Appendix BB
Sampling and weighting for National Diet and Nutrition Survey 2019 to 2023

BB.1 Sample design

The National Diet and Nutrition Survey (NDNS) sample covers the whole of the UK and is made up of a core sample and a Northern Ireland 'boost' sample.

The aim of NDNS is to achieve a representative sample of 500 adults (aged 19+ years) and 500 children (aged 1.5 to 18 years) per year from the core UK sample. In addition, the Northern Ireland (NI) boost sample is required to yield a sample of 100 adults and 100 children per year.

To achieve the even split between adults and children, the sample is divided into address types: 'basic' addresses, where up to 2 adults (aged 19 years and over) and 1 child (aged 1.5 to 18 years) are selected; and 'young person' addresses, where up to 2 children are selected from different NDNS age groups (1.5 to 3 years; 4 to 10 years; 11 to 18 years). The proportion of the sample assigned as 'basic' addresses was approximately 28% in years 12 to 14, and 33% in year 15, see table BB.2 below.

The participant selection method was adjusted from year 12 onwards to provide a more cost-efficient sample design which increased the potential number of sampled participants per household. Further changes to the sample design were then necessary from year 13 onwards to account for lower response rates due to the COVID-19 pandemic. These measures included increasing the number of addresses per Primary Sampling Unit (PSU) (from 29 in years 12 to 14 to 33 in year 15) and adjusting the address type split.

BB.1.1 Selecting addresses

The sample for each year was drawn from the Postcode Address File. Table BB.1 summarises the number of PSUs for the core, NI boost, reserve and shortfall samples. Table BB.2 summarises the number of addresses for the core, NI boost, reserve and shortfall samples.

The core (and NI boost) sample comprised the main, issued sample. The reserve sample refers to the PSUs which would be issued in full (all addresses) or in part (not all addresses) should the number of achieved (productive) cases and expected response rate fall below target in a given fieldwork year. The shortfall sample refers to a further set of PSUs and addresses which were drawn in years 14 and 15 to make up the shortfall due to the additional impact of COVID-19 on response rates and achieved numbers. This 'shortfall sample' was issued in England, Scotland and Wales between January and March 2022. Note that year 12 fieldwork was suspended in March 2020 due to the COVID-19 pandemic and did not resume; a total of 2,553 addresses were issued to interviewers prior to the suspension.

Table BB.1: Summary of selected PSUs

Sample type	Year 12	Year 13	Year 14	Year 15
Core UK sample	115 PSUs	115 PSUs	115 PSUs	115 PSUs
NI boost sample	17 PSUs	17 PSUs	17 PSUs	17 PSUs
Reserve sample	18 PSUs	18 PSUs	18 PSUs	20 PSUs
Shortfall sample	not applicable	not applicable	95 PSUs	60 PSUs

Table BB.2: Summary of selected addresses

Sample type	Year 12	Year 13	Year 14	Year 15
Core UK sample	3,335 addresses	3,335 addresses	3,335 addresses	3,795 addresses
NI 'boost' sample	493 addresses	493 addresses	493 addresses	561 addresses
Reserve sample	522 addresses	522 addresses	522 addresses	660 addresses
Shortfall sample	not applicable	not applicable	3,800 addresses	1,980 addresses
Number of addresses per PSU	29 (8 basic, 21 young person)	29 (8 basic, 21 young person)	29 (8 basic, 21 young person)	33 (12 basic, 21 young person)

In each PSU, the addresses were randomly selected. At each address, the interviewer established the number of households and, in cases where there were 2 or more, selected 1 household at random.

