are will the non-response weight be entirely accurate. If you as an analyst are examining characteristics that do vary between respondents and non-respondents within weighting

classes, then the weighted estimates you derive will be biased population estimates. This problem is most likely to occur when examining measures of social engagement such as voting/not voting, which are likely to be highly correlated (even within a weighting class) with whether a respondent agrees to take part in a survey or not.

Further information on this topic is available in Ian Plewis' slides on the ESDS website (op cit).

2.3 Post-stratification weights

Post-stratification weights (also known as population or calibration weights) are constructed after the other types of weights have been constructed and applied to the data. They are applied to make the data even more representative of the population. As for probability weights, information on the population is usually derived from the decennial Census of Population.

These weights allow for more accurate population totals of estimates, they reduce non-response bias further (over and above non-response weights), and improve precision

Whereas sample design (probability) and non-response weights result from a very simple computation (1/selection probability), post-stratification weights are mathematically complex, requiring iterative algorithms that maximise the fit of the data to the population. This procedure is called 'raking', and requires specialist software. The Office for National Statistics previously used a SAS based macro called CALMAR for the estimation of post-stratification weights but have recently switched to the Generalized Estimation System (GES) programme for many surveys.

For example, prior to 2007, calculation of the Labour Force Survey individual level post-stratification weights (see: Barton 2004) used CALMAR and involved 'raking' to 3 controls (derived from the Census and population projections):

5-year age group by sex within region
Local Authority
Single years 16-24 by sex
Population projections

The raking procedure iterates until the data best match all three controls, and computes the post-stratification weight accordingly.

The more robust and efficient Generalized Estimation System (GES) programme has been used since 2007 for the calculation of post-stratification weights in the Labour Force Survey. The methodology employed by the GES tool differs from CALMAR in that it calibrates the data in a single process rather than numerous iterations over three stages. The two methodologies have been shown to produce equivalent estimates for surveys such as the LFS that have large sample sizes. For more details see Palmer and Hughes (2008).

2.4 Concluding remarks

Few datasets supplied by the ESDS will contain distinct weight variables that correspond to these three types of weight. By multiplying the relevant weights together, one can create a single weighting variable that incorporates all three effects (or however many exist in a given study). Since this makes life easy for the secondary data user, this is typically done by the major data creators: the Office for National Statistics (ONS) and the National Centre for Social Research (NATCEN) and thereby supplied in the datasets provided by ESDS Government. The weights in these surveys will generally incorporate all the relevant weighting effects. This is not to say there is only ever one weighting variable in the dataset, there may be separate weights to make the data representative of individuals versus households, or to make the data representative of different geographical populations, e.g. the United Kingdom versus Great Britain.

While design weights are not usually changed (i.e. they remain valid in perpetuity), non-response weights and post-stratification rates may be subsequently altered to reflect new and better information becoming available to the data creator. The Office for National Statistics (ONS) is

currently shifting the basis upon which weighting classes for non-response weights and post-stratification weights are calculated for surveys conducted in the late 1990s from the 1991 census (and forward projection there from) to the 2001 census (and back projection there from). As a result ONS surveys, particularly the Labour Force Surveys, are periodically resupplied to ESDS Government with recomputed weight variables. When this occurs, anyone who has previously ordered the affected datasets will automatically be notified of the resulting new edition of the data and be resupplied with the revised data and documentation.

3. Weighting Variables for the ESDS Government Surveys

The information given below relates to the latest available data for individual surveys. You should refer to the survey documentation on the ESDS website[5] for the specific year(s) you are interested in, as the weighting may change slightly from year to year.

3.1 Annual Population Survey

This information is extracted from the APS user guide at:
http://www.esds.ac.uk/doc/6310/mrdoc/pdf/6310userguide.pdf

Since December 2005 the APS has included data from the Labour Force Survey (LFS) data and English Local Labour Force Survey (LLFS)..Before December 2005, the APS also included Annual Population Survey boost (APS(B)) data which covered a subset of the topics covered on the LFS and LLFS. All variables on the LLFS appeared on the APS dataset including those which are not on the APS (B).

The main purpose of the APS weights is to gross to the population. However, this is achieved through calibration to age/sex/region totals which means that the APS weights *indirectly* deal with some of the main areas of concern for non-response.

For datasets up to and including the January – December 2005 dataset the APS requires two weighting variables due to the different data sources (the APS and the LFS) which make up the final dataset. One weight is required when looking at core variables, and one weight when looking at either only non-core variables or a combination (e.g. a crosstab) of core and non-core variables. A summary of which weight to use is as follows:

PWAPS04a
This is used when looking at only core variables. These are those marked as X and Y in diagram 1 in the document in the first link above.

PWLFS04a
This weight is used when looking at either only variables which are non-core or looking combinations of core and non-core variables. These are those marked as Z in diagram 1 in the document in the first link above.

