Stats NZ
Tatauranga Aotearoa

Household Labour Force Survey sources and methods: 2016

New Zealand Government

Citation
Stats NZ (2017). *Household Labour Force Survey sources and methods: 2016*. Retrieved from www.stats.govt.nz.

ISBN 978-0-908350-90-2 (online)

Published in April 2017 by
Stats NZ Tatauranga Aotearoa
Wellington, New Zealand

Contact
Stats NZ Information Centre: info@stats.govt.nz
Phone toll-free 0508 525 525
Phone international +64 4 931 4600

www.stats.govt.nz

Contents

List of tables, figures, and abbreviations

List of tables

Appendix tables

List of figures

Appendix figures

Abbreviations list

ANZSIC	Australia and New Zealand standard industrial classification
CAI	computer-assisted interview
CAPI	computer-assisted personal interview
CATI	computer-assisted telephone interview
GSS	General Social Survey
HLFS	Household Labour Force Survey
NILF	not in the labour force
NPAD	National Postal Address Database (New Zealand Post)
NPD	non-private dwellings
NZDep	New Zealand Deprivation Index
NZIS	New Zealand Income Survey
NZSCO	New Zealand Standard Classification of Occupations
PPS	probability proportional to size
PNL	pre-notification letter
SRS	simple random sampling
PSU	primary sampling unit
TA	territorial authority

Purpose and introduction to the Household Labour Force Survey

Purpose

Household Labour Force Survey sources and methods: 2016 updates an older version of the same paper, Household Labour Force Survey sources and methods: 2015, to incorporate changes made from the June 2016 quarter, including the introduction of a new questionnaire. It provides Household Labour Force Survey (HLFS) customers with a technical description of the sample design for the HLFS. It also covers technical details on other aspects of the survey, including collection methodology, estimation and imputation methodology, and coding and processing details (from the June 2016 quarter).

We have phased in the new sample over two years, with the first one-eighth of the sample rotated into the field in the December 2014 quarter. See Changes affecting data comparability over time for major changes made to the HLFS, both in the past and as a result of this redevelopment. As of the September 2016 quarter the sample is entirely made up of this new sample.

About the HLFS

The HLFS was introduced in October 1985. It is a continuous national survey of households, which measures quarterly average levels of employment, unemployment, non-participation in the labour force, and the quarterly and annual changes in these levels. The purpose of the HLFS is to enable development and monitoring of the labour market and social policy, support research, and to help inform on the quality of employment and the health and general well-being of New Zealand's economy.

We collect responses from around 15,000 households every quarter, amounting to responses from approximately 30,000 individuals aged 15 and over. The HLFS is a rotating panel survey, which means we interview the same respondents over a set number of consecutive quarters and then replace them (on a rotating basis) by a new set of respondents. Households are selected using a multistage stratified, clustered design, and we interview all adults in selected households.

HLFS objectives

The objectives for the HLFS are to:

- Measure the levels, changes, and characteristics of employment, unemployment, and people not in the labour force (NILF) in New Zealand on a regular basis – using international guidelines and best practice. Specifically, the HLFS aims to:
    - produce reliable national and subnational estimates of the employed, unemployed, and NILF groups over time
    - produce reliable estimates of change in national and subnational estimates of the employed, unemployed, and NILF groups over time
    - collect and produce supporting information to describe, monitor, and explain the state of the labour market (eg hours of work, status in employment, duration of unemployment)
    - collect and produce information about the structure and characteristics of the employed, unemployed, and NILF groups.
- Collect and produce information on topics relevant to labour market data, such as qualification level and involvement in education and training. This may include regular and irregular attachment of topics covering labour market-related information.

Output from the HLFS

Stats NZ produces quarterly information releases (around four weeks after the end of the quarter) that contain information about the levels of employment and unemployment, by national and subnational breakdowns. Specifically, each quarter we produce a set of standard tables that contain the following figures:

- people employed, unemployed, and not in the labour force – actual, seasonally adjusted, and trend
- total people employed, unemployed, and not in the labour force – by age group, ethnic group, and regional council area
- people employed – by sex, and industrial activity of place of employment
- people employed – by sex and status in employment
- people employed full-time and part-time – seasonally adjusted
- the underutilised
- total usual and actual hours worked
- people underemployed by sex
- people employed, unemployed, and not in the labour force – by sex and formal study status
- people aged 15–24 years not in education, employment, or training (NEET).

Demographic characteristics we collect include age, sex, country of birth, ethnicity, educational attainment, and relationships of all household members (from which we derive family type). Other labour-related characteristics include industry, occupation, status in employment, full-time/part-time employment status, characteristics of previous job, methods of looking for work, and reasons for not being in the labour force. Labour force status is compiled in accordance with the International Labour Organization's definitions (see Labour force definitions for more details).

Labour force definitions

Employment relates to everyone in the working-age population who did one of the following during the reference week:

- worked for one hour or more for pay or profit in the context of an employee/employer relationship or self-employment
- worked without pay for one hour or more in work that contributed directly to the operation of a farm, business, or professional practice owned or operated by a relative (before April 1990 this was defined as 15 hours or more)
- had a job but was not at work due to illness or injury, personal or family responsibilities, bad weather or mechanical breakdown, direct involvement in industrial dispute, leave, or holiday.

Unemployment relates to everyone in the working-age population who, during the reference week, were not employed, were available for work, and:

- had actively sought work in the past four weeks ending with the reference week (only looking at job adverts is not counted as actively seeking) or
- had a new job to start within four weeks.

Not in the labour force relates to anyone in the working-age population who is neither employed nor unemployed (as defined above). This residual category includes people who:

- are retired
- have personal or family responsibilities (eg unpaid housework or childcare)
- attend educational institutions
- are permanently unable to work due to sickness, injury, or disability
- were temporarily unavailable for work in the survey reference week
- were not actively seeking work.

Scope and coverage

This chapter explains the target and survey populations, and the HLFS's usual residence guidelines.

Target population

The target population is the entire group from which you would ideally like to get information. The target population for the HLFS is the working-age population of New Zealand. We define this as "the non-institutionalised population 15 years and over, who usually live in New Zealand."

Specifically the target population excludes:

- people who have been living in New Zealand for less than 12 months, and who do not propose to stay in New Zealand for a total of 12 months or more
- long-term residents of homes for older people, hospitals, and psychiatric institutions (long-term is defined as six weeks or more)
- people in prison.

Survey population

The survey population consists of the group members (from the target population) who have a chance of being selected as part of the sample (ie they can be identified through the sampling frame). For the HLFS, we apply further exclusions to the target population to create the survey population (often due to cost and practical reasons), from which we then select the HLFS sample. These exclusions are a small percentage of the population and the bias introduced is minimal.

The survey population is the target population with the following exclusions. People:

- residing in non-private dwellings (eg people in hotels, motels, hostels, military camps)
- residing in non-permanent dwellings (eg people in tents or caravans not permanently sited)
- residing at wharves or landing places (eg people in ships, boats)
- residing on islands other than North, South, and Waiheke islands (eg people on Great Barrier, Kawau, Chatham, and Stewart islands).

Non-private dwellings

Dwellings are classified into two types – private and non-private.

Non-private dwellings (NPDs) include hotels, motels and guest accommodation, residential and community care facilities, hospitals, educational institutions, and prisons. A private dwelling is usually a self-contained housing unit that is not available for public use. The proportion of people staying in an NPD in New Zealand on census night in 2013 was 4.0 percent. The proportion living in NPDs (ie those who stated an NPD was their usual residence) was 2.3 percent in 2013.

The HLFS excludes NPDs from the survey population. This is because of the expense involved in surveying them and because the sampling and non-sampling errors for NPDs were much larger (partly due to poor response rates) when they were surveyed (before June 1995), compared with

the error from modelling the data (ie assuming the distribution of characteristics in the non-private sample is similar to that in the private sample).

We undertook analysis in 2013 to assess if the decision to not survey NPDs should be changed, and concluded we'd make no change at this time. Most NPDs contain very few residents (so excluding them does not result in significant bias), or they contain larger numbers of residents but the residents are only a small proportion of the people staying in these NPDs on any one night – this makes them costly to survey, due to low eligibility. Of NPD types with a high number of residents and a larger proportion of usual residents (eg educational institutions), we felt the potential bias (assessed using Census 2013 data) was not large enough (compared with the cost involved in surveying them) to warrant their inclusion in the sampling strategy for the HLFS.

Eurostat acknowledges the difficulties associated with sampling NPDs; it recognises that "for technical and methodological reasons it is not possible…to include the population living in collective households" (Eurostat, EU LFS Methods and Definitions 2001, p10). Their requirement for labour force surveys therefore is to provide results for private households only. Thus many of the labour force surveys run by European member states (including the United Kingdom) exclude communal establishments.

The sampling frames that the Australian, Canadian, and United States labour force surveys use are designed to represent the civilian non-institutionalised population. However, in Australia effort is made to include some NPDs by using a list sample that includes hotels and motels. The US survey (the Current Population Survey) also attempts to include such people by having a 'group quarter' stratum, containing housing units where residents share common facilities or receive formal care, on their sampling frame.

Avoiding double-counting

The HLFS surveys all people at a selected dwelling who consider themselves to usually reside at the selected dwelling. To usually reside at a dwelling, a person must consider they usually live at the dwelling.

In some situations a person may be unclear about where they usually reside. This generally occurs for people who live at more than one dwelling.

Usual residence guidelines

The HLFS guidelines below clarify where a person usually resides if they cannot decide this for themselves. These guidelines should reduce the probability of the HLFS double-counting people who live at more than one dwelling.

- People who are away overseas for more than six months, or plan to be away overseas for more than six months, do not usually reside at the surveyed dwelling.
- People who are temporarily staying at the surveyed dwelling do not usually reside at the surveyed dwelling (temporary is defined as staying less than six months).
- People who spend equal amounts of time residing at different addresses do usually reside at the surveyed dwelling.
- Dependent children in shared care usually reside at the place where they spend most nights. If they spend equal amounts of time at each residence, they usually reside at the surveyed dwelling.

- Dependent children who board at another residence to attend primary or secondary school usually reside at the dwelling of their parents/guardians.
- Tertiary students usually reside at the address where they live while studying. If they give up their residence in the holidays and return to the family home, their residence over the holiday period is the family home.
- If a person can't decide where they reside, they usually reside at the surveyed dwelling.

HLFS sample design

This chapter outlines the sample design for the most recently selected HLFS sample (first introduced in the December 2014 quarter).

Previous sample designs were established in 1986, 1991, 1996, and 2003. See Changes affecting data comparability over time for the main changes from the different sample designs.

The HLFS sample has a stratified design with two stages of clustering. Firstly we select a random sample of primary sampling units (PSUs) from each stratum (first stage of clustering), then we select a systematic sample of households from each PSU (second stage of clustering). Every person in a selected household aged 15 years and over is eligible for the survey (see Scope and coverage for exceptions).

PSUs are aggregations of one or more meshblocks, where meshblocks are the smallest geographical area unit in New Zealand. (In urban areas a meshblock is usually a block of residential area containing about 40 dwellings surrounded by streets; in rural areas a meshblock covers a much wider area because dwellings are sparsely spread.) PSUs constructed from the 2013 Census have an average of 70 occupied and under-construction dwellings.

Size and allocation to strata

Sample size

The new sample of the HLFS aims to achieve interviews with 15,000 households in the September 2016 quarter. One-eighth of the new sample of PSUs was rotated into the field in the December 2014 quarter, to be followed by another one-eighth in each subsequent quarter over a two-year period. Therefore the September 2016 quarter is the one in which the transition to the new sample will be complete.

To achieve the 15,000 interviewed households, we selected a sample of 20,165 total private dwellings (including occupied, under-construction, and vacant private dwellings) – from a total of around 1.8 million private dwellings in New Zealand – before the December 2014 quarter. This sample size of selected dwellings was based on three assumptions:

- We achieve interviews with 76 percent of the total number of selected dwellings (this includes roughly 90 percent of all private dwellings in the selected PSUs as being either occupied or under construction).
- A growth rate of dwellings in the selected PSUs of 1 percentage point each year.
- Five percent initial undercoverage, where the physical enumeration of the selected PSUs yields fewer dwellings than suggested by the census.

We selected 1,768 PSUs for the new sample, resulting in an average of 10.31 occupied or under-construction private dwellings per PSU. Because these selected PSUs will be used for both the HLFS and the General Social Survey (GSS) until the next census (2018), an equal number of dwellings is not selected from each PSU (see Panels for further discussion).

Note: Although the GSS uses the selected HLFS PSUs, the GSS sample will be a subset of these and have its own sample design.

Stratification

Stratification is the process of dividing the population (or survey frame) into homogeneous subgroups before sampling. Stratification is used for two different reasons. Firstly, to reduce sampling errors for survey estimates and to ensure that sample sizes for strata are of their expected size. Secondly, to target subgroups by disproportionate sampling (or over-sampling) certain strata.

Stratification for the new HLFS sample design includes five dimensions. PSUs are stratified by region, urban/rural status, a high-NILF (not in the labour force) status, groups based on New Zealand Deprivation Index values, and territorial authority (in that order).

The first four dimensions are explicit, or primary, strata (ie the sample is split by these groups and a random sample selected from each group), while the final dimension is implicit (PSUs are sorted by territorial authority within the primary strata and selected from the ordered list).

Region

New Zealand has 17 regional council areas, but some are combined for sampling purposes due to size. This stratification dimension or layer has 12 regions, with Gisborne and Hawke's Bay forming one combined region, and the West Coast, Marlborough, Nelson, and Tasman combined into one region. We use the same regions for disseminating the survey estimates. Regions are used as the first layer in the stratification hierarchy to ensure the sample's representativeness for regions – regional estimates are of high priority for the HLFS.

Urban and rural PSUs

The next layer of stratification is used to reduce fieldwork costs. Analysis of recent time and mileage data suggests that urban PSUs are cheaper than rural PSUs for fieldwork. Incorporating this layer into the stratification hierarchy, combined with disproportionate allocation (discussed below), results in cost savings for the HLFS.

The strata are formed based on main urban areas, using the Urban/rural profile (experimental) classification categories. We apply this classification at the meshblock level. If PSUs are made up of more than one meshblock, we use the modal urban/rural classification unless any meshblocks in the PSU are classified as highly rural or remote, in which case the PSU is classed as rural for the stratification.

NILF strata

Within region and urban/rural classification, we divide PSUs into two groups based on the percentage of individuals within the PSU who were not in the labour force (NILF) in the 2013 Census. Including this stratum improved variances for the estimates of both NILF and employed, and income.

PSUs with a density of NILF individuals greater than or equal to the 75th percentile (based on all PSUs in the population, approximately 36 percent) are placed in the 'high NILF' stratum. If splitting by low/high NILF density would result in any stratum containing fewer than 100 PSUs (within region and urban/rural classification), the strata formed by region and urban/rural classification were not split further by NILF density.

New Zealand Deprivation Index

Finally, strata are further split into groups based on New Zealand Socioeconomic Deprivation Index (NZDep) values. As with the NILF strata, including a socio-economic stratum is beneficial for the sampling errors of labour market estimates, because area deprivation is strongly correlated with labour market measures. Ensuring the sample has the correct area deprivation profile means there will be less sampling variance in the estimates. NZDep is calculated by the University of Otago at meshblock level.

For stratifying PSUs, we used the mean NZDep value where a PSU contains more than one meshblock. Superstrata (region by urban/rural status by high/low NILF status) are split equally – based on NZDep2013 values, up to a maximum of 10 groups, with a target size of 200 PSUs in each stratum. If this results in any stratum having fewer than 100 PSUs, the superstrata are not split.

Appendix table 1 shows the number of PSUs in each stratum.

Territorial authority

The ability to produce territorial authority (TA) estimates is one priority for the HLFS. However, there are 68 TAs in New Zealand. Explicitly stratifying by this many categories would be problematic due to the small number of PSUs in some TAs (ie if we forced selection of a PSU in each TA, this would result in a large variation of selection weights, increasing the sampling errors for national and other aggregate estimates). Therefore, we only implicitly stratify by this variable. That is, within the strata defined in Appendix table 1 (region by urban/rural status, by NILF status, by NZDep groups), PSUs are ordered by TA. We then select a systematic sample from the ordered list. This results in a high, although not certain, probability that at least one PSU from each TA will be selected. For the new sample we obtained at least one PSU from each TA.