BB.1.2 Selecting participants

As only around one-third of households contain children,[1] over-sampling of children was required to achieve equal numbers of adult (aged 19 years and over) and child (aged 18 months to 18 years) participants each fieldwork year. In order to achieve equal numbers of adults and children in the sample, the selected addresses in each PSU were randomly allocated to 1 of the 2 sample types: a 'basic' or 'young person' address. For years 12 to 14, there was a 'basic' sample of 8 addresses where interviewers randomly selected up to 2 adults and 1 child. In year 15, 12 addresses were randomly assigned to the 'basic' sample and the remaining 21 addresses in each year were assigned to the 'young person' sample. This involved screening to identify households with children and, at each of these addresses, up to 2 children, each from a different NDNS age group (18 months to 3 years, 4 to 10 years, 11 to 18 years), were selected for interview.

Child selection for households with three NDNS age groups represented changed for year 15 to simplify the process for fieldworkers. Through years 12 to 14 two different age groups were selected and then children were selected from those. For year 15 one child was selected at random first and the second child was selected from the

[1] https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/families/bulletins/familiesandhouseholds/2022#:~:text=In%202022%2C%20there%20were%2019.4,dependent%20children%20(2.9%20million)

remaining (different) age groups. While this only affected a small number of households it meant that the selection weight was calculated differently at year 15.

BB.2 Weighting the survey data

It is necessary to apply weighting factors to the data collected in the NDNS for 2 reasons: to remove any bias in the observed results which may be due to differences in the probability of households and individuals being selected to take part; and to attempt to reduce non-response bias. The survey was designed so that no more than 2 adults and up to 2 children were selected to take part from any 1 household. This meant that adults living in households with 1 or more other adults, and children in households with 1 or more other children were less likely to be selected than were adults or children in single adult or single child households.

In addition, there were a number of stages in the survey where it was possible for participants to drop out. If the people who refused to participate at a particular stage were systematically different from those who took part, then the sample would be biased.

Weighting factors were used to correct for both these cases. There were 2 stages to the weighting scheme: the first was to generate a set of design weights to correct for the unequal selection probabilities; and the second was to create a set of weights at each stage where dropping out makes a difference to adjust for non-response.

BB.3 Calculation of the general population weights

BB.3.1 Address selection weights

The sample design includes an adjustment for the number of addresses selected in each country (called add_selwt) as they are not selected in proportion to the UK population. All NI addresses, and therefore participants, are weighted down as a result due to the use of the NI boost. The applied weights put the 4 countries into their correct population proportions in the UK so that, for example, the percentage of the NDNS sample in NI is the same as the percentage of the UK population that is in NI. These weights are calculated as the number of addresses in each country divided by the number of addresses selected in each country.

BB.3.2 Dwelling unit[2] selection weights

Most addresses selected from the PAF contained a single dwelling unit, with a separate entrance. At addresses with more than one dwelling unit, only one was selected; interviewers carried out a selection procedure to identify which dwelling unit to include in the sample using a Kish grid.[3]

[2] A Dwelling Unit is an address or part of an address, which has its own front door which separates one part of the address from other parts (that is, only those who live behind the door have access to the area, it is not a communal part of the address). A Dwelling Unit need not be fully self-contained - for example, an address may contain four bedsitters, the inhabitants of whom share a bathroom. Each bed-sitter would count as a dwelling unit as long as it had its own front door.

[3] A Kish grid is a framework to ensure that the dwelling unit is selected without interviewer bias. The number of dwelling units is listed across the top of the grid, with a random number below to indicate which dwelling unit should be selected.

The dwelling unit selection weights (w_{du}) adjust for this selection at addresses with more than one dwelling unit. The weights were calculated as the number of dwelling units identified at the address.

The dwelling unit selection weights ensured that addresses containing more than 1 dwelling unit were not under-represented in the issued sample.

BB.3.3 Catering unit[4] selection weights

Most addresses selected from the PAF contained a single catering unit. At addresses with more than 1 catering unit, only 1 was selected; interviewers carried out a selection procedure to identify which catering unit to include in the sample using a Kish grid.[5]

The catering unit selection weights (w_{cu}) adjust for this selection at addresses with more than 1 catering unit. The weights were calculated as the number of catering units identified at the address.