The last letter of the weighting variable changes with each quarter, as it represents the next quarter. Every quarter there will be a new weight as the weight is calculated on the sample size and characteristics. So as each new dataset is available and is different to the previous one there is a new weight calculated for each quarter and this new weight is represented by the change in last letter on the weight variable. A spreadsheet of core/non-core variables is included with the documentation from ESDS for each year of the APS (an example of this is included in the spreadsheet entitled 'weighting summary' on the ESDS website).

All APS datasets after December 2005 contain only data from the LFS and LLFS (no APS(B) data) and so there is no need for two weights. The weight pwt07 is used for these surveys.

In 2007, ONS undertook a reweighting project, whereby APS and LFS data were reweighted using population estimates for 2007-2008 using the new Generalized Estimation System reweighting programme. As a result, reweighted editions of APS datasets were redeposited at UKDA during 2008 and have been added to the collection. For more details on the reweighting exercise see:

[5] http://www.esds.ac.uk

The next APS reweighting exercise to be undertaken will use 2009 population figures, and it is planned that revised datasets will be deposited accordingly during 2010. At this time data will be redeposited with a new set weights which will most likely be called pwt09.

3.2 Labour Force Survey

Since 1984 the LFS has been weighted (grossed) to produce population estimates and to compensate for non-response among sub-groups. Additionally, the earnings data is also grossed. As part of the 2009 reweighting exercise new weights were released for all LFS datasets from mid-2006 onwards. For datasets prior to this date the weights associated with the 2007 reweighting exercise still apply.

For more information on the 2009 reweighting exercise see:
http://www.esds.ac.uk/doc/6199/mrdoc/pdf/lfs_reweighting_2009.pdf

The 2009 Quarterly LFS datasets have two weights (Pwt09 and Piwt09), (1) Pwt09 is the weight for individual data - this compensates for non-response and grosses to population estimates. (2) Piwt09 is the weight for income data - this weights so that that the weight of a sub-group corresponds to that sub-group's size in the population and also weights to give estimates of the number of people in certain groups. This is restricted to employees' earnings: other income data are not (yet) weighted. NB: In 2010 Pwt09 and Piwt09 replaced the weights pwt07 and piwt07 because of the re-weighting exercise to bring LFS data in line with the population estimates from the 2009 mid-year estimates.

The QLFS household datasets contain individual level data for households, but have been designed for household analyses. They have one weight to gross to population estimates. The weight is the same for all household members. The weighting variable for quarterly household dataset (April to June 2009) is called phhwt09. See section 4 of the Household and Family Data User Guide[6] for more information (note this guide has not been updated to acknowledge that the 2009 reweighting exercise has taken place).

The QLFS longitudinal datasets (2-quarter and 5-quarter) contain one weight to compensate for non-response and to produce population estimates. The 2009 weighting variable (Two-Quarter Longitudinal Dataset, October 2008 - March 2009 and Five-Quarter Longitudinal Dataset, January 2008 - March 2009) is called LGWT. See the Longitudinal Datasets User Guide[7] for more information.

Since 2007, the LFS weights have been produced using a Generalised Regression (GREG) framework and the Statistics Canada Generalised Estimation System (GES). For more information see:
http://www.esds.ac.uk/doc/5715%5Cmrdoc%5Cpdf%5Clfs_reweighting_2009.pdf

Users should consult the survey documentation for information about the sample design, which involves a five-quarter rolling panel.

Latest userguide:
http://www.data-archive.ac.uk/doc/6199%5Cmrdoc%5Cpdf%5Cbackground.pdf

3.3 General Lifestyle Survey (GLF) (formerly General Household Survey)

In 2005, the GLF methodology changed to longitudinal data collection so that in 2006 dataset a proportion (68%) of the sample are people who were also interviewed the year before. The name of the survey was changed from the General Household Survey to the General Lifestyle Survey at this time.

The dual weighting scheme in the 2006 GLF is very similar to that employed since 2000 involving one weighting variable for two purposes (1) to compensate for non-response and, since the introduction of a longitudinal design, attrition in the sample (2) to gross up to match known population distributions in terms of region, age-group and sex. The 2006 weighting variable is called Weight06.

[6] http://www.esds.ac.uk/doc/6292/mrdoc/pdf/lfs_vol8_household2008.pdf
[7] http://www.data-archive.ac.uk/doc/6202%5Cmrdoc%5Cpdf%5Clongitudinal.pdf

See Appendix D of the GHS 2006 documentation produced by the Office for National Statistics for more information on the production of GLF weights. For more details on the Generalized Estimation System (GES) programme used to gross up survey estimates to match known population totals see Palmer and Hughes (2008).

Weight variable: Weight06

The data set is unweighted. Weight06 is the variable you should use to weight the data (see http://www.statistics.gov.uk/downloads/theme_compendia/GHS06/AppendixD2006.pdf). This weight applies to both household and individual level data.