Allocating PSUs to strata

Region

Any disproportionate sampling by stratum increases the sampling errors of national and other cross-strata estimates (eg age by sex breakdowns). Therefore, the sampling fraction in each region is proportional to that of the total population. That is, the number of PSUs selected from stratum h was:

$$n_h = n\frac{N_h}{\sum N_h}$$

where n is the total number of PSUs to select, and N_h is the number of PSUs in total in stratum h.

Table 1 shows the population proportion per region and the number of PSUs selected per region. (Note: The number of selected PSUs in each region is slightly different to that implied by the population proportions, due to allocation to other stratum layers, and also due to the overlap control methodology. See Overlap control.)

Table 1

Total number and number of PSUs selected per region

Region	Total number of PSUs	Percentage of total population	Number of PSUs selected	Percentage of sample
Northland	809	3.6	61	3.5
Auckland	6,549	29.2	512	29.0
Waikato	2,144	9.6	178	10.1
Bay of Plenty	1,400	6.2	101	5.7
Gisborne/Hawke's Bay	1,124	5.0	91	5.1
Taranaki	659	2.9	52	2.9
Manawatu-Wanganui	1,332	5.9	108	6.1
Wellington	2,634	11.7	204	11.5
Tasman/Marlborough/Nelson/West Coast	956	4.3	78	4.4
Canterbury	3,003	13.4	235	13.3
Otago	1,220	5.4	103	5.8
Southland	609	2.7	45	2.5
NATIONAL	22,439	100	1,768	100

Urban and rural PSUs

While any disproportionate sampling does affect sampling variances, it can also enable cost savings if we identify areas that are more expensive for fieldwork. For the HLFS, we found main urban areas to be cheaper for fieldwork than other areas, and so they are over-sampled in relation to their actual proportion in the population. That is, they are over-sampled such that:

$$\frac{n_{urban}}{N_{urban}} = 1.35\,\frac{n_{non-urban}}{N_{non-urban}}$$

where 'urban' refers to PSUs that are classified as being in main urban areas. Table 2 shows the total population of PSUs by main urban/other classification, and the resulting sample sizes in each stratum.

Table 2

Total number and number of PSUs selected
By urban/rural classification

Urban/rural classification	Total number of PSUs	Percentage of total population	Number of PSUs selected
Main urban area	16,289	72.6	1,365
Other	6,150	27.4	403

Not in the labour force strata

Proportionate allocation to the NILF strata is also employed (see Region for a description of proportionate allocation). The resulting sample sizes in each stratum are in table 3.

Table 3

Total number and number of PSUs selected
By NILF strata

NILF stratum	Total number of PSUs	Percentage of total population	Number of PSUs selected
Low NILF	15,428	68.8	1,216
High NILF	5,174	23.1	403
Not split	1,837	8.2	149

New Zealand Deprivation Index

We use Neyman allocation to the NZDep strata, based on unemployment. This method allocates the sample to strata based on the strata variances and similar sampling costs in the strata. A Neyman allocation scheme provides the most precision for estimating a population mean given a fixed total sample size. Neyman allocation assigns sample units (PSUs) within each stratum, proportional to the product of the population stratum size (N_h) and the within-stratum standard deviation (S_h), so we can achieve minimum variance for the population mean estimator. The sample size in any stratum is worked out according to the following equation:

$$n_h = n\frac{N_h S_h}{\sum N_h S_h}$$

where $n_h$is the sample size for stratum h, n is the total sample size, N_h is the population size for stratum h, and S_h is the standard deviation of stratum h of the variable for which you are trying to maximise survey precision (in this case unemployment).

Using unemployment results in an over-sampling of more highly deprived areas, which in turn leads to a higher proportion of the achieved sample being Māori (improved Māori estimates being a key priority for the HLFS), alongside improving the precision of the unemployed estimate. Simulations indicated that sampling errors of national labour market estimates were not negatively affected by this slight disproportionate allocation. Appendix tables 2–13 show the combined final allocation.

Selecting PSUs and targeting Māori

Māori statistics are of high importance for the HLFS. Targeting Māori in the sample design allows estimates of these statistics to be produced that are of sufficient quality to support policy analysis.

Within strata, we selected PSUs with probability proportional to size (PPS). (Previous samples used a simple random sample of PSUs within strata.) The size measure used was a combination of the size of the PSU (number of dwellings in the PSU), along with a targeting factor based on the density of Māori in the PSU:

$$s_{hk} = r_{hk}\left(\sqrt{0.02 + p_{hk}}\right)$$

where r_{hk} is the number of occupied or under-construction dwellings in PSU k of stratum h, and p_{hk} is the proportion of adults in the PSU that are Māori. This means that PSUs with a higher proportion of adults who are Māori will have a higher selection probability. By over-sampling Māori, we increase the estimated proportion of Māori in our sample (before non-response) from 11.8 percent to 14.2 percent.

The optimal targeting factor for over-sampling Māori is the square root of the proportion of adults in the PSU who are Māori. More heavily targeting Māori results in an increase in weight variation that more than offsets the increased number of Māori in the sample. The design effects for national estimates of NILF and employed are estimated to be around 1.04–1.06 with this level of targeting. The design effect for unemployed is estimated to be unaffected, since areas of high Māori density are more likely to have higher unemployment (and therefore a higher variation in the number unemployed).

Overlap control

At Stats NZ, we select PSUs for any household survey through the overlap control system. This is because it is desirable to minimise the sample overlap, as measured by the proportion of selected PSUs in common across surveys. That is, we try to reduce respondent burden as much as possible by minimising the number of respondents selected to take part in more than one survey.

In any overlap control scheme, stratum size and stratum sampling fractions may conspire to make sample overlap unavoidable. In such a case it is desirable to have a scheme that minimises the extent of overlap as part of its normal operation. The overlap control we use implements the conditional sampling methods of Chowdhury et al. (2000), and Bell (2011) – fusing the main ideas from these two methods, then extending the system to deal with frame reformation (changing the PSU boundaries), as in Lu (2012). This methodology allows us to share the household frame with other government departments and minimise overlap across organisations.

The overlap control system requires a list of PSUs that has each inclusion probability and stratum code identified, together with a TA code (see Territorial authority for why we include a TA code).

Sample rotation

Rotation groups

The HLFS is a rotating panel survey. This means that households are contacted a set number of times to take part in the survey before they are replaced by new households. A staggered design is used, so that a fixed proportion of the sampled dwellings are rotated out of the survey each quarter and replaced by a new set. The main reason for using a rotating panel design is that the precision of estimates of change over time are improved when there is overlap in the sample. That is, better estimates of quarter-on-quarter and quarter-on-same-quarter-a-year-ago change can be produced with this rotation pattern, since 70 percent or more of the sample is common in adjacent quarters. Another reason is that we can produce longitudinal datasets, which may be used for analysis of individuals' changing circumstances.

PSUs are sorted by strata, and then randomly within strata, before being assigned systematically to one of eight rotation groups.

Each quarter, we replace one-eighth of the sample, or one rotation group, with a new panel of dwellings from the same PSU. At the end of any quarter, one rotation group will have been in the survey for one quarter, another for two quarters, another for three quarters, and so on. The last rotation group will have been in the survey for its eighth and final quarter. Thus there is always an overlap of seven-eighths of the sample from one survey quarter to the next. This improves the quarterly estimates of change in the labour market.

We emphasise that it is the address (dwelling), not the particular people who live there, that is selected for eight quarters. Therefore, it is possible that while a certain address has been in the sample for more than one quarter, the people interviewed at that address may be in the sample for the first time. It is also possible for people to drop out of the sample before their eighth quarter if they move home.

Figure 1 shows the transition to the new sample. The first rotation group used in the new sample was rotation group 4 in quarter 117 (December 2014), with panel number 3 of the new sample being the first introduced. This is followed by panel numbers 5 and 1. By quarter 124 (September 2016) the sample in-field was made up entirely of the new sample of PSUs.

Figure 1

Rotation groups

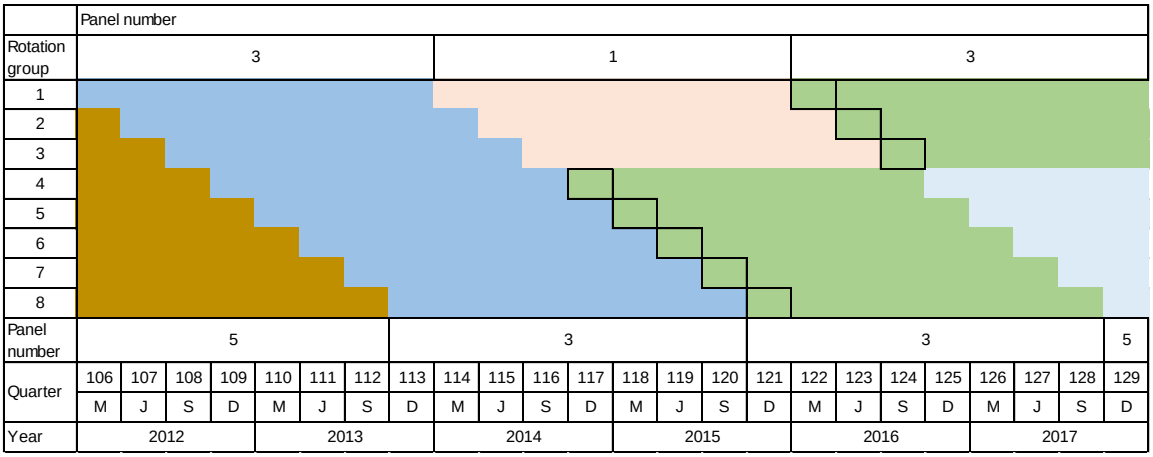

Panels

After selection, all dwellings within the selected PSUs are enumerated (ie interviewers visit PSUs to create an up-to-date listing of dwellings – see Enumeration for details). These dwellings are then systematically allocated to groups (panels), based on the PSU size and the total number of dwellings required. A PPS selection of PSUs would ideally involve taking panels of equal size across PSUs to achieve an equal selection probability design. (Note: Targeting Māori rules out the possibility of obtaining an equal probability design.)

For the selected PSUs to service both the HLFS and GSS until the 2018 Census, we require that each PSU be divided into no fewer than five full panels. This is at odds with the requirement that we draw a fixed number of dwellings from each PSU, since a good number of PSUs are too small. In practice, roughly 30 percent of selected PSUs will have exactly five panels, and less than the desired per-PSU take in each panel. In contrast, 70 percent of PSUs will yield more than the desired per-PSU take in each panel.

Note: The per-PSU take is described by the number of occupied or under-construction dwellings, so PSUs with large numbers of vacant dwellings will yield greater numbers of dwellings in total. Figure 2a shows the expected (before enumeration) number of dwellings per panel – the average is 10.31 (occupied or under-construction dwellings).

Dwellings within panels are geographically spread across a PSU. One panel is surveyed for eight quarters before it is replaced by another panel to be surveyed for the next eight quarters. When a

panel is replaced, it is never replaced with a neighbouring panel. This ensures that we are not replacing the current set of respondents with their next-door neighbours.

Figure 2a

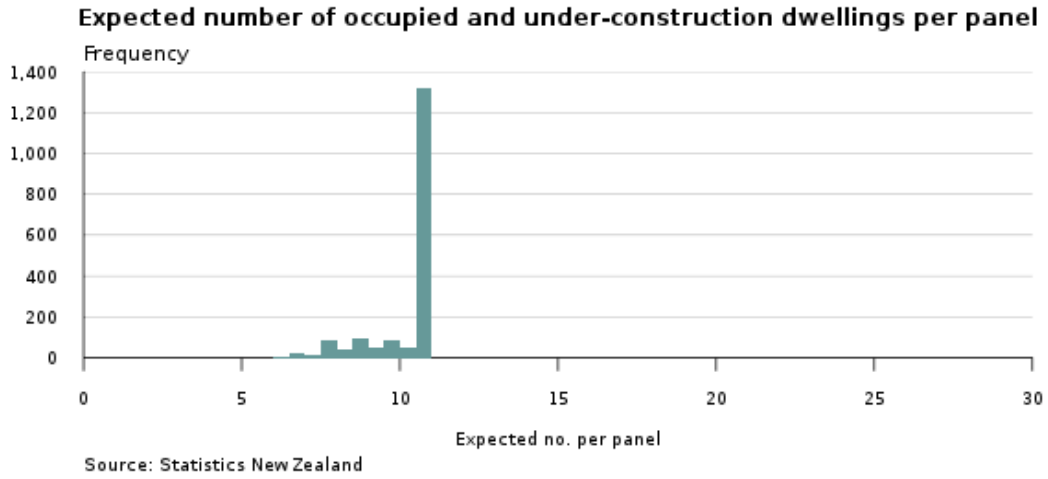

Figure 2b

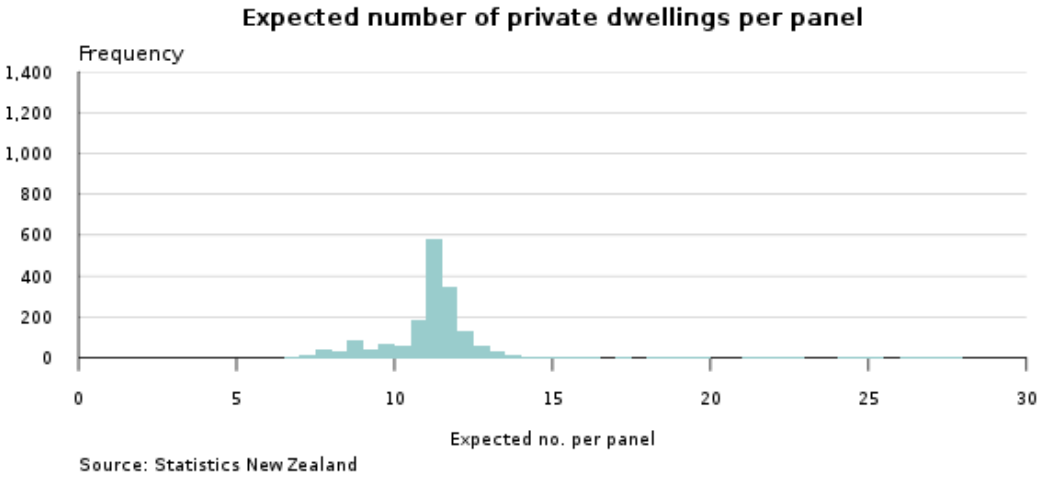

Sample reselection

Following a census, it is usual to reform the set of PSUs that comprise the sampling frame. Reformation is the process whereby we aggregate meshblocks to form PSUs. Reformation ensures PSUs are balanced in size as much as possible. PSUs can grow substantially over time, and reformation provides an opportunity to deal with those that have grown too large (or have become too small).

Following this reformation, it is usual practice to select a new sample of PSUs for the HLFS from the updated sampling frame. This avoids reusing panels, and provides an opportunity to improve the sample by using the reformed PSUs and the new sample design, which can both improve the precision of survey estimates.

Collection methodology

This chapter explains the collection methodology of the HLFS, including the questionnaires used, the different modes of collection, and the enumeration of PSUs.

Questionnaires

The HLFS has two separate questionnaires – a household questionnaire and a personal questionnaire. One member of the household is interviewed for the household questionnaire. This questionnaire determines the relationship information of everyone in the household, and who is eligible for a personal questionnaire. Flow charts for the current questionnaire can be found at Stats NZ Store House.

Stats NZ manages the questionnaire (ie the content), in consultation with the main external customers. The current questionnaire (with modified and additional content) was introduced in the June 2016 quarter.

Prior to the introduction of the current questionnaire, we appended supplementary survey topics to the HLFS personal questionnaire in some quarters. The New Zealand Income Survey was appended annually each June quarter. See Table 7: Supplements to the HLFS for details of other supplementary survey topics previously appended to the HLFS main interview. With the introduction of the current questionnaire, these supplementary survey topics are embedded within the questionnaire to allow for a more seamless experience for respondents, and to avoid asking the same questions more than once. Income content (to replace the New Zealand Income Survey) was included in the June 2016 quarter. See Household surveys programme 2016-20: Responding to New Zealand's information needs for proposed future additional content.

The HLFS allows interviewers to take responses from proxies if a respondent is unavailable or unable to answer the questions themself. A proxy can be any other eligible adult in the household. See Proxies for the frequency of proxy responses. With the change to the way supplementary content is included in the questionnaire, the HLFS proxy rules will now also apply to any supplementary content. In the past, supplementary content used separate proxy rules.