The catering unit selection weights ensured that addresses containing more than 1 catering unit were not under-represented in the issued sample.

BB.3.4 Household selection weights

Overall household selection weights (hh_selwt) were calculated as the product of the dwelling unit selection weights (w_{du}) and catering unit selection weights (w_{cu}). The composite selection weights were trimmed at 3 to avoid any large values.

BB.3.5 Child selection weights

To ensure that that children in households with 2 or more children were correctly represented in the survey estimates, selection weights (wt_ind) were calculated reflecting their probability of selection. In basic sample addresses a single child or young person aged 18 months to 18 years (hereafter referred to as children for simplicity) could be selected for participation, so selection weights were calculated as the number of children within the household. In young person sample addresses up to 2 children could be selected. In years 12 to 14, these were selected from different NDNS age groups (18 months to 3 years; 4 to 10 years; 11 to 18 years). Selection weights were thus calculated for years 12 to 14 as the number of children within the household in their age group, multiplied by 1.5 if children from all 3 age groups were present in the household. For year 15 the selection weights were calculated in the same way as for years 12 to 14 unless children from all 3 child age groups were present in the household, in which case they were calculated as the number of children in the household multiplied by 1 plus the ratios of the number of children in the household in each of the other 2 age groups to the number of children not in that age group. These weights were trimmed at 4 to avoid any large weights.

---

[4] A Catering unit is a group of people who eat food that is bought and prepared for them (largely) as a group. Occasionally a household will be found to consist of more than one catering unit. Although people may share accommodation and even be related, they may not be in the same catering unit. For example, young adults sharing a house with their parents may shop, cook and eat by themselves, in which case the parents would be in one catering unit and the children in another.

[5] A Kish grid is a framework to ensure that the dwelling unit is selected without interviewer bias. The number of dwelling units is listed across the top of the grid, with a random number below to indicate which dwelling unit should be selected.

BB.3.6 Adult selection weights

Adult selection weights (also wt_ind) were simpler. They were calculated as the number of adults in the household divided by the number of adults selected. Participants from single-adult households therefore had a weight of 1. In practice these were made simultaneously with the child selection weights and received the same trim at 4 to avoid large weights.

BB.3.7 Combining the selection weights

The total person selection weights (ind_selwt) for the general population sample of adults and children were then calculated as:

ind_selwt = hh_selwt x wt_ind

These were trimmed at 4 to avoid large overall selection weights.

The final selection weights (Selwt) were calculated as:

Selwt= add_selwt x ind_selwt

The selection weights for all responding adults and children were re-scaled so that the weighted sample size is the same as the achieved sample size. Therefore, the final selection weights adjust for differences in the chance of selection for addresses, households and individuals.

BB.3.8 Individual calibration weighting

Calibration weighting was used to ensure that the weighted distribution of participants matched Office for National Statistics (ONS) 2021 mid-year population estimates for sex/age groups and region. The aim of the calibration weighting is to reduce bias resulting from the selection of individuals.

Calibration was conducted separately for adults and children and separately for each country, even being trimmed separately to avoid outlying weights within adults versus children. The starting weight was the final selection weight (selwt). The separate post-calibration weights were adjusted before combination into a single weights variable such that children are weighted down to represent the appropriate proportion of the UK population. This overall weight is then rescaled to a mean of 1.

Thus, the final calibrated individual weight (wti_Y1215) adjusts for address, dwelling unit, catering unit and household selection, and for the age,sex and region profiles of participating individuals.

BB.3.9 Biomedical visit weights

Participants completing at least one dietary recall were asked whether they were willing to have a visit from a biomedical fieldworker. The main focus of this visit was to take a blood sample, along with other measurements depending on the age of the participant.