The GLF sample is based on private households, which means that the population totals used in the weighting need to relate to people in private households. These totals are taken from population projections for local authorities based on mid-year estimates and adjusted to exclude residents of certain institutions.. There have been revisions to some local authority population estimates which impact upon weights used in the GLF and it likely there will be further alterations in the future.

Latest userguide:
http://www.esds.ac.uk/findingData/snDescription.asp?sn=5804

3.4 British Crime Survey

The BCS has been weighted since 1982. The BCS 2008-9 includes four weights. **Indivwgt** should be used for individual based analysis (attitudinal questions and estimates of personal crime rates). **Hhdwgt** should used for household based analysis (estimates of household crime rates). For incident-based analysis, the weight **weighti** should be used. For analysis confined to 16-24 year olds a weight based on 16-24 year olds from the main sample and those in the young adults boost sample should be used (**ypcwgt**).

There are three main reasons for weighting the BCS (1) to compensate for unequal selection probabilities (2) to compensate for differential response rates (3) to ensure that quarters are equally weighted for analyses that combine data from more than one quarter. In the 2008-9 BCS, the components of the weights are:

- w_1 : weight to compensate for unequal address selection probabilities in each Police Force Area;
- w_2 : inner city versus non inner-city non-response weight;
- w_3 : dwelling unit weight;
- w_4 : individual selection weight;
- numinc : series of incidents weight

More information on each component see the BCS 2008-9 Technical report[8].

The table below shows the weighting components that are included in the four BCS weights.

Components included in each of the BCS weights

Weight	Component
Individual weight (Indivwgt)	$W_1*w_2*w_3*w_4$
Household weight (Hhdwgt)	$W_1*w_2*w_3$
Incident weight (weighti)	Numinc and the hh or individual weight components depending on offence code
Youth weight (ypcwgt)	$W_1*w_2*w_3*r^i$

i: r is the number of the number of adults aged between 16 and 24 years in the household

Since 2001, the Home Office have applied additional calibration weights once they receive the data so that the (weighted) data reflects the population profile by age and

[8] http://www.esds.ac.uk/doc/6367/mrdoc/pdf/6367_bcs_2008-09_technical_report_vol1.pdf

sex within Government Office Regions (see section 7.6 of the BCS 2008-9 technical report).

In sweeps of the BCS which also included an ethnic boost, the boost is only included when examining results by ethnic group. The boost is excluded from all other analysis.

3.5 Scottish Crime Survey

In April 2008 the Scottish Crime and Justice Survey (SCJS) replaced the Scottish Crime and Victimisation Survey (SCVS) which had replaced the Scottish Crime Survey (SCS) in 2004.

The SCJS is weighted for three reasons:

1. To correct the sample for unequal probabilities of selection that arose from various aspects of the sample design
2. To correct the sample for differing response rates by sub-groups within the sample
3. To gross up the sample data to allow the results to be expressed as population values

The survey has a number of different weights which should be applied in different circumstances. For example, the 2008-9 SCJS has the following weights:

Weight	Files*	Description
WGTGHHD	RF and VFF	Gross household weight (grossed to population)
WGTGINDIV	RF and VFF	Gross individual weight (grossed to population)
WGTGINC_SCJS	VFF	Gross incident weight SCJS crimes (The values are the products of the appropriate household or individual weight and the number of incidents (the incident count), capped at five)
WGTGHHD_SC	SCF	Self-completion household weight (grossed to population)
WGTGINDIV_SC	SCF	Self-completion household weight (grossed to population)

RF = Respondent form. VFF = Victim form file. SCF = Self-completion form file

Separate weights are calculated for the self completion form (SCF) because of the higher levels of non-response compared to the respondent form. It is thought that the sensitive nature of questions in the SCF is responsible for this higher non-response. More details on the calculation of the weights in the SCJS can be found in section 8 of the SCJS 2008-2009 Technical report:
http://www.data-archive.ac.uk/doc/6362%5Cmrdoc%5Cpdf%5C6362_scjs0809_techreportwithannexes.pdf

3.6 British Social Attitudes Survey

The BSAS has been weighted since 1983. In 2005 the BSAS moved to a more sophisticated set of weights that included two new components to correct for non-response and to calibrate the sample to regional sex and age population profiles. As was the case for surveys prior to 2005 the weights also take into account differing selection probabilities.

The 2008 survey has a weight called wtfactor which must be used in all analysis – the data is not preweighted.

When reporting time-series analysis, there is a small possibility that the change of weighting scheme (in 2005) could disrupt the time-series. As a precaution, NATCEN recommend that when reporting time-series analysis figures from 2005 onwards the calculations should be rerun using the old weighting structure (oldwt) to check that this does not present a radically different picture. The figures produced using the new weights (wtfactor) should still be the ones used in reporting, but any substantial differences should be mentioned in a note.