Households with only people aged 75 and over are only interviewed in their first quarter of participation in the survey (and in June quarters – see Estimation and imputation for more details).

Collection modes

Interviews are conducted face-to-face using computer-assisted personal interviewing (CAPI) for most first interviews. Most respondents are then interviewed for their second and subsequent interviews over the telephone, using computer-assisted telephone interviewing (CATI), (although respondents can request CAPI mode if they prefer).

Telephone interviews are conducted by interviewers from our centralised telephone unit in Auckland, while face-to-face interviewers visit respondents at home. Interviewers (both telephone and face-to-face) are employed by Stats NZ and also work on our other household surveys, and collect for the consumers price index.

The questionnaires we use for face-to-face and telephone interviewing are almost identical. They are produced using the BLAISE CAI software package, which was created by Statistics Netherlands.

Prior to the introduction of the current questionnaire (in the June 2016 quarter), respondents also had the option of filling in a paper self-completion questionnaire. Additionally, as part of the non-response follow-up process, we sent a paper self-completion questionnaire automatically if we received no response within the week the case is assigned for fieldwork. Since the June 2016 quarter we no longer produce a paper self-completion questionnaire. In the future we will explore introducing an internet self-completion questionnaire.

Interviews are conducted over a quarter in one-week periods starting on a Sunday. Most information obtained relates to the week before the interview ('the reference week'). Before the interview, we send a pre-notification letter (PNL) to every address in the selected sample explaining their address has been selected for the HLFS and an interviewer will be calling. Within the PNL, we assure respondents that their information will be kept secure and confidential; we also give them answers to a list of frequently asked questions.

See Frequently asked questions about our surveys for details.

Enumeration

Traditionally, when a new sample of PSUs is selected, interviewers visit those PSUs before we select the sample of dwellings within those PSUs, to create an up-to-date listing of dwellings. We call this process pre-enumeration.

Once this has occurred, each PSU is periodically 'check-enumerated' by an interviewer – to ensure the dwellings list is up-to-date. In the HLFS, this occurs once every two years in each PSU as a new rotation group (selection of dwellings) is selected into the sample.

Pre-enumeration

We identify the PSUs where the quality of New Zealand Post's National Postal Address Database (NPAD) listing is expected to be high enough that no pre-enumeration is required.

The general methodology is to:

- look at how well NPAD compares with census counts and enumeration counts from the HLFS stored on the Household Survey Frame
- within the HLFS PSUs, match addresses between the HLFS enumeration and NPAD to check the addresses are actually the same addresses.

Over time, the comparison between NPAD and census counts loses relevance. We identify areas (using building consents) where growth has occurred, and look to see if this growth is also captured by New Zealand Post. If not, the PSU should be pre-enumerated. This mirrors the strategy being developed for ongoing check enumeration.

Additionally, pre-enumeration only occurs in PSUs that are within 20kms of the closest interviewer. For PSUs outside this range, we use in-office enumeration to clean NPAD address listings – these listings are used for the initial selection of dwellings. This listing is supplemented in the first quarter of interviewing by a check-enumeration.

Check-enumeration

For the new sample, instead of check-enumerating all PSUs, a subset of PSUs is identified – to check-enumerate based on changes to NPAD and/or building consent data. The NPAD data is also used to provide an initial listing of addresses (rather than interviewers starting from scratch). Geospatial methods are used to automate the production of map images, and to order the NPAD listing in a way that is sensible for interviewers to traverse the PSUs when completing enumeration.

Coding and processing the data

This chapter explains the editing and coding used in the HLFS, and provides a list of variables found in the dataset.

Edit checks

Editing aims to ensure that collected data meets certain quality requirements. This includes:

- providing internal consistency
- improving the validity of the findings
- reducing non-sampling errors.

With the introduction of computer assisted interviewing (CAI) over 2005 and 2006, we perform edit checks for the HLFS as the data is entered. Edit checks, including range and consistency checks, are programmed into the electronic questionnaire and are triggered according to set rules as the interviewer enters responses.

Range checks involve checking that answers to questions are within a specified range. If the respondent's answer is outside this range, the interviewer can tell immediately that there is a mistake. For example, in 2014 the date-of-birth question may have a specified range of 1904–2014. This assumes no one in New Zealand is over 110 years old.

Consistency checks involve checking the relationship between a set of answers. For example, if a person's age is 14 years, based on the date of birth supplied, but they are identified in a parent role to a 35-year-old household member in the relationship table, this would trigger an edit.

Questionnaire edits can be classified as 'hard edits' or 'soft edits'. We use hard edits only where it can be identified that the response is definitely an error. For example, a person's date of birth cannot be in the future. An interviewer cannot move past a hard edit. Soft edits indicate to the interviewer that a response may be an error, but these can be suppressed and moved past by the interviewer.

Post-collection editing of income content is also carried out in June quarters. Outliers are examined for income from wage and salary, self-employment, superannuation, WINZ, student allowance, and IRD, and hourly earnings for wages and salaries. If the data entered is clearly wrong (using other information collected in the survey) then the value is edited with a corrected value.

Coding

The HLFS collects information about a person's current or previous employer, and current or previous job characteristics. We use this information to create statistics about industry and occupation by coding the person's response to standard classifications. Country of birth is coded to the 1999 New Zealand Standard Classification at interview time. In addition, we perform some coding of ethnicity data in the HLFS (where a person indicates an 'other' ethnic group and gives further detail), and a small amount of qualification data is coded. From the June 2016 quarter onwards, the following variables are also coded where there are any text responses provided in the 'Other – please state' category of the questionnaire:

- methods used to find work
- main activity of those not in the labour force
- main reason respondent was away from work in reference week
- main reason respondent worked fewer hours in main job
- main reason respondent worked fewer hours in second job
- main reason respondent worked fewer hours in other job(s)
- main reason respondent was not looking for work
- main reason respondent does not want a job
- main reason respondent was not available for work in the reference week
- main reason respondent left their last job
- main reason respondent who wants more hours is not working more hours
- methods used to get more hours of work.

From the September 2016 quarter onwards, the HLFS also makes use of rolling-over coded data from previous quarters. This means that more use is made of data from a previous quarter for the current quarter's interview. For example, we are able to confirm that a person's job or qualifications have not changed since the last time we spoke to them. For coding this means that we can use or 'rollover' the coded value from the previous quarter's interview rather than send the information through to be manually coded again.

Industry

Information about a person's employer (eg name and address) is collected in the HLFS questionnaire along with a description of the main activity of the place in which the person works. We collect this for people who are currently employed, and for the previous employment spell for those not currently employed but who had worked in the previous five years (in September quarters only). Where a person has more than one job, the information is collected for their main job and second job.

Once the data is collected in the questionnaire, we code it.

Since the September 2009 quarter, we've coded industry statistics in the HLFS to the Australian and New Zealand Standard Industrial Classification 2006 (ANZSIC06). Previously, HLFS used the 1996 version. Outputs of industry data are further classified using the New Zealand Standard Industrial Output Categories (NZSIOC) classification.

Occupation

We collect information about a person's tasks and duties in their main job and second job (where applicable), and their description of their occupation in the HLFS questionnaire – for those who are currently employed and, where someone is not currently employed but had worked in the previous five years, for their previous employment spell (in September quarters only).

Since the September 2009 quarter, we've coded occupation statistics to the Australian and New Zealand Standard Classification of Occupations 2006 (ANZSCO). Before this, we classified occupation data using the New Zealand Standard Classification of Occupations 1999.

Ethnicity

In the HLFS we collect ethnicity data about all household members. Respondents are able to provide up to 14 ethnic groups. Where 'other' ethnic groups are given and further text description is supplied, we code the data using the Ethnicity New Zealand Standard Classification 2005.

Variables available in the dataset

The following is a list and brief description of variables available in the HLFS datasets. Flow charts for the current questionnaire can be found at Stats NZ Store House.

Table 4

Variables available in the dataset

Variable	Description
Household_code	Household identifier number
Person_code	Household_code followed by person number – unique number for each person within a household
Quarter	Survey quarter that the dataset relates to (eg 112 = September 2013 quarter)
HQStatusCode	Response status code for the household questionnaire
Week	Week of the quarter that the case is allocated to (1–13)
DVEStatus	Eligibility status – only those with a DVEStatus of 'Eligible responding' are assigned a weight
DVSex	Sex of person – collected for all household members
Sex_Imputed	Indicator to show whether DVSex has been imputed or not
Sex_Donor_ID	Person_code of the donor used for imputing DVSex
DVDOB	Date of birth
DVAge	Age of respondent held in single year format (calculated from date of birth, or age in years if date of birth is not provided)
DVAgeGp1	5-year age band (maximum category = 80+)
DVAgeGp2	Broad age bands (under 15, 15-29, 30+)
Age_Imputed	Indicates whether age has been imputed or not
Age_Donor_ID	Person_code of the donor used for imputing age
DVCOB	Country of birth
DVYrsInNZ	Number of years person has lived in New Zealand if born overseas
EthEuropean	Indicator of European ethnicity
EthMaori	Indicator of Māori ethnicity
EthPacific	Indicator of Pacific ethnicity
EthAsian	Indicator of Asian ethnicity
EthMELAA	Indicator of Middle Eastern/Latin American/African ethnicity

Variable	Description
EthOther	Indicator of 'Other' ethnicity
DVEthnic	Combines the above six high level classifications above into a character string without any duplication
DVEth1-DVEth14	Ethnic group(s) individual identified with stored at level four of the Ethnicity New Zealand Standard Classification 2005
Ethnic_Imputed	Indicator to show that ethnicity has been imputed
Ethnic_Donor_ID	Person_code of the donor used for imputing ethnicity
DVDescent	Indicator of Māori descent
DVFamCode	Family group code identifies parents and children of different family nuclei, non-family members, and single-person households
NucleusNumber	Shows what family nucleus household members belong to (ie people with the same nucleus number belong to one family nucleus) – a household may contain more than one family nucleus
DVChildEmpStat	Indicates whether children aged 15–17 are employed full-time or not
DVFam_FamNucYN	Indicates whether the person is in a family nucleus or not
DVFam_NumInFamNuc	The number of people in a family nucleus
DVFam_WithPartner	Indicates whether the person is with a partner or not
DVFam_NumPartner	Number of people in a partner role within a family nucleus
DVFam_ParentRole	Indicates whether the person is in a parent role or not
DVFam_ChildRole	Indicates whether the person is in a child role or not
DVFam_NumChild	Number of children in the family nucleus
DVFam_DepStat	Dependency status of a child
DVFam_NumDepChild	Number of dependent children in a family nucleus
DVFam_NumIndepChild	Number of independent children in a family nucleus
DVFam_NumChildDepNK	Number of children of unknown dependency in a family nucleus
DVHHSize_code	Total number of people in the household
DVHHType	Household composition
DVHHTen	Indicator of whether the dwelling is owned or held in family trust by any member of the household
DVHHOver75	Identifies households where all members are aged 75 years or over
DVRegCouncil	Regional council area
DVRegion	Regional council area (combining Gisborne and Hawke's Bay, and Nelson/Tasman/Marlborough/West Coast Regional councils) – these Regional Council areas are used for weighting and final outputs
DVMeshblock	Meshblock 2015 classification
DVTA	Territorial Authority 2015
DVUrban	Urban/rural area status 2015
HQRespondent	Indicator showing if the person was the HQ respondent
PersonSource	Shows the 'source' of the person this quarter in relation to previous quarter (eg from dependant data, a new HQ respondent that has joined the household, or a new household member)
ShortTermVisitor	Indicates whether the person was identified as a 'short-term visitor' in New Zealand or not
AnyoneTempInNZ	Indicates whether someone in the household is not permanently living in New Zealand
HHRotatingIn	Indicator showing whether the household is rotating in this quarter
HHRotatingOut	Indicator showing whether the household is rotating out this quarter
RotnGrp	The rotation group that the household belongs to
RotnNum	The rotation number of the household (ie first quarter, second quarter etc)

Variable	Description
IncomeQtr	Indicates whether it is an income quarter or not
SeptQtr	Indicates whether it is a September quarter or not
SupplementInd	Indicates whether the current quarter is a supplement quarter
Mode	The mode that the case was completed in (ie CATI or CAPI)
HQInterviewDate	Date the HQ interview took place
RefWeekStartDate	Start date of the reference week
RefWeekEndDate	End date of the reference week
StillHere	Indicates whether a person recorded in a previous quarter is still living in the same household
PQStatusCode	Response status code for the PQ
ProxyYN	Indicator to show if a proxy was used
ProxyWho	Person number of individual in the household who provided the proxy response
DVLFS	Labour force status of individuals aged 15 or over – see Output from the HLFS for the definition of labour force status
DVUnderUtilise	Underutilisation status of individual – includes the unemployed labour force, the underemployed labour force, and the potential labour force
OwnBus	Indicates whether the respondent did any work in their own business last week
PaidJob	Indicates whether the respondent did any work in a paid job last week
FamBus	Indicates whether the respondent did any unpaid work in a family business last week
AwayAllJobs	Indicates whether the respondent was away from all types of work last week because of sickness, holidays, or some other reason
AwayOwnBus	Indicates whether the respondent was away from their own business last week
AwayPaidJob	Indicates whether the respondent was away from a paid job last week
AwayFamBus	Indicates whether the respondent was away from an unpaid job in a family business last week
WhyAway	Main reason why an employed person was away from all work last week
AwayOthJobs	Indicates whether the respondent has any other work that they were away from last week
AwayOthOwnBus	Indicates whether the other work the respondent was away from last week was their own business
AwayOthPaidJob	Indicates whether the other work the respondent was away from last week was a paid job
AwayOthFamBus	Indicates whether the other work the respondent was away from last week was an unpaid job in a family business
AwayOthNone	Indicates whether the other work the respondent was away from last week was not a business, paid job, or family business
HasOwnBus	Indicates whether the person has their own business
HasPaidJob	Indicates whether the person has a paid job
HasFamBusJob	Indicates whether the person has an unpaid job in a family business
WorkType	The type of work the respondent was doing or was away from in the reference week
DVMJH	Identifies people who had more than one job last week
OwnBusNum	Number of own businesses the respondent had last week
OwnBusNumNR	Indicator of a 'Don't know' or 'Refusal' response to the question asking number of businesses the respondent had last week
PaidJobNum	Number of paid jobs the respondent had last week

Variable	Description
PaidJobNumNR	Indicator of a 'Don't know' or 'Refusal' response to the question asking number of paid jobs the respondent had last week
FamBusNum	Number of unpaid jobs in family businesses the respondent had last week
FamBusNumNR	Indicator of a 'Don't know' or 'Refusal' response to the question asking number of unpaid jobs in family businesses the respondent had last week
NumJobs	Total number of jobs the respondent had last week
DVEmpStatMain	Employment status in main job
DVOccMain	Occupation in which people are employed in their main job
DVIndMain	Industry in which people are employed in their main job
DVEmpStatSec	Employment status in second job
DVOccSec	Occupation in which a holder of multiple jobs is employed in their second job
DVIndSec	Industry in which a holder of multiple jobs is employed in their second job
DVFTPT	Full-time/part-time status for those employed – full-time employment is classified as usually working for 30 hours or more, part-time employment is usually working fewer than 30 hours per week
DVUHrMain	Hours usually worked each week in main job
DVUHrSec	Hours usually worked each week in second job
DVUHrOth	Hours usually worked each week in other job(s)
DVUHrTot	Hours usually worked each week in all jobs
UHr_Imputed	Indicator to show that usual hours have been imputed
UHr_donor	Person_code of the donor used for imputing usual hours
SameHrsMain	Whether respondent worked the same hours, more hours, or fewer hours than usual in main job last week
SameHrsSec	Whether respondent worked the same hours, more hours, or fewer hours than usual in second job last week
SameHrsOth	Whether respondent worked the same hours, more hours, or fewer hours than usual in other job(s) last week
DVAHrMain	Hours actually worked last week in main job
DVAHrSec	Hours actually worked last week in second job
DVAHrOth	Hours actually worked last week in other job
DVAHrTot	Hours actually worked last week in all jobs
AHr_Imputed	Indicator to show that actual hours have been imputed
Ahr_donor	Person_code of the donor used for imputing actual hours
ReasFewHrsMain	Main reason an employed person worked fewer hours than usual hours in their main job last week
ReasFewHrsSec	Main reason an employed person worked fewer hours than usual hours in their second job last week
ReasFewHrsOth	Main reason an employed person worked fewer hours than usual hours in their other job(s) last week
DVJobTenW	Number of weeks an employed person has been working for their current employer or in their current business or have been self-employed in their main job
DVJobTenC	Length of time (in categories) employed people have been working for their current employer or in their current business or have been self-employed in their main job
DVDaysWrkMain	Days of the week the respondent worked in their main job