Not all those interviewed went on to have a biomedical visit. This may have been the source of further non-response bias. For data relating to biomedical visits, a logistic regression model was fitted for all individual respondents, weighted by individual weight (wti_Y1215). The outcome variable was whether or not a biomedical visit was undertaken, and the models were made separately for adults and children. The adult model has the following as covariates: age group by sex, region, household

size, ethnicity and tenure. The child model has the following as covariates: age group by sex, region, household size, ethnicity, tenure and number of children in the household.

The weights for non-response to the biomedical visit (w5ad for adults and w5ch for children) were calculated as the reciprocal of the predicted probability of a biomedical visit being undertaken, estimated from the respective regression model.

The weights were trimmed by 1% at the tail to remove extremely high weights. The weights for the biomedical visit sample were calculated as wt_nurse_ad = wti_Y1215 x w5ad for adults and wt_nurse_ch = wti_Y1215 x w5ch for children. The weights resulting from the adults and child models were then scaled up to the population for their respective age groups and combined into one variable, thus weighting children down to represent the appropriate proportion of the UK population. These weights were re-scaled so that the weighted sample size for the biomedical visit is the same as the achieved sample size. They adjust for selection, non-response and population profile for the sample that had a biomedical visit.

BB.3.10 Blood weights

Almost all adults who had a biomedical visit were eligible to have a blood sample taken but not all those eligible agreed or were able to do so. A logistic regression model was fitted, weighted by the biomedical visit weight (wt_nurse). The outcome variable was whether or not a usable blood sample was analysed, and the models were made separately for adults and children. The adult model has the following as covariates: age group by sex, region, household size, number of children in the household, marital status, academic qualifications and regularity of alcohol consumption. The child model has the following as covariates: age group by sex, region, household size and number of children in the household.

The weights for non-participation for the blood sample (w6ad for adults and w6ch for children) were calculated as the reciprocal of the predicted probability of a blood sample being obtained, estimated from the respective regression model.

The weights were trimmed by 1% at the tail to remove extremely high weights. The weights for the blood sample were calculated as wt_blood_ad = wt_nurse x w6ad for adults and wt_blood_ch = wt_nurse x w6ch for children. The weights resulting from the adults and child models were then scaled up to the population for their respective age groups and combined into one variable, thus weighting children down to represent the appropriate proportion of the UK population. These weights were re-scaled so that the weighted blood sample size was the same as the achieved sample size.

BB.3.11 Spot urine weights

Those aged 4 years and over who had either a biomedical or face-to-face interviewer visit were eligible to provide a spot urine sample, but not all those eligible agreed or did so. A logistic regression model was fitted for all NDNS participants, weighted by the individual weights. The outcome variable was whether or not a usable spot urine sample was analysed, and the models were made separately for adults and children. The adult model has the following as covariates: age group by sex, region, ethnicity and tenure. The child model has the following as covariates: age group by sex

region, household size, number of children in the household and working status of representative person. These models are doing more than 1 job – modelling non-response on face-to-face contact (either by an interviewer or a biomedical fieldworker) as well as the urine sample non-response and validity, all in one step.

The weights were trimmed by 1% at the tail to remove extremely high weights. The weights for participation for the urine sample (w7ad for adults and w7ch for children) were calculated as the reciprocal of the predicted probability of a urine sample being obtained, estimated from the respective regression model. The final spot urine weights were calculated as wt_urine_ad = wti_Y1215 x w7ad for adults 19+ and wt_urine_ch = wti_Y1215 x w7ch for children. The weights resulting from the adults and child models were then scaled up to the population for their respective age groups (which are different for spot urine) and combined into 1 variable, thus weighting children down to represent the appropriate proportion of the UK population. These weights were re-scaled so that the weighted spot urine sample size was the same as the achieved sample size.

BB.3.12 Dietary recall weights

Individuals were asked to recall their dietary intake for 4 separate days. Not all those interviewed went on to record dietary recall days. This may have been the source of further non-response bias. For data relating to diet recalls, a logistic regression model was fitted for all individual respondents, weighted by individual weight (wti_Y1215). The outcome variable was whether or not at least 1 recall day was recorded, and the models were made for adults and children together. The non-response model has the following as covariates: age group by sex, region, household size, ethnicity, number of children in the household, tenure, working status of representative person, survey year, quintiles of population density and quintiles of Index of Multiple Deprivation within country.