Latest userguide:
http://www.data-archive.ac.uk/doc/6390%5Cmrdoc%5Cpdf%5C6390userguide.pdf

3.7 Scottish Social Attitudes Survey

The SSAS is weighted to (1) account for differing selection probabilities because only one person in the household is interviewed, (2) to account for the addresses in remote and rural parts of Scotland having a greater chance of selection due to the rural boost and (3) to account for non-response . One weight is used (WtFactor in 2007).

The weights in the 2005 SSAS were the first to include a component to correct for non-response and are considered superior to the weights used prior to 2005 for this reason. The new weights (WTFACTOR) should therefore be used in all reported analysis. However, when reporting time-series analysis, there is a small possibility that the change of weighting scheme could disrupt the time-series. The 2007 dataset also includes a variable based on the old weighting structure (OLDWT). It is recommended by the Scottish Centre for Social Research that **when reporting time-series analysis** – and particularly when presenting 'head line' frequencies without more detailed analysis – **the 2007 figures should be rerun using the old weighting structure** (OLDWT) to make sure that this does not present a radically different picture. The figures produced using the new weights (WTFACTOR) should still be the main ones used in reporting.

Latest userguide: **http://www.esds.ac.uk/doc/6262/mrdoc/pdf/6262userguide.pdf**

3.8 Northern Ireland Life and Times Survey

All analyses of the adult data should be weighted in order to allow for disproportionate household size. In 2008 the weighting variable is called WTFACTOR. The only exceptions are the few household variables (for example, tenure and household income), which do not need to be weighted.

Latest userguide:
http://www.data-archive.ac.uk/doc/6296%5Cmrdoc%5Cpdf%5C6296userguide.pdf

3.9 Young Peoples Social Attitudes

As with the British Social Attitudes Survey (BSAS), the YPSA data were weighted to take account of the relative selection probabilities of the BSAS adult respondent at the two main stages of selection: address and household. In this respect the young people's data were weighted in the same way as the adult data. The weight on the 2003 dataset is called YPWT.

Latest userguide for 2003 survey
http://www.esds.ac.uk/doc/5250/mrdoc/pdf/5250userguide.pdf

3.10 Living Costs and Food Survey (previously known as the Expenditure and Food Survey)

In 2008 the Living Costs and Food Survey (LCF) replaced the Expenditure and Food Survey (EFS). More information on the LCF can be found on the Office for National Statistics web site.

The LCF is weighted to adjust for non-response and to gross to population estimates. The non-response component is calculated using 2001 Census-linked data and the grossing component is calculated using population projections based on the 2001 Census. More information on the calculation of these weights can be found in Family Spending: A report on the 2007 Expenditure and Food Survey:
http://www.statistics.gov.uk/downloads/theme_social/family_spending_2007/familyspending2008_web.pdf

The 2008 LCF dataset contains two weights: weighta and weightq. Weighta is an annual weight and weightq is a quarterly weight. The quarterly weight was introduced because sample sizes vary from quarter to quarter as a result of re-issuing addresses where there had been a non-contact or refusal to a new interviewer after an interval of a few months, so that there are more interviews in the later quarters of the year than in the first quarter. Spending patterns are seasonal and quarterly grossing counteracts any bias from the uneven spread of interviews through the year.

For recent documentation please see the following link.
http://www.esds.ac.uk/findingData/snDescription.asp?sn=6385

The ESDS Government Introductory guide to the EFS also contains information on the weights in this survey. The guide can be downloaded from:

http://www.esds.ac.uk/government/docs/efsguide.pdf

3.11 Family Expenditure Survey

Since 1998/99 the FES data has used one weight which adjusts for non-response and grosses to population estimates.

The 2000-2001 weighting variable is called "weight". Appendix F of the 2000 FES Report 'Family Spending'[9] contains further details of the weights.

3.12 Heath Survey for England

Weighting variables are year specific owing to the variable sample design and the survey topic. For example, in 2000 weights are added for different probabilities of selection in care homes - see the 2000 User Guide[10]. Similarly, in 2002, the survey included a boosted sample of children and young people and mothers of infants aged under 1. For analysis of the HSE in 2002 no weights need to be applied if only using the adult sample. However, if using the boost sample (on its own or together with the adult sample) a sample design weight which accounts for unequal probabilities of selection needs to be applied (tablewt). Other years of the HSE that include boost samples are 1999 and 2004 (ethinc minority groups), 2000 and 2005 (older people – including some instiutional coverage) and 1997 (children and young people.

In 2003, non-response weighting was introduced to the HSE data. Although the HSE has generally presented a good match to the population, this decision was taken to keep up with the recent changes on many large-scale government sponsored surveys, and with the aim of reducing the possible biases.