Variable	Description
DVDaysWrkSec	Days of the week the respondent worked in their second job
DVDaysWrkOth	Days of the week the respondent worked in their other job(s)
Union	Whether an employee belongs to a union in their main job
EmpAgreeYN	Whether the respondent had a written employment agreement in their main job
EmpAgreeType	Type of employment agreement the respondent had in their main job
DVEmpRel	Type of employment relationship employees have in their main job
PermJob	Whether the respondent's main job is a permanent job
FixedTerm	Whether the respondent is on a fixed-term contract in their main job
Project	Whether the respondent is working until a task or project is finished in their main job
TempReplace	Whether the respondent was hired to temporarily replace another worker in their main job
Agency	Whether the respondent is paid by or through an employment agency in their main job
Casual	Whether the respondent's main job is a casual job
Seasonal	Whether the respondent's main job is a seasonal job
PrefPerm	Whether a respondent holding a temporary, non-seasonal main job would prefer a permanent job
PrefSelfEmp	Whether a respondent who is self-employed in their main job would prefer to continue being self-employed or to have a paid job working for someone else
PrefOngoing	Whether a respondent who has a seasonal main job would prefer to continue in seasonal work or to get a job that is ongoing
PrefOwnBus	Whether a respondent who has their own business as their main job would prefer to continue in their own business or in a paid job working for someone else
DVUnderEmp	Identifies people who are in part-time employment, would prefer to work more hours, and would be able to do so within four weeks of the interview date
WantMoreHrs	Whether the respondent would like to work more hours than they usually do
HrsWanted	Number of hours a week those who would like more hours would like to work in total
WhyNotMoreHrs	Main reason an employed person worked fewer hours than they would like
DVMoreHrsLook	Whether the respondent looked at job advertisements in the last four weeks to try and get more hours of work
DVMoreHrsEmp	Whether the respondent asked their employer for extra hours in the last four weeks to try and get more hours of work
DVMoreHrsOthEmp	Whether the respondent asked another employer in the last four weeks to try and get more hours of work
DVMoreHrsEmpAg	Whether the respondent contacted an employment agency in the last four weeks to try and get more hours of work
DVMoreHrsRel	Whether the respondent contacted friends or relatives about a job in the last four weeks to try and get more hours of work
DVMoreHrsOther	Whether the respondent did something else in the last four weeks to try and get more hours of work
MoreHrsAvail	Whether the respondent could start working more hours in the next four weeks if more hours were available

Variable	Description
DVNEET	Ten categories of labour force and education status, with three categories indicating those respondents who are not in employment, education, or training
JobToStart	Whether the respondent has a job to start in the next four weeks
LookForWork	Whether the respondent has been looking for paid work in the last four weeks
DVLook	Whether the respondent looked at job advertisements in the last four weeks to find work
DVContactEmp	Whether the respondent contacted an employer in the last four weeks to find work
DVContactEmpAg	Whether the respondent contacted an employment agency in the last four weeks to find work
DVContactWINZ	Whether the respondent contacted Work and Income in the last four weeks to find work
DVContactRel	Whether the respondent contacted friends or relatives in the last four weeks to find work
DVSetUpBus	Whether the respondent had taken steps to set up a business in the last four weeks
DVLookOther	Whether the respondent had done something else to find work in the last four weeks
LookFTPT	Whether the respondent had been looking for full-time or part-time work
LookFTPT_Imputed	Indicator to show whether looking for full-time or part-time work has been imputed or not
LookFTPT_donor	Person_code of the donor used for imputing looking for full-time or part-time work
AvailLastWk	Whether the respondent could have started work last week, if a job had been available
WhyNotAvail	The main reason the respondent was not available for work last week
Avail4Wks	Whether the respondent could start work in the next four weeks, if a job was available
HowLongLook	How long the respondent has been looking / looked for work
WksLookNum	The number of weeks the respondent has been looking / looked for work
MthsLookNum	The number of months the respondent has been looking / looked for work
YrsLookNum	The number of years the respondent has been looking / looked for work
DVDurUnemp	The length of time a respondent without employment has been looking for work
WantJob	Whether a respondent who is not currently looking for work, would like to have a paid job
WhyNotJob	Main reason the respondent does not want a job
WhyNotLook	Main reason the respondent did not look for a job
Retired	Whether the respondent was retired last week
IntendLook	Whether the respondent intended to look for work within the next two years
WhenLook	The time within which the respondent intends to look for work
DVTimeSinceWork	The length of time since the respondent last worked in a job or business
DVOccPrev	The occupation of the respondent in their last job
DVIndPrev	The industry of the last job held by the respondent

Variable	Description
DVPrevEmpStat	The employment status of the last job held by people of working age who were not employed in the reference week, but have been employed within the last five years
WhyLeft	The main reason the respondent left their last job
DVMainAct	The main activity of the respondent in the reference week, for those respondents who were not in the labour force
SchoolYN	Whether the respondent was still at school
SchQualYN	Whether the respondent had a school qualification
DVSchQual	The highest secondary school qualification obtained by the respondent
DVPSQual1-DVPSQual20	All the formal post-school qualifications obtained by the respondent
DVQualYrs1-DVQualYrs20	The number of years since each post-school qualification was gained
DVHQual	Highest qualification of the respondent
DVStudy	The study status of the respondent
StudyYN	Whether the respondent was doing any study last week
DVInHHOver75	Whether the respondent is in a household where all in-scope members are aged 75 years or older
SelWgt	Selection weight of the household (used in the calculation of the final weight)
RespWgt	Non-response weight (not used in the calculation of the final weight)
Finalwgt	Final weight assigned to individual
Finalwgt_1-Finalwgt_100	Replicate weights (used to calculate sampling errors)

Estimation and imputation

This chapter outlines the estimation for the HLFS, including assigning eligibility to households, weighting households for selection, non-response, and undercoverage, imputation of missing responses, time series estimation, and data suppression and rounding procedures we use in the information releases.

Assigning eligibility

We assign each dwelling selected for the HLFS one of five eligibility statuses (see table 14 in Appendix 2) derived from the household and personal questionnaire data. A dwelling is classified as eligible if it has at least one eligible member. Eligibility status (DVEStatus) is derived from labour force status, the household response status code, and the personal response status code (see tables 15–17 in Appendix 2).

Diagrammatically this is represented by figure 3.

Figure 3

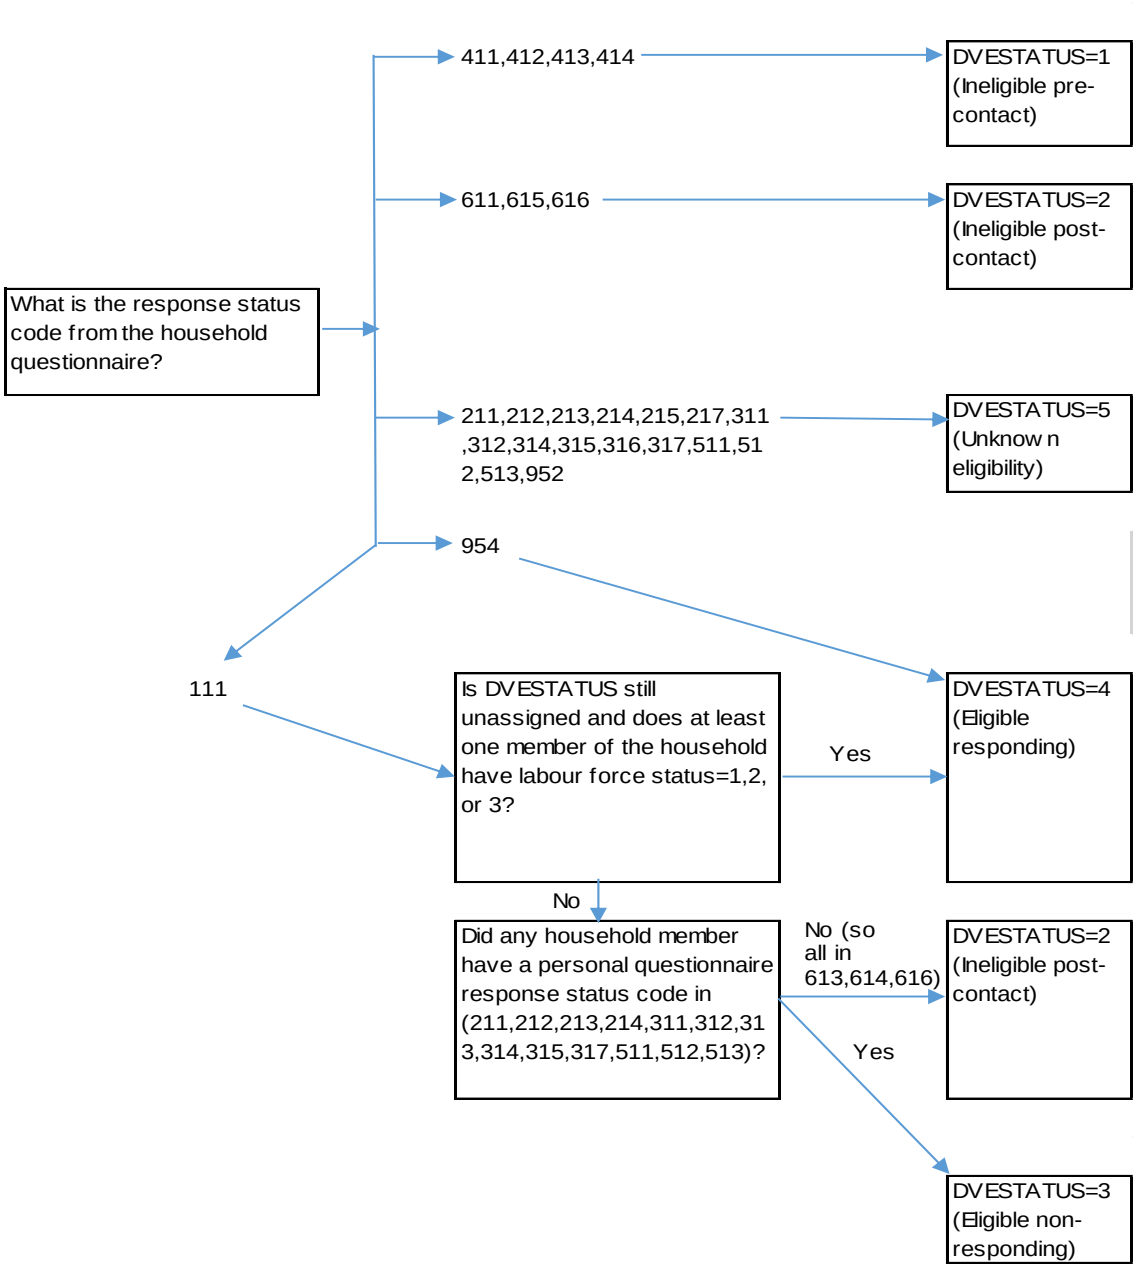

Weighting

Sample selection weights

The HLFS collects information on a sample of the population. To enable us to infer from this sample to the entire population we must weight the sample data. This entails assigning each responding a weight, which can be thought of as the number of people in the population that each individual represents.

PSU selection weight

The first stage of the weighting is the selection weight (also called a design weight). The overall selection weight for a household is made up of the PSU selection weight and the household selection weight. We calculate the selection weight for each PSU as the inverse of the probability

of selection, so PSUs with a lower probability of selection receive a higher selection weight. Within strata, PSUs are selected with probability proportional to size. This means that larger PSUs have a higher probability of being selected. See Appendix 3 for more details.

One consequence of introducing a new sample is that during the transition there are effectively two independent samples in operation. To account for this, we employ a transition factor, which we multiply the PSU weights by. See Appendix 3 for details.

Household selection weight

We next multiply the PSU selection weight by a household selection weight to give the overall selection weight. The household selection weight accounts for the sampling of households within PSUs – we calculate it as the inverse of the selection probability, where the selection probability is the number of selected addresses in the PSU divided by the total number of addresses in the PSU. See Appendix 3 for details on the number of selected addresses in each PSU.

Non-response adjustment

We can modify the design weight by non-response factors to account for some subgroups being less likely to respond than others. There are various strategies for doing this, including adjustment cell reweighting, propensity score reweighting, and calibration.

Before we introduced a regional benchmark (see Calibration), the HLFS employed a cell adjustment technique, where the cells were based on region and month. That is, within each region-by-month group (or cell), we multiplied the design weight by a rate-up factor to adjust for people who do not respond to the HLFS. The rate-up factor was based on the inverse of the weighted response rate within that group. That is, within each cell, we inflated the weights of respondents to account for the non-respondents.

However, with the introduction of the regional benchmark, this step became redundant. Therefore, we no longer include a non-response adjustment stage within the weighting schema. However, calibration also adjusts for non-response (along with undercoverage), and is included in the weighting for the HLFS.

Calibration

The final stage of weighting for the HLFS is the calibration to benchmarks (auxiliary information), which are the expected counts of people in the total target population. This adjusts for undercoverage of the target population and undercounting of some groups in the population due to differential response rates. We set the calibration weights to sum to a set of benchmarks.

We use a form of calibration known as 'integrated weighting'. This ensures that all individuals in the same household have the same weight, and that household statistics derived from person-level data match the same statistic calculated directly from household-level data. See Appendix 3 for more details on the calibration methodology.

The benchmarks we use for the HLFS are five-year age groups by sex, the number of Māori adults by two age groups (age 15–29, 30+), and 12 regions. These benchmarks refer to the target population (the civilian, usually resident, non-institutionalised population aged 15 years and over). Ratios are first calculated as the proportion of the estimated resident population on census night (census usually resident population count plus net census undercount plus residents

temporarily overseas) that is within the target population. We then apply these ratios to the estimated resident population at a given date to produce the population benchmarks.

After each census, we revise the population estimates, and the historical HLFS data is reweighted to these updated population totals. Reweighting (or rebasing) exercises give the opportunity to use the most up-to-date population estimates within HLFS weighting. We recalibrate historical data, which was initially weighted to population estimates, to the newly available revised population estimates.

Imputation

Imputation is the process where we impute values for people who have missing values in their survey questionnaire.

We impute for people who belong to eligible responding dwellings and have missing values for their sex, age, ethnicity, looking for full-time or part-time work (if unemployed or not in the labour force), and usual and actual hours worked in all jobs. In June quarters, we also impute income from jobs (main and second, where income from jobs includes wages and salaries, self-employment, and business's income), and income from government transfers (income received from Inland Revenue, Work and Income, ACC, and Student Allowance) and superannuation.

All variables are imputed using nearest-neighbour donor imputation. The software used is the Canadian Census Editing and Imputation System (CANCEIS). A donor is selected by finding a respondent with similar values to the recipient on some set of matching variables, where these matching variables are correlated to the missing values. An investigation was carried out to determine appropriate matching variables for age and ethnicity imputation and for hours imputation. In June quarters, the same donor is used for hours and income, and if income is imputed where hours were reported (and vice versa), the hours variables in the income module are overwritten using the donor's values (or income is overwritten where income was reported and hours were not), to maintain consistency for hourly earnings. Additional matching variables are used for hours and income imputation in June quarters.

The most important reason for imputation is to prevent the loss of data where possible. Records with missing values should not be ignored. Consistent bias could come from certain types of respondents who are difficult to collect data from. Assigning values at the micro-level allows us to conduct analysis as if the dataset were complete. Imputation aims to produce distributions that reflect as near as possible the distributions of the total population, but it does have the potential to distort the data or introduce additional biases.

We use another form of imputation for people aged 75 and over (75+). If a household has only people aged 75+ when interviewed in its first quarter of participation, then we do not interview respondents in subsequent quarters. Instead, their current quarter responses are imputed by carrying over the data from their most recent interview. The exception to this is quarters where we also collect income data (June quarters) or other additional information that requires the re-interviewing of 75+ households. In these quarters, 75+ households are again interviewed, and we use this data for imputation in subsequent quarters.