The weights for non-response to the dietary recalls (w8) were calculated as the reciprocal of the predicted probability of at least one dietary recall being completed, estimated from the respective regression model.

The weights were trimmed at the 1% tail to remove extremely high weights. The weights for the diet recalls data were calculated as wt_diet = wti_Y1215 x w8. These weights were re-scaled so that the weighted blood sample size was the same as the achieved sample size.

BB.3.13 Physical Activity Questionnaire weights

Only those aged 4 years and over with at least 3 dietary recalls were eligible to complete the Physical Activity Questionnaire (PAQ), but not all those eligible agreed or did so. A logistic regression model was fitted, weighted by the diet weights. The outcome variable was whether or not the PAQ was completed. Adults and children were given separate weights here, and the threshold for being considered an adult for the PAQ decreased from 19 to 16 years.

For those aged 16 years and over, the following were included as covariates: age group by sex, region, number of children, tenure and marital status. For children aged 4 to 15 years, the following were included as covariates: age group by sex, region, household size and tenure. These models are doing more than 1 job:

modelling non-response on the number of dietary recalls (as well as non-response to the PAQ itself in 1 step.

The PAQ weights (w9) were calculated as the reciprocal of the predicted probability of PAQ completion, estimated from the regression model. The weights were trimmed by 1% at the tail to remove extremely high weights. The final PAQ weights were calculated as wt_RPAQ = wt_diet x w9 for adults and wt_CPAQ = wt_diet x w9 for children.

BB.4 Calculation of the Northern Ireland (NI) weights

Since the UK weights for years 12 to 15 contain all cases who responded to each respective section of NDNS, the NI weights are simply the UK weights but for NI cases only. Outliers were trimmed, and the weights were rescaled to a mean of 1. These weights were checked against the NI mid-year population estimates of age and sex and are indeed still representative of the general population.

Note that weights for NI have also been created for years 10 to 11 in a very similar way, rescaled from the weights for years 9 to 11. Years 9 to 11 followed a similar approach to years 12 to 15 but with a different way of making the selection weights to reflect a different selection process and no need for diet weights as participation in that element was a condition to be considered a productive NDNS participant, so for years 9 to 11 the individual weight does the same job.

BB.5 Effect of the weights on the precision of the estimates

A design factor (DEFT) for each weight has been calculated to provide an approximate guide to the effect of the weighting on the precision of estimates. The DEFT is calculated as the average squared weight divided by the square of the average weight.

For instance, the DEFT of 1.36 for the individual weight indicates that the standard error of estimates is assumed to increase by 36%, with a corresponding loss of precision. Consequently, these weighted estimates have the same level of precision as an estimate based on a simple random sample, unweighted, of around 54% of the size of the actual sample. This is known as the effective sample size.

BB.5 Selecting the appropriate weight

The weight to use depends on the part of NDNS to be analysed:

- For overall analysis and analysis of the original survey, use the individual weight (wti_Y1215)
- For analysis of measurements and questions from the biomedical visit (except spot urine and blood samples), use the biomedical weight (wt_nurse)
- For analysis of the blood samples and associated findings, use the blood weight (wt_blood)
- For analysis of the spot urine samples and associated findings, use the urine weight (wt_urine)
- For analysis of the dietary recalls, use the diet weight (wt_diet)
- For analysis of the PAQ, use the adult weight for those aged 16 years and over (wt_RPAQ) and the child weight for those aged 4 to 15 years (wt_CPAQ)

The names are self-explanatory. The individual and biomedical weights are unlikely to be used unless it is for very specific analyses but they form the basis for subsequent weights such as the blood and diet weights.