The 2008 HSE follows the same general weighting strategy as developed in 2003. Four sets of non-response weights have been generated and these described in the table overleaf:

[12] http://www.statistics.gov.uk/downloads/theme_social/Family_Spending_2001-02_revised/Family_Spending_revised.pdf

[9] http://www.statistics.gov.uk/downloads/theme_social/Family_Spending_2000-01/Family_Spending_2000-01.pdf

[10] http://www.esds.ac.uk/doc/4487%5Cmrdoc%5Cpdf%5Ca4487uab.pdf

Description of weights in the Health Survey for England

Weight	Decription	When to use
Wt_hhld	household weight that corrects the distribution of household members to match population estimates for sex/age groups and GOR	Use during household analysis
Wt_int	Weights that include the wthhld component and a component to correct for bias resulting from individual non-response within households	Use during individual level analysis
Wt_nurse	Corrects for non-response to the nurse visit	Use on all analysis of questions asked during the nurse visit.
Wt_blood	A blood weight has been generated for all adults who had a nurse visit, were eligible for and agreed or were able to give a blood sample.	Use on all analysis of questions asked relating to blood samples
Wt_continine	A saliva weight has been generated for all adults and children that are aged 4-15yrs who had a nurse visit and were eligible for a saliva sample.	Use during analysis of questions asked relating to saliva samples.
Wt_hhld_acc	Only respondents in a sub-sample of the selected core addresses were eligible to be selected to wear an accelerometer. This required an additional set of calibration weights	Use during analysis of the accelerometer data
Wt_int_acc	Calibration weights for the analysis of the interview data from accelerometer sample.	Use during individual level analysis from the accelerometer sample
Wt_nurse_acc	Calibration weights for the analysis of the nurse data from accelerometer sample.	Use during analysis of nurse data from the accelerometer sample
Wt_blood_acc	Calibration weights for the analysis of the blood data from accelerometer sample.	Use during analysis of blood data from the accelerometer sample
Wt_continine_acc	Calibration weights for the analysis of the saliva data from accelerometer sample.	Use during analysis of saliva data from the accelerometer sample

Not all respondents were eligle or agreed to a nurse visit. Of those who did have a nurse visit not all agreed to give a blood or saliva sample. It is important to note that if using the nurse data you should use only the variable wt_nurse as this overrides the individual weight wt_int. Similarly, if using the blood or saliva data you should use the blood/saliva weight variables as these include all other weighting components. The accelerometer variable includes a component calculated in the same way as wthhld and another weighting component that adjusts for the fact that not all households and not all members of selected household are eligible for the accelerometer part of the survey.

For more information on the weights in the HSE and their calculation see the methods and documentation report:

http://www.ic.nhs.uk/webfiles/publications/HSE/HSE08/Volume_2_Methods_and_Documentation.pdf

Latest user guide
http://www.data-archive.ac.uk/doc/6397%5Cmrdoc%5Cpdf%5C6397userguide.pdf

3.13 Survey of English Housing

In April 2008 the Survey of English Housing (SEH) merged with the English House Condition Survey (EHCS) to form the new English Housing Survey (EHS). The final fieldwork year for the SEH was 2007/08. To find out more go to the EHS section of the Communities and Local Government web site.

The SEH has been weighted since 1994/95 to produce population estimates and to compensate for different response rates among households. The 2007-2008 dataset has two weight variables (H4b and H4bt), both of which combine weights for non-response and grossing. H4b weights for non-response and grosses to households in England (in 000s) and h4bt: weights for non-response and grosses to tenancy groups in England (in 000s). For further information see the following document. http://www.esds.ac.uk/doc/5021/mrdoc/pdf/5021userguide2.pdf

There are several stages for grossing. The first is to use the sampling fraction and response rate. Broadly, if the end result of sampling and non-response is that there is an interview for one in a thousand households, the grossing factor is one thousand. The initial grossing compensates for different response rates among households that were more or less difficult to find at home, measured by the number of calls needed to make contact. Households that were harder to contact receive a bigger grossing factor than those that were easier to contact (see "Sampling fraction and response rate" below).

The remaining stages adjust the factors so that there is an exact match with population estimates, separately for males and females and for broad age groups. An important feature of the SEH grossing is that this is done by adjusting the factors for whole households, not by adjusting the factors for individuals. The population figures being matched are those for the household population and exclude people who are not covered by the SEH that is those in bed-and-breakfast accommodation, hostels, residential care homes and other institutions. There is a final stage which applies only to private tenancy groups. This compensates for the small dropout between the main stage of the survey and the private renters module.

Latest userguide 2007/8: http://www.esds.ac.uk/doc/6399/mrdoc/pdf/6399userguide.pdf

3.14 National Travel Survey

A weighting strategy for the NTS was developed following a recommendation in the 2000 National Statistics Quality Review of the NTS. For the first time, the 2005 NTS results were based on weighted data. The weighting methodology has been applied to data back to 1995 and all NTS figures for 1995 onwards which are published or released are now based on weighted data. As well as adjusting for non-response bias, the weighting strategy for the NTS also adjusts for the drop-off in the number of trips recorded by respondents during the course of the travel week; for uneven recording of short walks by day of the week and for the short-fall in reporting of long distance trips. Therefore, there are several sets of weights which apply to different levels of the database; household, trip and long distance journey**. It is important to select the correct weights for each analysis. Initial results should be checked against published data to ensure weights are being applied correctly.**

Following the introduction of the weighting strategy there are now two samples which can be used for analysis. Analysis of travel data is based on the diary sample (contains all fully co-operating households) which includes weights that adjust for non-response and, at the trip-level, adjust for drop-off in recording observed during the seven day travel week. Analyses at household, individual and vehicle level are based on the interview sample (contains fully and partially co-operating households).