Doing this introduces cost savings and reduces respondent burden while estimates remain largely unchanged – the labour force status of people aged 75+ tends to be relatively stable. Such households make up approximately 9 percent of the first-time-in rotation group.

Time-series estimates

Many time series have a recurring seasonal pattern that obscures the underlying behaviour of the series. Seasonal adjustment is the process of estimating and removing the varying seasonal effects from a time series in order to reveal non-seasonal features. This ensures that the underlying movements in the time series are more visible.

We seasonally adjust these series in the HLFS:

- people employed, unemployed, and not in the labour force, by sex
- employment rate, labour force participation rate, and unemployment rate, by sex
- levels and quarterly and annual changes of people employed, unemployed, and not in the labour force, full-time and part-time work, and total actual and usual hours
- labour force and education status of those aged 15–24 years (including not in employment, education, or training).

Seasonal adjustment is produced using the X-13ARIMA-SEATS Version 1.1 package developed by the U.S. Census Bureau. Each quarter, the seasonal adjustment process is applied to the latest and all previous quarters. This means that seasonally adjusted estimates for any of the previously published quarters are revised and may change slightly.

In some cases, two or more time series make up a composite time series. For example, in the HLFS, male employed and female employed make up the total employed series (ie male employed + female employed = total employed). We can also derive total employed using the part-time and full-time series (ie part-time + full-time = total employed). The HLFS adjusts each series separately. That is, the components of total employed (male and female employed) are adjusted separately and then summed to produce the seasonally adjusted total employed. The seasonal patterns of the component series will influence the values of the composite series. For example:

- there is an expectation that seasonally adjusted male and seasonally adjusted female will sum to seasonally adjusted total employed
- there is no expectation that seasonally adjusted full-time and seasonally adjusted part-time will sum to seasonally total employed

The standard seasonal model assumes that the actual (observed) series is composed of three different components:

- the trend cycle
- the seasonal component (this includes calendar effects)
- the irregular component.

For the HLFS, we assume the relationship between these components is multiplicative. This means the seasonal effect varies in size with the level of the series. The seasonally adjusted series refers to the actual series with the seasonal component removed.

The seasonal component is the seasonal pattern found in many subannual economic series. It is reasonably stable in terms of annual timing, direction, and magnitude. It can be caused by, for example, natural factors (eg seasonal weather patterns), administrative measures (eg start and end dates of the school year), and social/cultural/religious traditions (eg fixed holidays such as Christmas).

The trend cycle reveals the smooth, relatively slowly changing features in a time series. It merges any cyclical movements present with the long-term underlying trend. The underlying trend is often associated with some basic characteristic of the economy, such as population growth, and reveals the underlying direction of movement in a series. Cyclical movements are composed of cumulative, reversible, short-, or medium-run movements. They are characterised by alternating periods of expansion and contraction as they reflect general economic activity.

The irregular component is the part of the observed value that is not included in the trend cycle or the seasonal effects (or in estimated trading day or holiday effects). Its values are unpredictable for timing, impact, and duration, and can arise from influences such as sampling error, non-sampling error, unseasonable weather, natural disasters, or strikes. Much of the testing done in seasonal adjustment is seeing if there is any structure left in the irregular component. Many of the quality diagnostics are based on comparing the variability of the other identified components with that of the irregular component.

We monitor our data to make sure our seasonal adjustment is robust. The X-13ARIMA-SEATS programme is highly customisable and can produce a wide variety of possible adjustments for any particular input series. Consequently, the programme produces diagnostics that are useful in assessing the quality of the chosen adjustment. We publish these diagnostics quarterly within the labour market statistics' releases.

See Seasonal adjustment in Statistics New Zealand for more information on methodology.

Adjusting for the Survey of Working Life

In the March 2008 and December 2012 quarters, we found evidence of a response bias which we suspected was due to running a supplement to the HLFS in those quarters (the Survey of Working Life (SoWL)). To adjust for the effect of the SoWL on responses, we applied a prior adjustment, which is an explicit form of treating an outlier. Using a regression model, we estimated a replacement value for the March 2008 and December 2012 quarters. We prior-adjusted these time series to account for the SoWL:

- female employed
- male employed
- total actual hours
- total usual hours
- female not in the labour force
- male not in the labour force.

Adjusting for the Easter effect in the total actual hours series

There is an expectation that fewer actual (worked) hours are reported for the quarter that Easter falls in, due to public holidays and a tendency for employed individuals to take leave over the period. If Easter consistently falls in one quarter, this effect would be accounted for by the seasonal component. However, Easter dates change; it can fall in the March quarter or the June quarter. We found evidence for an Easter effect on total actual hours – the number of hours reported for a given quarter depended on whether Easter fell within that quarter. To account for the changing date of the holiday, we applied an Easter effect to this time series, using the Genhol utility (an executable that can be used alongside the X-13ARIMA SEATS programme).

Data suppression and rounding procedures

In the information release each quarter, we suppress cells with estimates of less than 1,000. They appear as 'S' in the tables. These estimates are subject to sampling errors too great for most practical purposes.

Figures in the release are also rounded – to the nearest hundred or nearest thousand for seasonally adjusted and trend estimates. This may result in a total disagreeing slightly with the sum of the individual items shown in any table. Where figures are rounded, the unit is shown as (000) for thousands. We calculate quarterly and annual changes for figures on unrounded numbers. However, quarterly and annual percentage-point changes for rates are done on rounded rates.

Reliability of the estimates

In this chapter we describe sampling and non-sampling errors that affect the reliability of the estimates, and present some estimates that are indicators of aspects of non-sampling errors.

Sampling errors

Sampling error occurs because the sample survey collects information from only a fraction of the population, rather than the entire population. The extent of the error depends on many factors, including:

- sample size – increasing the sample size (all else being equal) reduces sample error; however, there is a point beyond which we gain little by further increasing the sample size
- variability of the characteristics of interest – the greater the variation in the population the greater the sampling error
- sample design – designs that use information known about the population (and that relate to the characteristics of interest) reduce the sample error.

We calculate sampling errors using the replication-based delete-a-group jackknife method. Replication methods for variance estimation draw multiple replicates (or subsamples) from the full sample by following a specific resampling scheme. The parameter of interest is estimated from each replicate, and we use the variability among the replicate estimates to estimate the overall variance of the parameter estimate. The jackknife method deletes groups of PSUs from the full sample to create the replicates, and modifies the original weights to obtain replicate weights. We calculate replicate weights at each stage of the estimation (design weights and calibrated weights). One hundred replicates are created. Deleting groups of PSUs, instead of one PSU at a time, saves a considerable amount of computing time and allows us to calculate any sampling errors relatively quickly. Once the replicate weights are created, the variance can be calculated using the following formula:

$$V(\hat{y}) = \frac{G-1}{G}\sum_{g=1}^{G}\left[\left(\hat{y}_{(g)} - \hat{y}\right)^2\right]$$

where G is the total number of groups and g the replicate group.

Sampling errors (with 95% significance) are published quarterly for each cell in the published tables and for estimates of change between adjacent quarters. Figures 4–6 show time series of these confidence intervals for employed, unemployed, and not in the labour force estimates.

Figure 4

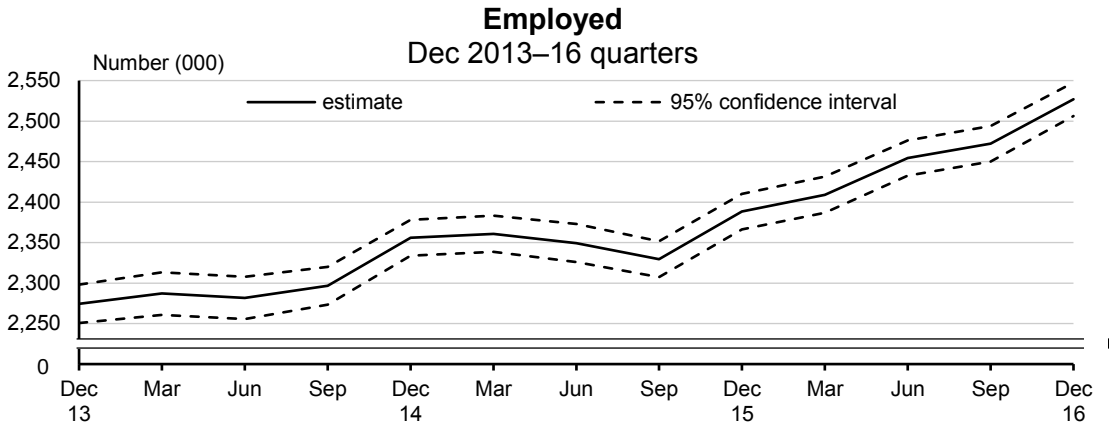

Source: Statistics New Zealand

Figure 5

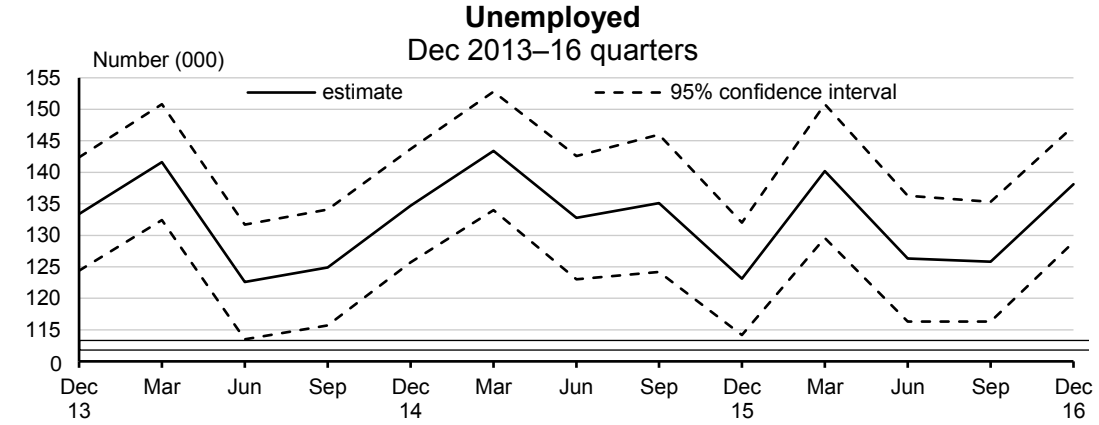

Source: Statistics New Zealand

Figure 6

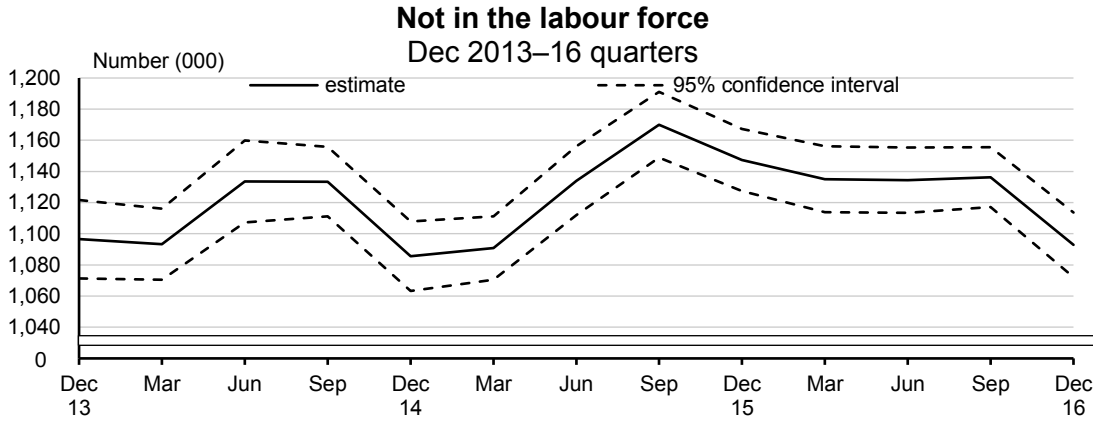

Source: Statistics New Zealand

Non-sampling errors

Non-sampling error arises from inaccuracies in collecting, measuring, and processing the data, and can arise at any stage of the survey process. Measuring non-sampling error is much more difficult than measuring sampling error (see Sampling errors). In many cases the reasons for the non-sampling error are not known, whereas sampling error is a direct result of the survey design and is under the control of the sample designer.

Non-sampling error may be related to:

- the definition of the target population – excluding groups within the scope of the survey or including groups outside the scope will cause non-sampling errors
- selection of the sampling frame – a frame that does not match the target population causes bias if the excluded units have different characteristics from the included ones (coverage error).
- the sample design should give everyone in the target population a known chance of selection
- non-response – if people who do not respond are different to those who do respond, then estimates from the achieved sample may be biased estimates of population values – this can be corrected to some extent by weighting the sample
- questionnaire development – if the questionnaire does not ask exactly what is required, in a way that is easy for the respondent to understand and answer, the resulting data quality will not be as high as required
- social desirability effects – respondents may be unwilling to admit to socially undesirable behaviour, but this should be minimal in a questionnaire about labour force characteristics
- collection of information (eg mode effects, proxies) – bias from mode effects may arise when questions are asked or interpreted differently between different modes (eg computer-assisted face-to-face mode and a paper self-completion questionnaire). Error from using proxies may arise if the quality of proxy data is not as high as that collected from the actual intended respondents
- inadequate interviewer training – if interviewers are less successful at gaining responses, or influence the answers of respondents, bias and variability can be introduced into the survey; a respondent should answer the same way irrespective of the interviewer (and regardless of whether it is by phone, or face-to-face)
- data coding and entering – for example, data may be classified incorrectly or mistakes made inadvertently when editing the data.

Coverage error is minimised by using an accurate and up-to-date sampling frame. An out-of-date sampling frame can have the following effects on data quality:

- bias in survey estimates if new elements differing from elements already on the frame are not added to the frame
- increase in the variance of estimates if auxiliary information used for stratification or estimation is inaccurate.

Non-response bias is introduced when non-respondents have different characteristics to those who respond. We make every effort to minimise non-response – including call-backs, alternative mode options, interview times, pre-notification letters for CAPI interviews, assurances of

confidentiality, interviewer training, and questionnaire design. Additionally, we use weighting adjustments to try and adjust for the bias (see Estimation and imputation).

We also aim to minimise error that could arise from questionnaire development, collecting information, and interviewer effects, by careful design of questionnaires (including cognitive testing and pilot surveys before introducing a new questionnaire), and intensive training and supervision of interviewers. We only use proxy interviews when it is not possible or practicable to obtain the answers from the actual intended respondent.

With ongoing changes in technology and the economy, questionnaire design also needs to evolve to keep up-to-date with these changes to information sources, methods of capturing responses, terminology, new business concepts, legislation, and shifts in the economy or new industries.

Processing errors (including data coding and editing errors) are minimised by using high quality editing and imputation procedures. We thoroughly test all programming code used in the editing and coding of answers. We also have a substantial array of checks (see Coding and processing the data) in place to identify and fix most (but not necessarily all) errors. Additionally, survey edits are incorporated into the computer assisted interviews so the respondent can be asked about suspect responses.

For the HLFS, we report on three estimates that are indicators of aspects of non-sampling errors. These are response rates, undercoverage rates, and proxy rates. They are discussed in more detail in the sections below.

See Non-sampling error in economic surveys at Statistics New Zealand for more detail on non-sampling error. Much of this content applies to the HLFS.

Response rates and achieved sample characteristics

The response rate indicates what percentage of eligible households responded to the survey. Non-response bias is one source of non-sampling error – the response rate provides a measure of one component of non-response bias. The other component is the difference between the true answer to a question for respondents and non-respondents. Even if the response rate is high, large differences in the true answer may lead to substantial non-response bias. However, this component is not often available for us to measure.

The response rate is a survey estimate and is reported each quarter. It is generally consistent over time, but can be affected by changes in coding practices. The response rate can be calculated as:

$$\frac{D}{D + C + E\left(\frac{C + D}{B + C + D}\right)}$$

where:

$$A = sum\ of\ the\ design\ weights\ from\ all\ ineligible\ pre - contact\ households$$

$$B = sum\ of\ the\ design\ weights\ from\ all\ ineligible\ post - contact\ households$$

$$C = sum\ of\ the\ design\ weights\ from\ all\ eligible\ non - responding\ households$$

$$D = sum\ of\ the\ design\ weights\ from\ all\ eligible\ responding\ households$$

$$E = sum\ of\ the\ design\ weights\ from\ all\ unknown\ eligibility\ households$$

Prior to the June 2016 quarter, the HLFS had a target response rate of 90 percent. Achieved response rates for the last four years have been around 85 percent on average.