The following weighting variables are available:

W1 - Unweighted diary sample - this gives unweighted results for the diary sample only. (This is equivalent to the results produced before the weighting strategy was introduced and can be used to generate unweighted sample sizes for analysis of the diary sample. It is effectively the same as the 'status' variable mentioned above)
W2 - Diary sample household weight - apply to all analysis of the diary sample at household, individual and vehicle level.
W3 - Interview sample household weight - apply to all analysis of the interview sample at household, individual and vehicle level.
W4 - LDJ weight incorporating household weight - apply to all analysis at long distance journey (LDJ) level
W4xhh - LDJ weight excluding household weight
W5 - trip/stage weight - apply to all analysis of trip/stage data
W5xhh - Trip/stage weight excluding household weight
No weighting variable - if no weighting variable is applied, this gives unweighted results for the interview sample.

For most analyses at household, individual and vehicle level, w3 should be applied. For most analyses of travel patterns, w5 should be applied to trip/stage data and w2 should be applied at the individual level in order to calculate rates.

Examples of applying weights:

- To generate trip rates, apply w5 to trip data and apply w2 to individual data (i.e. Diary sample)
- To calculate household car ownership - apply w3 to the household data (Interview sample)
- To calculate the proportion of driving licence holders - apply w3 to the individual data (Interview sample)
- To determine the unweighted sample size for trip rate analysis, apply w1 to the trip data and to the individual data (Diary sample)
- To determine the unweighted sample size for household car ownership or driving licence figures, apply no weights. (Interview sample)

The 2002-2006 NTS userguide provides more information on NTS weights (see Non-response and drop-off weighting section) and is available for download at:

http://www.esds.ac.uk/doc/5340%5Cmrdoc%5Cpdf%5C5340userguide.pdf

Further information on the weighting methodology, together with analysis comparing weighted and unweighted data, is available in the Methodology section at: www.dft.gov.uk/transtat/personaltravel.

In addition to the above, it is important to note that special weights for 'short walks' should be applied when analysing data relating to trips of less than one mile. Because trips of less than one mile in distance are recorded only on the seventh day of the travel week, these trips must be weighted by a factor of seven when analysed. Also for consistency with earlier surveys 'series of calls' trips are excluded from analysis of stage and trip counts and time. Therefore, one of a number of 'short walkweights' must be applied to any tabulations using trip or stage counts, distance or time. Several 'short walkweights' have been provided and page 4 of the 2002-2006 user guide provides more information on these.

3.15 National Food Survey

The weighting used in the National Food Survey is for Northern Ireland. Prior to inclusion of Northern Ireland (1996) there was no weighting. The weight accounts for the deliberate oversampling of Northern Ireland and for differential response rates among different household types. This is described in detail in the NFS User Guide[11]. The datasets for 1996 onwards contain an Excel file called nfsweights.xls which gives the weights that users should add to the files if using the NI data.

Weights for NFS Data 1996-2000 can be found in the following link
http://www.esds.ac.uk/doc/4512/mrdoc/excel/nfsweights.xls

3.16 Family Resources Survey

Since 1992 the FRS has used one weighting variable for two purposes (1) to gross to population (2) to compensate for non-response - in the FRS 2007-2008 the weighting variable is called Gross3. The 1994-1995 to 2001-2002 datasets were re-released due to the inclusion of a new (interim) grossing factor introduced to make adjustments to the FRS for low income households in Scotland. These datasets contain two weighting variables: Gross1 is the original variable and Gross2 is the new variable. From 2003-04 onwards there have been revisions to the grossing scheme. Revised grossing factors, incorporating both the new grossing regime and the revised population counts, have been calculated for all the years for which full-year FRS data is available, from 1994-95 onwards. - see the Grossing Review information in the FRS User Guide 1[12] for more information.

Latest FRS userguide:
http://www.esds.ac.uk/doc/6252/mrdoc/pdf/frs_2007_08_introduction_doc.pdf

[11] http://www.esds.ac.uk/doc/4512/mrdoc/pdf/a4512uab.pdf
[12] http://www.esds.ac.uk/doc/6252/mrdoc/pdf/frs_2007_08_introduction_doc.pdf

3.17 Time Use Survey

The TUS uses weighting for a variety of reasons. There are different weights on the different files (individual questionnaire file, worksheet file, household questionnaire file and diary file). For more information go to the Time Use 2000 User Guide[13].