Figure 7

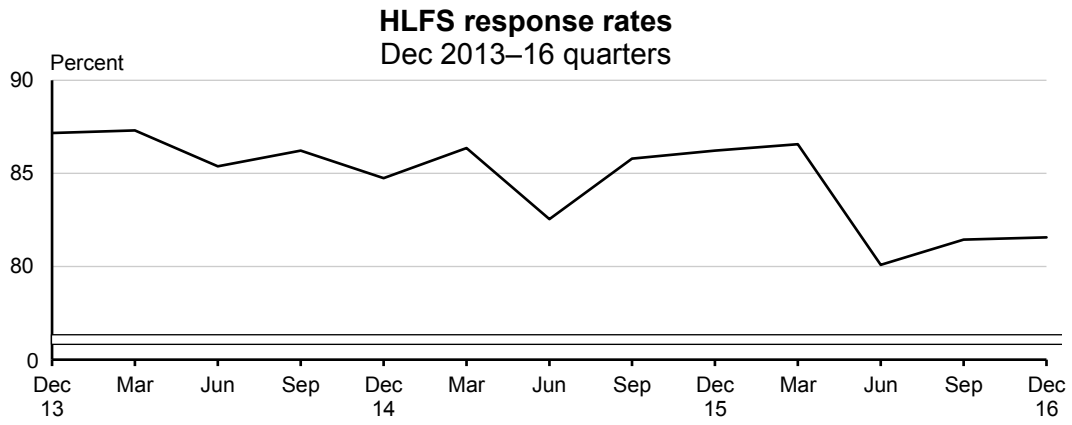

Source: Statistics New Zealand

The response rate is sensitive to how we classify household eligibility. It can vary when the overall sample and the achieved sample are constant. The HLFS therefore now has a target based on the percentage of the total sample classified as 'eligible responding households' (also called the achieved sample rate). This is calculated as:

$$\frac{D}{A + B + C + D + E}$$

The HLFS has a target achieved sample rate of 76 percent. The achieved sample rate has been around 74 percent on average over the last four years.

Figure 8

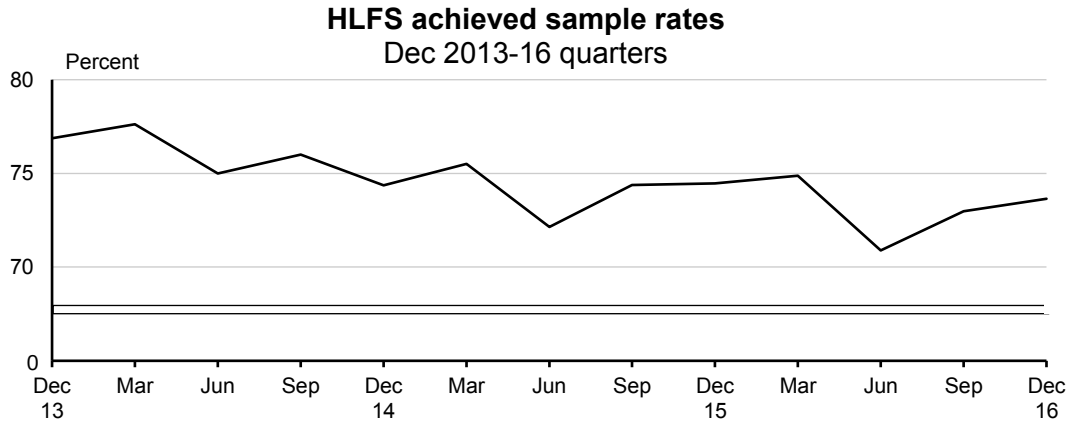

Source: Statistics New Zealand

The HLFS now also reports the achieved sample characteristics every quarter, pre- and post-calibration weights. Obtaining a sample that represents the population is essential when it comes to producing reliable labour market estimates.

Non-response is made up primarily of households where interviewers fail to make contact (full non-contact households) and those contacted but who refuse to take part in the survey (full refusal households). The refusal rate is calculated as:

$$\frac{R}{D + C + E\left(\frac{C + D}{B + C + D}\right)}$$

where $R = sum\ of\ the\ design\ weights\ from\ all\ full\ refusal\ households$

Similarly, the non-contact rate can be calculated as:

$$\frac{NC}{D + C + E\left(\frac{C + D}{B + C + D}\right)}$$

where $NC = sum\ of\ the\ design\ weights\ from\ all\ full\ non - contact\ (verified\ and\ non - verified)\ households$

Figure 9 shows the non-contact rate has been around 15 percent on average over the last three years, while the refusal rate has been around 3 percent.

Figure 9

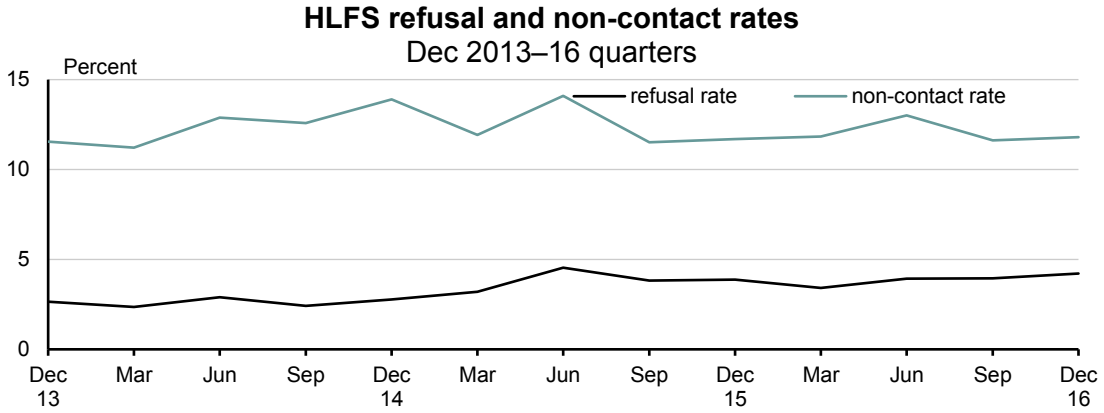

Source: Statistics New Zealand

Imputation

We impute values for sex, age, ethnicity, looking for full-time or part-time employment, and actual and usual hours worked where they are missing for people who belong to eligible responding dwellings (see Estimation and imputation). Very few respondents have sex missing. Figure 10 shows the imputation rates for age and looking for full- or part-time employment.

Figure 10

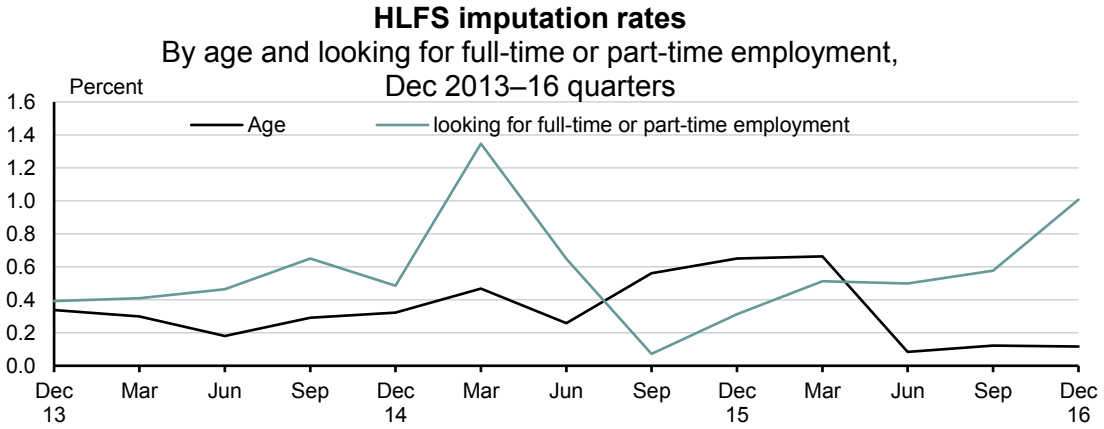

Source: Statistics New Zealand

Imputation aims to produce distributions that reflect as near as possible the distributions of the population. However, we do not have the true population data to be able to compare the survey distribution against. Therefore, we can only compare the distributions of the original (pre-imputation) and final (post-imputation) data. The imputation process should not make dramatic changes to the distribution of the original data.

Table 5 shows the distributions of full-time and part-time employment status after imputation of hours worked remain very similar to the distributions before imputation.

Table 5

Pre- and post-imputation distributions

Quarter	Employment status	Pre-imputation distribution	Post-imputation distribution
		Percent	
Jun-16	Full-time	78.19	78.05
	Part-time	21.81	21.95
Sept-16	Full-time	78.32	78.03
	Part-time	21.68	21.97
Dec-16	Full-time	78.76	78.48
	Part-time	21.24	21.52
Source: Stats New Zealand			

Proxies

The HLFS allows interviewers to take responses from proxies if a respondent is unavailable or unable to answer the questions themselves. Currently the HLFS has reasonably relaxed rules around who can and cannot proxy for another individual (see Collection methodology for more details).

While using proxies increases the response rate (and therefore may reduce non-response bias), it may introduce measurement error, since we expect capable self-respondents to provide higher quality data than proxy respondents. People may remember first-hand events more accurately than second-hand information, because they experience them in a more vivid and detailed form than events they merely hear or read about (Sudman et al, 1994, in Tourangeau, 2000).

However, the literature surrounding proxy and self-respondents provides mixed results, with some studies finding no difference in responses between the two, and others finding significant differences in answers. Where significant differences were found, they do not exhibit a consistent bias of proxy responses relative to self-respondents and/or administrative data (eg Moore, 1988; Martin & Butcher, 1982; Tamborini & Kim, 2013; Hill, 1997; Boehn, 1989).

Although the evidence is mixed, we expect proxies may not be as accurate as self-responses. Therefore, the HLFS monitors the rate of proxy responses – to gauge the quality of responses. The proxy rate is calculated as the percentage of respondents who had someone else respond on their behalf divided by the total number of respondents.

Figure 11 shows the proxy rates by quarter for the HLFS. Generally, the proxy rate is lower in quarters where a supplement is run than in non-supplement quarters – previously there have been more rules about an individual acting as a proxy for most supplements.

Figure 11

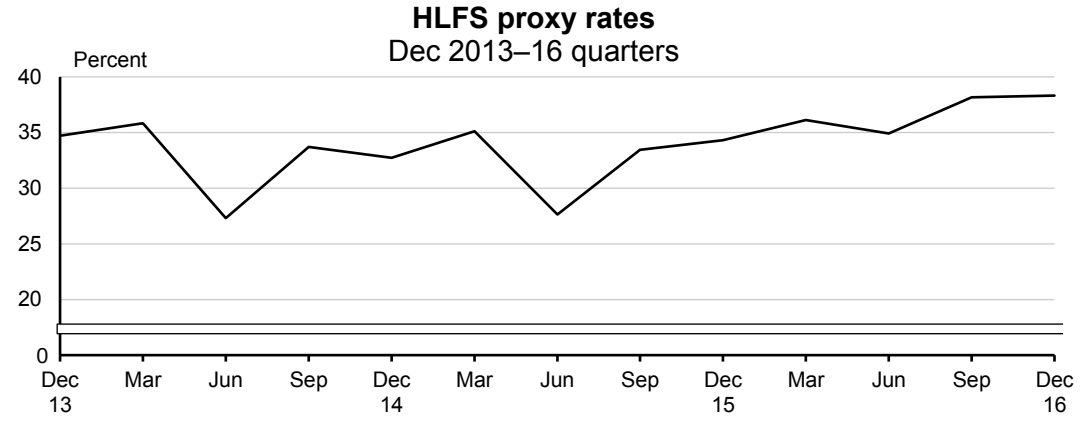

Source: Statistics New Zealand

Undercoverage

Undercoverage is another source of non-sampling error. We estimate the undercoverage rate as the difference in survey estimates before and after calibration. This indicates the level of possible undercoverage, but does not measure its effect on survey estimates.

Typically, estimates before calibration are too small. This represents undercoverage. The undercoverage rate indicates how representative the pre-calibrated sample is. The higher the rate, the less representative the pre-calibrated sample is. Figure 12 shows the overall undercoverage rate for the HLFS has been, on average, around 17 percent, and that males have a higher rate than females. Figure 13 shows the youngest age group (15–34 years) is underrepresented in the sample before calibration.

Figure 12

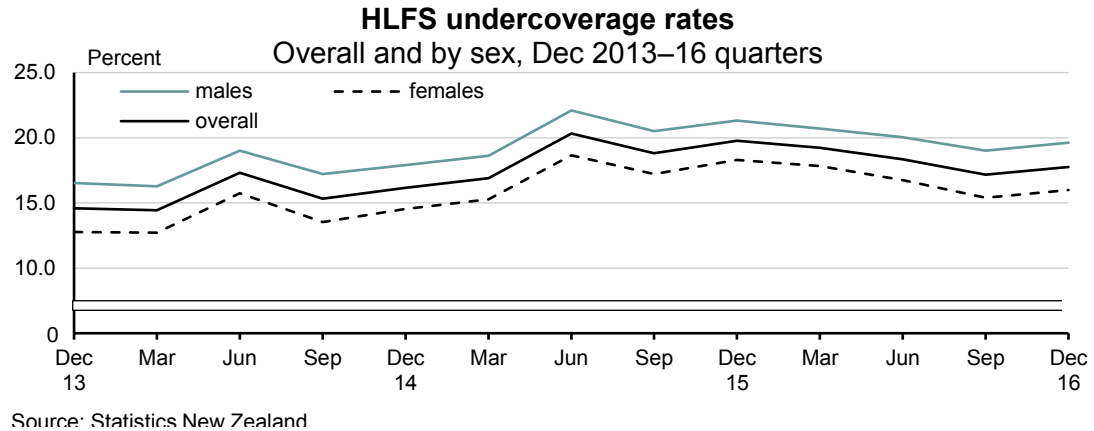

Source: Statistics New Zealand

Figure 13

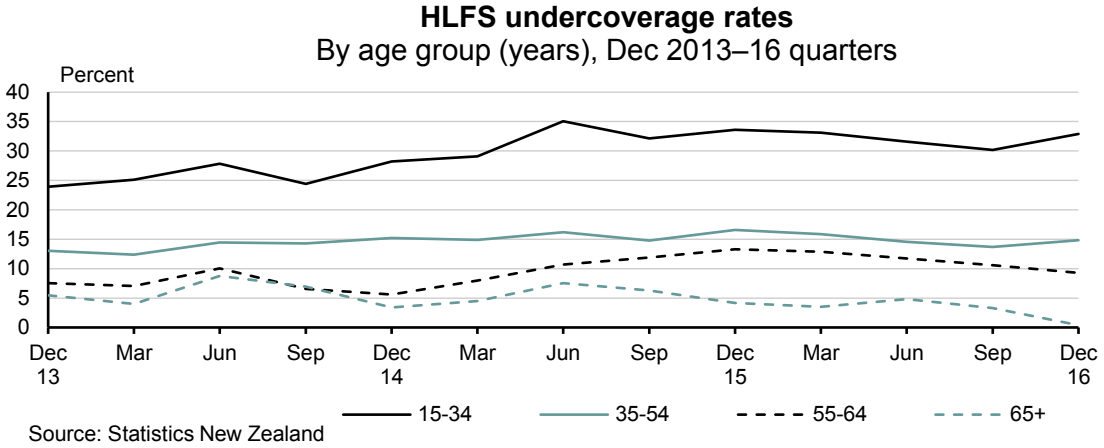

Source: Statistics New Zealand

Changes affecting data comparability over time

In this chapter we detail the major changes made to the HLFS, both in the past and as a result of the 2016 redevelopment.

Despite the over-arching desire for long-term comparability of the series, careful and continuing maintenance and development of a survey is essential to maintain the integrity of the data and the efficiency of the collection. Since it began in 1985, the HLFS has had changes to or development of the frequency of collection, collection and sampling methods, concepts and definitions, classifications, and time-series analysis techniques.

Changes due to the latest questionnaire re-design

The 2016 redevelopment of the HLFS is the first substantial change to the survey since it was first introduced in December 1985. The first quarter to use the redeveloped HLFS questionnaire was the June 2016 quarter.

The changes to the HLFS improve the quality and accuracy of some of the key labour market indicators – for example, the questionnaire design has been greatly improved, which should reduce non-sampling errors. However, some of the improvements mean we need to make changes to previously published estimates in order to retain a consistent time series. In some cases, this backdating might not be possible and there will be a conceptual break in the time series.

For information on which series are affected, see Household Labour Force Survey – Revisions to labour market estimates, and for an overview of the key changes made to the questionnaire, see Household Labour Force Survey – summary of 2016 redevelopment.

Changes due to the latest sample re-design

The major innovations to the sample design from the latest redesign are:

- refining the stratification and allocation to urban vs rural
- changing to PPS sampling based on Māori density
- moving from Kish allocation to proportional allocation by regions
- introducing a new overlap control method, including management of overlap with selected surveys from other organisations
- using 'not in the labour force' (NILF) densities and NZ Deprivation Index (NZDep) quintiles in the stratification.

The redesign includes proportional allocation to 12 regions (the regions currently used in HLFS dissemination) instead of Kish allocation to 14 regions; over-sampling main urban areas only and at a lower rate than the current over-sampling of all urban (including secondary and minor) areas; proportional allocation to a new stratification layer based on the proportion within PSUs not in the labour force; and finally, a Neyman allocation (based on unemployment) to a stratification layer based on NZDep quintiles.

All this results in slightly fewer strata than under the previous sample design (down from 119 to 108 strata), and less disproportionate allocation of PSUs to strata. Additionally, the new overlap control methodology allows us to employ systematic sampling by TA. Finally, targeting Māori no

longer occurs through a Māori stratification layer but instead through the PPS selection of PSUs, which incorporates the Māori density within PSUs in the size measure.

Note: Previously, we over-sampled Māori through disproportionately sampling a high Māori stratum, which was formed of PSUs where 12 percent or more of adults identified as Māori in the census.

Historical changes

Table 6 lists the key changes made to the HLFS since its inception, while table 7 shows the supplements we've run.

Table 6

Main historical changes to the HLFS

Quarter	Date	Description of change
2	Mar 1986	Interviewer coding of industry using NZSIC at a 2-digit level introduced. Occupation estimates coded using NZSCO68 (interviewer coded to a 2-digit level).
8	Sep 1987	Data collection practices changed to minimise respondent burden for older respondents. Respondents aged 65+ now only supplied information (after their first interview) if their labour market position altered, and if they met other criteria in the previous interview.
19	Jun 1990	Sample size doubled to 24,000 households (1,080 more PSUs sampled) to produce more accurate regional estimates and enable us to publish monthly national estimates (previous results published quarterly). Minor changes to questionnaire.
19	Jun 1990	Coding for industry changed to central coding using NZSIC to a 3-digit level. Coding for occupation changed to central coding of NZSCO68 to 3-digit level.
22	Mar 1991	Occupation estimates changed to central coding of NZSCO90 to 3-digit level.
24	Sep 1991	Sample size reduced to 16,000 households; HLFS reverted to quarterly collection. Estimates produced for 10 regional council areas rather than 15.
25	Dec 1991	Changes made to ethnic classification: 'Other Pacific' category added (previous categories were: European, NZ Māori, Samoan, Cook Island Maori, Niuean, Tongan, Chinese, Indian, and Other).
33–36	Dec 1993	Sample redesigned and phased in over four quarters. 1,752 PSUs now sampled, with nine households surveyed in each (2,400 PSUs in previous design; 6.5 households per PSU). Improved regional estimates (regional councils now the basis of stratification), for 12 regional council areas rather than 10.
39	Jun 1995	Non-private households excluded from survey population (still included in target population).
42	Mar 1996	Occupation estimates changed to central coding of NZSCO95 to 3-digit level.
45	Dec 1996	HLFS processing system migrated from mainframe to LAN environment.

Quarter	Date	Description of change
45	Dec 1996	Coding for industry changed to 4-digit NZSIC code – 4th digit ensured a unique concordance to ANZSIC at 3-digit level.
51	Jun 1998	Integrated weighting replaced post-stratification (applied back to June 1995 quarter; later backcast to include quarters 2–38).
53	Dec 1998	ANZSIC industry estimates from HLFS published – back to December 1996 quarter. NZSIC industry coding phased out.
54-57	Mar 1999	Redesign implemented (over four quarters). 1,760 PSUs sampled with nine households surveyed in each. No rural strata cross regional council boundaries. Asian stratum added.
63	Jun 2001	Total actual hours worked series revised (correcting misalignment between survey and calendar quarters).
70	Mar 2003	Industry coding changed from NZSIC/ANZSIC dual coding to ANZSIC-only coding (centrally coded to 3-digit level) – may have introduced a slight discontinuity to industry series.
74-81	Mar 2004	Redesign implemented (over eight quarters). Several questionnaire changes made to determine formal study status, and enable us to derive an 'extended unemployment' estimate.
79	Jun 2005	Data collection converted from pen and paper interviewing (PAPI) to computer-assisted personal interviewing (CAPI). The rotating-in group for this quarter was interviewed using CAPI. Testing for a CAPI effect based on this one rotation group revealed no significant effect on the unemployment rate, labour force participation rate, estimate of underemployed, or the estimates of actual and usual hours of work.
80	Sep 2005	CAPI group in quarter 79 converted to computer-assisted telephone interviewing (CATI). All other rotation groups stayed with PAPI. CATI group had a statistically significant higher unemployment rate than PAPI groups. Other variables tested not affected. Immediate investigation could not explain the increase, so we dropped results from CATI group from September 2005 quarter estimates.
81	Dec 2005	CATI group in quarter 80 remained CATI and two new groups converted to CATI. The newly rotated group converted to CAPI. Final results from December quarter indicated the two new CATI groups had slightly elevated unemployment rate; original CATI group remained high. We concluded the high figure for original CATI group was a characteristic of that group – leaving it out biased the result downwards. We revised the previous September 2005 estimate in the December 2005 release.
82	Mar 2006	No additional groups converted to CAI (either personal or telephone) – 50 percent of the sample interviewed using CAI and the rest using PAPI. The CAI effect on the unemployment rate exhibited in the previous two quarters disappeared.
85	Dec 2006	Whole sample now interviewed using CAI modes.
89	Dec 2007	Changes made to ethnicity classification, using the single/ combination output method – caused a break in ethnicity series.
96	Sep 2009	Dual coding using ANZSCO06 and NZSCO99 for occupation introduced, and dual coding using ANSZIC06 and ANZSZIC96 for industry introduced.

Quarter	Date	Description of change
113	Dec 2013	Imputation of 75+ households in quarters subsequent to their first time introduced (excludes June quarters, see Estimation and imputation for details).
117–124	Dec 2014	Sample redesigned and phased in over eight quarters – includes updated stratification variables, changes to allocation, and change from simple random sampling of PSUs within strata to probability proportional to size sampling within strata.
117	Dec 2014 (re-release Mar 2015)	Regional benchmark introduced and non-response adjustment step in the weighting removed.
123	June 2016	Introduction of new questionnaire. Some of the more significant changes come as a result of: • improved accuracy in identifying active jobseekers – looking at job advertisements on the internet is now correctly classified as not actively seeking work • better identification of both the number of people employed and their employment status due to the ways in which people are asked about employment • the change in the target and survey populations to include defence force personnel living in private dwellings
Source: Stats New Zealand		

Table 7

Supplements to the HLFS

Quarter	Date	Supplement attached
44	Sep 1996	Education and training
47	Jun 1997	New Zealand Income Survey (annual basis)
52	Sep 1998	New Zealand Childcare Survey
58	Mar 2000	Survey of Older People
66	Mar 2002	Cultural Experiences Survey
85	Dec 2006	Household Use of Information and Communication Technology (ICT) Survey
86	Mar 2007	Survey of Dynamics and Motivation for Migration in New Zealand
90	Mar 2008	Survey of Working Life
96	Sep 2009	New Zealand Childcare Survey
97	Dec 2009	ICT Survey
108	Sep 2012	ICT Survey
109	Dec 2012	Survey of Working Life

References

Bell, P (2011). *Conditional Selection for Business and Household Surveys*. Department of Public Health, University of Otago, Wellington. New Zealand Socioeconomic Deprivation Index.

Boehm, LE (1989). Reliability of Proxy Response in the Current Population Survey. In Proceedings of the American Statistical Association, Survey Research Methods Section, 486–489.

Chowdhury, S, Chu, A, & Kaufman, S (2000). *Minimizing overlap in NCES surveys.* Proceedings of the Survey Methods Research Section, American Statistical Association, 174-179. Retrieved from www.amstat.org/sections/srms/Proceedings/papers/2000_025.pdf

Deville, J-C & Sārndal, C-E (1992). Calibration estimators in survey sampling. *Journal of the American Statistical Association, 87*(418), 376–382.

Eurostat (2001). *The European Union labour force survey methods and definitions*. Luxembourg: EC. PDF available from EU labour force survey – methodology.

Hill, DH (1987). *Response errors in labor surveys: Comparisons of self and proxy reports in the survey of income and program participation*. Third Annual Research Conference, Washington D.C.

Lu, K (2012). Minimizing sample overlap with surveys using different geographic units. Applied Statistics Education and Research Collaboration (ASEARC) – Conference papers. PDF available from http://ro.uow.edu.au/asearc/

Martin, J, & Butcher, B (1982). The quality of proxy information: Some results from a large scale study. *The Statistician, 31*(4), 293–319.

Moore, J (1988). Miscellanea, self/proxy response status and survey response quality, a review of the literature. *Journal of Official Statistics, 4*(2), 155–172.

Tamborini, C, & Kim, CH (2013). Are proxy interviews associated with biased earnings reports? Marital status and gender effects of proxy. *Social Science Research, 42*(2), 499–512.

Tourangeau, R, Rips, L, & Rasinski, K (2000). *The psychology of survey response*. Cambridge University Press.

Appendix 1: Total PSUs and number of PSUs selected per stratum

Appendix table 1

Stratification dimensions and total number of PSUs within each stratum

region	primary urban (1=urban, 0=not urban)	high-NILF (1=high NILF, 0=low NILF, 2=not used)	Final Stratum / NZ Dep group	region	primary urban	high-NILF	Final Stratum / NZ Dep group
01 (Northland)	0	0	001	809	533	284	142
			002				142
		1	003			249	125
			004				124
	1	0	005		276	139	139
		1	006			137	137
02 (Auckland)	0	0	007	6549	501	389	195
			008				194
		1	009			112	112
	1	0	010		6048	4675	468
			011				468
			012				468
			013				468
			014				468
			015				468
			016				468
			017				468
			018				468
			019				463
		1	020			1373	229
			021				229
			022				229
			023				229
			024				229
			025				228

Appendix table 1 (cont)

Stratification dimensions and total number of PSUs within each stratum

region	primary urban (1=urban, 0=not urban)	high-NILF (1=high NILF, 0=low NILF, 2=not used)	Final Stratum / NZ Dep group	region	primary urban	high-NILF	Final Stratum / NZ Dep group
03 (Waikato)	0	0	026	2144	1131	737	246
			027				246
			028				245
		1	029			394	197
			030				197
	1	0	031		1013	800	200
			032				200
			033				200
			034				200
		1	035			213	107
			036				106
04 (Bay of Plenty)	0	0	037	1400	493	316	158
			038				158
		1	039			177	177
	1	0	040		907	640	214
			041				214
			042				212
		1	043			267	134
			044				133
05 (Gisborne / Hawke's Bay)	0	2	045	1124	239	239	120
			046				119
	1	0	047		885	611	204
			048				204
			049				203
		1	050			274	137
			051				137
06 (Taranaki)	0	2	052	659	340	340	170
			053				170
	1	2	054		319	319	160
			055				159

Appendix table 1 (cont)

Stratification dimensions and total number of PSUs within each stratum

region	primary urban (1=urban, 0=not urban)	high-NILF (1=high NILF, 0=low NILF, 2=not used)	Final Stratum / NZ Dep group	region	primary urban	high-NILF	Final Stratum / NZ Dep group
07 (Manawatu-Whanganui)	0	0	056	1332	649	394	197
			057				197
		1	058			255	128
			059				127
	1	0	060		683	453	227
			061				226
		1	062			230	115
			063				115
08 (Wellington)	0	0	064	2634	307	202	101
			065				101
		1	066			105	105
	1	0	067		2327	1942	216
			068				216
			069				216
			070				216
			071				216
			072				216
			073				216
			074				216
			075				214
		1	076			385	193
			077				192

Appendix table 1 (cont)

Stratification dimensions and total number of PSUs within each stratum

region	primary urban (1=urban, 0=not urban)	high-NILF (1=high NILF, 0=low NILF, 2=not used)	Final Stratum / NZ Dep group	region	primary urban	high-NILF	Final Stratum / NZ Dep group
09 (Nelson / Tasman / Marlborough / West Coast)	0	0	078	956	626	485	243
			079				242
		1	080			141	141
	1	2	081		330	330	165
			082				165
10 (Canterbury)	0	0	083	3003	1013	804	201
			084				201
			085				201
			086				201
		1	087			209	105
			088				104
	1	0	089		1990	1646	206
			090				206
			091				206
			092				206
			093				206
			094				206
			095				206
			096				204
		1	097			344	172
			098				172
11 (Otago)	0	0	099	1220	532	420	210
			100				210
		1	101			112	112
	1	0	102		688	491	246
			103				245
		1	104			197	197
12 (Southland)	0	2	105	609	287	287	144
			106				143
	1	2	107		322	322	161
			108				161

Appendix table 2

Number of PSUs selected for Northland region

By stratum

Region	Stratum		
	Urban/other	**High/low NILF**	**NZDep**
Northland 61	Areas other than main urban 34	Low NILF 19	9
			10
		High NILF 15	7
			8
	Main urban 27	Low NILF 14	14
		High NILF 13	13

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 3

Number of PSUs selected for Auckland region
By stratum

Region	Stratum		
	Urban/other	**High/low NILF**	**NZDep**
Auckland	Areas other than main urban	Low NILF	8
512	31	22	14
		High NILF	
		9	9
	Main urban	Low NILF	26
	481	367	32
			35
			34
			36
			38
			37
			42
			44
			43
		High NILF	13
		114	15
			16
			19
			23
			28

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 4

Number of PSUs selected for Waikato region
By stratum

Region	Stratum		
	Urban/other	High/low NILF	NZDep
Waikato	Areas other than main urban	Low NILF	14
178	75	49	15
			20
		High NILF	12
		26	14
	Main urban	Low NILF	22
	103	82	17
			20
			23
		High NILF	10
		21	11

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 5

Number of PSUs selected for Bay of Plenty region
By stratum

Region	Stratum		
	Urban/other	High/low NILF	NZDep
Bay of Plenty	Areas other than main urban	Low NILF	10
101	28	20	10
		High NILF	
		8	8
	Main urban	Low NILF	12
	73	53	20
			21
		High NILF	8
		20	12

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 6

Number of PSUs selected for Gisborne/Hawke's Bay region

By stratum

Region	Stratum		
	Urban/other	High/low NILF	NZDep
Gisborne/Hawke's Bay	Areas other than main urban	Not split	8
91	16	16	8
	Main urban	Low NILF	16
	75	53	18
			19
		High NILF	9
		22	13

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 7

Number of PSUs selected for Taranaki region

By stratum

Region	**Stratum**		
	Urban/other	**High/low NILF**	**NZDep**
Taranaki	Areas other than main urban	Not split	12
52	23	23	11
	Main urban	Not split	13
	29	29	16

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 8

Number of PSUs selected for Manawatu-Wanganui region
By stratum

Region	Stratum		
	Urban/other	High/low NILF	NZDep
Manawatu-Wanganui	Areas other than main urban	Low NILF	11
108	45	28	17
		High NILF	7
		17	10
	Main urban	Low NILF	17
	63	44	27
		High NILF	7
		19	12

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 9

Number of PSUs selected for Wellington region
By stratum

Region	Stratum		
	Urban/other	High/low NILF	NZDep
Wellington	Areas other than main urban	Low NILF	6
204	17	11	5
		High NILF	
		6	6
	Main urban	Low NILF	12
	187	157	13
			14
			18
			17
			20
			17
			24
			22
		High NILF	12
		30	18

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 10

Number of PSUs selected for Tasman/Marlborough/Nelson/West Coast region
By stratum

Region	Stratum		
	Urban/other	High/low NILF	NZDep
Tasman/Marlborough/Nelson/West Coast	Areas other than main urban	Low NILF	13
78	42	33	20
		High NILF	
		9	9
	Main urban	Not split	15
	36	36	21

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 11

Number of PSUs selected for Canterbury region

By stratum

Region	Stratum		
	Urban/other	High/low NILF	NZDep
Canterbury	Areas other than main urban	Low NILF	15
235	67	52	10
			11
			16
		High NILF	8
		15	7
	Main urban	Low NILF	10
	168	135	15
			15
			15
			18
			19
			21
			22
		High NILF	16
		33	17

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 12

Number of PSUs selected for Otago region
By stratum

Region	Stratum		
	Urban/other	High/low NILF	NZDep
Otago	Areas other than main urban	Low NILF	12
103	38	30	18
		High NILF	
		8	8
	Main urban	Low NILF	18
	65	47	29
		High NILF	
		18	18

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix table 13

Number of PSUs selected for Southland region
By stratum

Region	Stratum		
	Urban/other	High/low NILF	NZDep
Southland	Areas other than main urban	Not split	7
45	18	18	11
	Main urban	Not split	11
	27	27	16

Note: NILF is not in labour force stratum; NZDep is New Zealand Deprivation Index.