- There are 2 individual questionnaire weights: both weights compensate for non-response and are calibrated to UK population characteristics for age-group, sex and region. The difference between the two weights is that one grosses to the UK population and the other does not. (1) wtpq_ug is the ungrossed weight which weights to the achieved sample size (2) wtpq_gr is the grossed weight which weights to UK population of those aged 8yrs or more living in private households.
- There are 2 worksheet weights: as individual weights (1) wtwrk_ug is ungrossed (2) wtwrk_gr is grossed.
- There are two diary weights: as individual weights but also compensates for differential sampling of weekdays and weekends (1) wtdwh_ug is ungrossed weight (2) wtdwh_gr is grossed.
- There are six household questionnaire weights: as individual weights but two separate weights for each of following:
  - households with dairy-keepers (1) wtdh_ug is ungrossed (2) wtdg_gr is grossed
  - households with worksheet-keepers (3) wtwh_ug is ungrossed (4) wtwg_gr is grossed
  - households with diary and worksheet-keepers (5) wtdh_ug is ungrossed (6) wtdg_gr is grossed

3.18 Opinions (formerly the ONS Omnibus Survey)

The Opinions survey weights for unequal probabilities of selection caused by interviewing only one adult per household, or restricting the eligibility of the module to certain types of respondent. The weighting system also adjusts for some non-response bias by calibrating the Opinions sample to ONS population totals using the Generalized Estimation System (GES) programme. The February 2007 dataset has two weights (indwgt and hhwgt). Indwgt should be applied if the unit of analysis is the individual because the weight makes the sample representative of British adults. Hhwgt should be applied if the unit of analysis is the household reference person or spouse.

For recent documentation see the following link.
http://www.esds.ac.uk/findingData/snDescription.asp?sn=5813

For a copy of the Opinions Technical Report contact the Omnibus team on Omnibus@ons.gov.uk

[13] http://www.esds.ac.uk/doc/4504/mrdoc/pdf/4504userguide1.pdf

4: References and resources

4.1 Bibliography

References cited:

Barton, J. (2001) Appendix D: Living in Britain see section 4.3

Barton, J. (2004) Weighting the Social Surveys, slides from presentation, available online at http://www.ccsr.ac.uk/esds/events/2004-03-12/slides.shtml

Palmer, N. Hughes, M. Labour Force Survey: Reweighting and seasonal adjustment review 2008. *Economic and Labour Market Review.* 2(6) p33-42

4.2 General reading on sampling and sampling weights

Barnett, V. (2002) *Sample Survey Principles and Methods* London: Hodder Arnold

Barton, J. (2001) 'Developing a weighting and grossing system for the GHS' *Survey methodology bulletin 49*

Butcher, B. (1984) 'Grossing Up - when and how' *Survey Methodology Bulletin* 14

Elliot, D (1991) *Weighting for Non-response: A Survey Researcher's Guide* OPCS Social Survey Division

Elliot, D (1996) 'The Presentation of Weighted Data in Survey Report Tables' *Survey Methodology Bulletin 38*

Elliot, D (1999) *Report of the Task Force on Weighting and Estimation* GSS Methodology Series 16 London: Government Statistical Service

Foster, K. (1998) *Evaluating nonresponse on Household Surveys* GSS Methodology Series 8: London Government Statistical Service

Lynn, P. (2004) 'Weighting' in Kimberly Kempf-Leonard *Encyclopedia of Social Measurement*, pp 967-974. London: Academic Press.

P|E|A|S (Practical Exemplars and Survey Analysis)
http://www.restore.ac.uk/PEAS/about.php

ESDS Guides:
Topic Guides: ESDS Government produces an annual topic-oriented guide to the major cross-sectional surveys. In 2003 this was based on Employment and the Labour Market. The guide contains a summary of weighting schemes used in the surveys and clickable links to relevant documentation for individual surveys. This is available on the ESDS web pages from: http://www.esds.ac.uk/government/docs/ .

Other key documents: This page also contains a link to the GSS' 1999 Report of the Taskforce on Weighting and Estimation. The appendix of this document reviews contemporary weighting schema for a range of surveys.

Survey specific resources
All surveys have documentation available. This should be obtained with the data and consulted before the using the datasets.

General Lifestyle Survey: Appendix D of the 2007 GLF report contains guidance on how weights have been produced for the GHS, and their effect on results. This can be found at http://www.statistics.gov.uk/StatBase/Product.asp?vlnk=5756&Pos=2&ColRank=1&Rank=272

Labour Force Survey: Weights are available separately for different purposes including Individual analyses and Income on the QLFS general file, Household level analyses in the Household file and for users of the longitudinal data. Information on these are available in the appropriate documentation available from the UK Data Archive. The most recent of these can be found at the following locations:

- QLFS (Individual and Income data):

http://www.esds.ac.uk/doc/6199/mrdoc/pdf/background.pdf

- Household data:
  http://www.esds.ac.uk/doc/6292/mrdoc/pdf/lfs_vol1_background2009.pdf
  http://www.esds.ac.uk/doc/6292/mrdoc/pdf/lfs_vol8_household2008.pdf
-
- Longitudinal data:
  http://www.esds.ac.uk/doc/6201/mrdoc/pdf/background.pdf
  http://www.esds.ac.uk/doc/6201/mrdoc/pdf/longitudinal.pdf

Guidance on the effect of regrossing in the light of updated population estimates is available in the LFS documentation:

http://www.esds.ac.uk/doc/6199/mrdoc/pdf/lfs_reweighting_2009.pdf

Living Costs and Food Survey: A description of the weighting scheme used in the LCF is available in Appendix B6 of 'Family Spending 2009 edition' a report on the 2008 Living Costs and Food Survey. This is available online at:
http://www.statistics.gov.uk/downloads/theme_social/Family-Spending-2008/FamilySpending2009.pdf

Appendix

*Weighting Data in SPSS***

The WEIGHT command simulates case replication by treating each case as if it were actually the number of cases indicated by the value of the weight variable. You can use a weight variable to adjust the distribution of cases to more accurately reflect the larger population or to simulate raw data from aggregated data.

Example
A sample data file contains 52% males and 48% females, but you know that in the larger population the real distribution is 49% males and 51% females. You can compute and apply a weight variable to simulate this distribution.

```
*weight_sample.sps.
***create sample data of 52 males, 48 females***.
NEW FILE.
INPUT PROGRAM.
- STRING gender (A6).
- LOOP #I =1 TO 100.
- DO IF #I <= 52.
- COMPUTE gender='Male'.
- ELSE.
- COMPUTE Gender='Female'.
- END IF.
- COMPUTE AgeCategory = trunc(uniform(3)+1).
- END CASE.
- END LOOP.
- END FILE.
END INPUT PROGRAM.
FREQUENCIES VARIABLES=gender AgeCategory.
***create and apply weightvar***.
***to simulate 49 males, 51 females***.
DO IF gender = 'Male'.
- COMPUTE weightvar=49/52.
ELSE IF gender = 'Female'.
- COMPUTE weightvar=51/48.
END IF.
WEIGHT BY weightvar.
FREQUENCIES VARIABLES=gender AgeCategory.
```

- Everything prior to the first FREQUENCIES command simply generates a sample dataset with 52 males and 48 females.

File Operations

- The DO IF structure sets one value of *weightvar* for males and a different value for females. The formula used here is: *desired proportion/observed proportion*. For males, it is 49/52 (0.94), and for females, it is 51/48 (1.06).
- The WEIGHT command weights cases by the value of *weightvar*, and the second FREQUENCIES command displays the weighted distribution.

Note: In this example, the weight values have been calculated in a manner that does not alter the total number of cases. If the weighted number of cases exceeds the original number of cases, tests of significance are inflated; if it is smaller, they are deflated. More flexible and reliable weighting techniques are available in the Complex Samples add-on module.

Example

You want to calculate measures of association and/or significance tests for a crosstabulation, but all you have to work with is the summary table, not the raw data used to construct the table. The table looks like this:

	Male	Female	Total
Under $50K	25	35	60
$50K+	30	10	40
Total	55	45	100

You then read the data into SPSS, using rows, columns, and cell counts as variables; then, use the cell count variable as a weight variable.

```
*weight.sps.
DATA LIST LIST /Income Gender count.
BEGIN DATA
1, 1, 25
1, 2, 35
2, 1, 30
2, 2, 10
END DATA.
VALUE LABELS
Income 1 'Under $50K' 2 '$50K+'
/Gender 1 'Male' 2 'Female'.
WEIGHT BY count.
CROSSTABS TABLES=Income by Gender
/STATISTICS=CC PHI.
```

- The values for *Income* and *Gender* represent the row and column positions from the original table, and *count* is the value that appears in the corresponding cell in the table. For example, 1, 2, 35 indicate that the value in the first row, second column is 35. (The *Total* row and column are not included.)
- The VALUE LABELS command assigns descriptive labels to the numeric codes for *Income* and *Gender*. In this example, the value labels are the row and column labels from the original table.
- The WEIGHT command weights cases by the value of *count*, which is the number of cases in each cell of the original table.
- The CROSSTABS command produces a table very similar to the original and provides statistical tests of association and significance.

Crosstabulation and significance tests for reconstructed table

Income * Gender Crosstabulation

		Gender		
		Male	Female	Total
Income	Under $50K	25	35	60
	$50K+	30	10	40
Total		55	45	100

Symmetric Measures

		Value	Approx. Sig.
Nominal by Nominal	Phi	-.328	.001
	Cramer's V	.328	.001
	Contingency Coefficient	.312	.001
N of Valid Cases		100	

** This is extracted from Chapter 4 p.83-84 SPSS Programming and Data Management, 3rd Edition A Guide for SPSS and SAS® Users Raynald Levesque and SPSS Inc.