Source: Stats NZ

Appendix 2: Eligibility tables

Tables in this chapter describe the eligibility status codes assigned to households using labour force status, scope code, and participation code.

Appendix table 14

Eligibility status categories

Category	Description	Condition to be met
1	Ineligible pre-contact	Response status code for household questionnaire is in (411,412,413,414)
2	Ineligible post-contact	Response status code for household questionnaire is in (611,615,616) **OR** Response status code for personal questionnaire for **all** household members is in (613,614,616)
3	Eligible non-responding	Response status code for household questionnaire is in (111) **AND** Response status code for personal questionnaire for **any** household member is in (211,212,213,214,311,312,313,314,315,317,511,512, 513) **AND** No household member had DVLFS in (1,2,3)
4	Eligible responding	(Response status code for household questionnaire is in (111) **AND** At least one household member has response status code for personal questionnaire is in (111,211,212,213,214,311,312,313,314,315,511,512, 513) **AND** DVLFS in (1,2,3)) **OR** Response status code for household questionnaire is in (954)[1]
5	Unknown eligibility	Response status code for household questionnaire is in (211,212,213,214,215,217,311,312,314,315,316,317, 511,512,513,952[2])
Source: Stats NZ		

[1] 954 refers to 75+ households. This is a temporary code that may change number (or not be needed) once we move to a new collection platform.

[2] 952 refers to the 'Do not interview' code that only HLFS uses (eg a gang house that is unsafe to visit). It is a 'completed' code. This code rolls -over from quarter to quarter (unlike code 213). This is a temporary code that may change number (or not be needed) once we move to a new collection platform.

Appendix table 15

Labour force status categories

Category	Description
1	Employed
2	Unemployed
3	Not in the labour force
9	Labour force status unidentified
Source: Stats NZ	

Appendix table 16

HQ response status codes

Category	Description
111	Complete: Response
211	Complete: Unable to contact~
212	Complete: Confirmed away for the survey period~
213	Complete: Cannot contact – Health and safety~
214	Complete: Cannot contact – Natural disaster~
215	Complete: Address not visited/phoned~
217	Complete: Unable to access building~
311	Complete: No interview – Illness~
312	Complete: No interview – Bereavement~
314	Complete: No interview – Language~
315	Complete: No interview – Disability~
316	Complete: Unable to re-contact~
317	Complete: Insufficient information to continue
411	Complete: Dwelling vacant/empty~
412	Complete: Dwelling under construction~
413	Complete: Non-private dwelling~
414	Complete: Dwelling derelict/demolished/non-dwelling~
511	Complete: Refusal~
512	Complete: Refusal during interview~
513	Complete: Refusal during interview
611	Complete: Out of scope~
615	Complete: Not main residence (holiday/second home)~
616	Complete: Out of scope
Source: Stats NZ Note: ~ means interviewers are able to add comments.	

Appendix table 17

PQ response status codes

Category	Description
111	Complete: Response
211	Complete: Unable to contact~
212	Complete: Confirmed away for the survey period~
213	Complete: Cannot contact – Health and safety~
214	Complete: Cannot contact – Natural disaster~
311	Complete: No interview – Illness~
312	Complete: No interview – Bereavement~
313	Complete: No interview – Deceased~
314	Complete: No interview – Language~
315	Complete: No interview – Disability~
317	Complete: Insufficient information to continue
511	Complete: Refusal~
512	Complete: Refusal during interview~
513	Complete: Refusal during interview
613	Complete: Person found ineligible in interview~
614	Complete: Person removed from survey~
616	Complete: Out of scope

Source: Stats NZ

Note: ~ means interviewers are able to add comments.

Appendix 3: Weighting

This appendix explains in more detail the stages of weighting used for the HLFS.

PSU selection weight

The selection weight d_i for each PSU i is calculated as the inverse of the probability of selection π_i so that PSUs with a lower probability of selection receive a higher selection weight.

$$d_i = {}^{1}/_{\pi_i}$$

Within strata, PSUs are selected with PPS. That is, the selection probability of PSU i of stratum h is:

$$\pi_{ih} = n_h \frac{s_{ih}}{\sum s_{ih}}$$

where n_h is the number of PSUs to be selected from stratum h, and s_{ih} is the size of PSU i of stratum h. In addition, size is defined as:

$$s_{ih} = r_{ih}\sqrt{0.02 + p_{ih}}$$

where r_{ih} is the number of occupied or under-construction dwellings in PSU i of stratum h, and p_{ih} is the proportion of adults in the PSU who are Māori (see HLFS sample design for details).

Previously (ie in older designs), we selected a simple random sample of PSUs from each stratum. In this case, the selection weight was calculated as a ratio estimator because PSU sizes can vary considerably. The ratio estimator used was:

$$w_{hi} = \frac{\sum_{i=1}^{N_h} x_{hi}}{\sum_{i=1}^{N_h} \delta_{hi} x_{hi}}$$

where

$$\delta_{hi} = \begin{cases} 1, & \text{if PSU } i \text{ of stratum } h \text{ is selected} \\ 0, & \text{otherwise} \end{cases}$$

and x_{hi} is some known quantity such as PSU size. Consider the sample estimate of some stratum total:

$$\hat{Y}_h = \frac{\sum_{i=1}^{N_h} x_{hi}}{\sum_{i=1}^{N_h} \delta_{hi} x_{hi}} \sum_{i=1}^{N_h} \delta_{hi} y_{hi}$$

If $x_{hi} = y_{hi}$ then $\hat{Y}_h$ will be estimated without any error. More generally, as long as the relationship between the x_{hi} and y_{hi} is modelled well by a regression through the origin, the ratio estimator can reduce sampling variability considerably. The stronger the relationship the greater the improvement.

For the HLFS, x_{hi} was typically the PSU size, which is a count of the number of occupied private dwellings and dwellings under construction at the time of the census.

Household selection weight

We multiply the PSU selection weight by a household selection weight to account for the selection of households within PSUs. For the new sample, the number of addresses to select per PSU is calculated as the total desired sample size divided by the number of selected PSUs.

However, as discussed in HLFS sample design, we require a minimum of five panels in any PSU. Some PSUs do not contain enough addresses to make five panels of the implied number of addresses (given the total desired sample size) to select per PSU. Therefore, we select fewer addresses per panel in these PSUs (ie the sampling interval is set to 5). This means the actual number of selected addresses falls short of the desired sample size. We therefore run the exercise again, this time excluding those small PSUs – the number of addresses to select per PSU is re-calculated in the remaining PSUs, as:

$$\frac{\text{desired total sample size} - \text{selected sample size in the small PSUs}}{\text{remaining number of PSUs}}$$

This process is repeated until the actual number of addresses selected is the same as the desired total sample size. We then calculate the sampling interval (or household selection weight) as the total number of addresses in a PSU divided by the number of selected addresses in the PSU.

A simple example is presented below, where there are 3 PSUs and the desired total sample size is 35 addresses. PSU number 1 does not have enough addresses to meet the implied number of addresses per PSU (35/3=11.67), assuming 5 panels are needed per PSU. Therefore, the sampling interval in PSU number 1 is set to 5, with 10 addresses per panel, and the remaining number of desired addresses (35-10=25) sets the new implied number of addresses per PSU over the two remaining PSUs (25/2=12.5). Using this new implied number of addresses per PSU, PSU number 2 does not have enough addresses (again, assuming 5 panels are needed per PSU), and so the sampling interval in this PSU is also set to 5, with 12 addresses per panel. The new implied number of addresses for the final PSU is (35-10-12=13/1=13) and the third PSU is large enough to accommodate 5.4 panels of size 13.

Appendix table 18

Simple example of number of addresses selected per PSU

PSU number	**PSU size (no. addresses)**	**Sampling interval (K)**	**No. addresses per panel**	**New sampling interval (iteration 1)**	**New no. addresses per panel (iteration 1)**	**New sampling interval (iteration 2)**	**New no. addresses per panel (iteration 2)**
1	50	5	10	5	10	5	10
2	60	5.14	11.67	5	12	5	12
3	70	6	11.67	5.6	12.5	5.4	13
Total sample size	…	…	33.33	…	34.5	…	35

Symbol: … not applicable
Source: Stats NZ

For the old sample, we calculated the sampling interval as:

- 6.36 for PSUs with a size of less than 100
- the PSU size divided by 16 for PSUs with a size of 100 or more.

That is, for PSUs with a size of less than 100, every 6.36th household was selected for one panel. For larger PSUs, panels of size 16 were used, resulting in a sampling interval of the PSU size divided by 16.

Calibration

The calibration methodology we employ is generalised regression (GREG), which we implement using GREGWT. GREGWT is a SAS macro, which was developed by the Australian Bureau of Statistics. The GREG estimator is derived as follows:

We want to adjust the initial weights a_i (selection weights multiplied by any non-response adjustment) to get new weights $w_i = a_i g_i$ that meet our benchmark constraints:

$$\sum_{i \in p} w_i \, x_i = t_{x_i}$$

where x_i is an auxiliary variable available for all sampled units i, and the total of this variable for the population (t_{x_i}) is known. We want the new weights to be as close to the initial weights as possible. The GREG estimator is the set of new weights that meet the benchmark constraints above while minimising the generalised least-squares distance function given by:

$$F^{GLS} = \sum_i c_i (w_i - a_i)^2 / a_i = \sum_i c_i a_i \, (g_i - 1)^2$$

Using this distance measure, there is an analytic solution to the minimisation problem, namely:

$$w_i = [1 + [(t_X - \hat{t}_X)' (X'WX)^{-1} X'W]_i] a_i$$

where $t_X$is the vector of population totals for the auxiliary variables, $\hat{t}_X$ is the estimate of these totals from the sample, X is the matrix of auxiliary variables, W is a diagonal (weight) matrix whose i,i element is a_i/c_i. Clearly the g-factor is the term in the outer brackets on the right-hand side. For more details refer to Deville and Sārndal (1992).

Weights over the transition

Periodically, we draw a new sample for the HLFS. The new sample does not replace the old sample immediately. Rather, we use the HLFS's rotating design to introduce the sample gradually, minimising disruption to estimates of change in the process. That is, one of the main reasons the HLFS employs a rotating design is to ensure a significant number of respondents are common in adjacent quarters, which results in more reliable estimates of quarterly movement.

One consequence of introducing a new sample is that during the period of transition there are effectively two independent samples in operation; let's call them A and B – A being the existing, or old, sample, and B being the new sample. Employing a ratio estimator, the probability of selection of a particular PSU in stratum h is approximated as:

$$PR(\text{PSU k selected}) = \frac{(\text{number of dwellings in PSUs sampled in stratum h})}{(\text{number of dwellings in stratum h})}$$

In the first quarter of the transition (for the most recent transition this is the December 2014 quarter), we retained roughly 7/8ths of the selected PSUs from sample A, so the numerator in this expression becomes smaller, while the denominator remains unchanged. Therefore, the design weights (being the inverse of the expression) should grow by a factor of roughly 8/7. (Note that if all PSUs were exactly the same size, then the factor increase would be exactly 8/7).

In contrast, in the first quarter of the transition we use only 1/8th of sample B, so the calculated design weights come out roughly eight times as large as the full-sample weights. Thus, if we naively combined samples A and B and summed the weights, we would get a figure roughly double the working-age population (assuming no non-response or undercoverage). To remedy this, we then take:

$$a * sum(A) + (1 - a) * sum(B), 0 < a < 1$$

The variance of this quantity is:

$$a^2 Var\big(sum(A)\big) + (1 - a)^2 Var\big(sum(B)\big) + 2a(1 - a)Cov\big(sum(A), sum(B)\big)$$

Let's instead assume we wish to calculate some total t, and we have as estimates t_A and t_B by summing the weights from samples A and B for people with the characteristic of interest. Let us further assume that:

$$t_A = p_A Y \; and \; t_B = p_B Y$$

where Y is a known population total and $p_.$ is an estimated proportion. Then (assuming zero covariance):

$$Var[at_A + (1 - a)t_B] \approx (aY)^2 p_A \frac{(1 - p_A)}{n_A} + \big((1 - a)Y\big)^2 p_B \frac{(1 - p_B)}{n_B}$$

Assume that p_A and p_B, being estimates of the same proportion, are roughly equivalent, then minimising (by differentiating with respect to a, and setting to zero, and taking another short cut by skipping the second-order conditions) the variance above reduces to solving the following:

$a/n_A = (1 - a)/n_B$

Which gives:

$$a = n_A/(n_A + n_B)$$

The complete A and B samples are chosen to be roughly the same size, so in the first quarter of transition we have:

$a\sim = 7/8$ in the first quarter, $a\sim = 6/8$ in the second quarter, etc.

That is, the selection weights for sample A grow by roughly 8/7 in the first quarter of transition, and then are multiplied by roughly (actually, exactly in practice) 7/8. Thus, the selection weights for sample A in the first quarter of transition will be nearly equal to the full-sample selection weights.

Note: This solution minimises variance for cross-sectional estimates. We also investigated the solution for minimising variance for change, but that was not worth implementing.

Appendix 4: Sample composition over the transition

During the previous transition that began in March 2004, we noted some unusual movements in the working-age population, by region. That said, movements that appeared unusual up until immediately after the transition appear less so now we have added more data to the ends of the series. Nevertheless, suggestions were put forward that we did not calculate the design weights over the transition in the manner described, so the method was independently programmed and we compared the results.

Results showed the method appears to have been applied correctly, and the unusual movements during the transition were likely to be due to differences in the composition of the selected PSUs.

To look at this over the current transition, we ran simulations to obtain sampling distributions of the old and new samples.

Results of the simulations show the old and new samples match the census within sampling error. The new sampling process yields estimates with considerably smaller sampling variance before calibration, though there is much less difference after calibration. Both before and after calibration the new sample shows a reasonable overcount of unemployed, and undercount in NILF (well inside sampling error).

The old sample shows an overcount across the board so all categories are pushed down. Either way, based on the simulation results, it would not be surprising to see the unemployed series drift up after introducing the new sample, and NILF drift down, all things being equal.

The caveats are that census labour force status doesn't necessarily agree well with HLFS labour force status, there is additional sampling variation in the dwelling selection within PSUs (this represents one of many possible panelling scenarios), and time has passed since the 2013 Census.

See Appendix table 19 for the results.

Appendix table 19

Simulation results

Employment status	Results				
	Census[1]	Without calibration		With calibration	
		Old	New	Old	New
Estimate					
Employed	1,927,996	1,953,709	1,909,668	1,923,488	1,930,500
Unemployed	144,484	147,235	147,997	144,523	148,487
Not in labour force	960,427	984,605	945,143	965,035	949,412
Unknown	146,795	147,566	148,375	146,657	151,303
Total	3,179,702	3,233,114	3,151,183	3,179,702	3,179,702
Relative sampling error (%)					
Employed	…	3.9	2.3	1.0	0.9
Unemployed	…	6.3	5.2	5.2	5.0
Not in labour force	…	4.1	2.9	1.7	1.5
Unknown	…	8.6	7.2	8.0	7.3
Total	…	3.6	2.0	…	…
Relative error in point estimates (%)					
Employed	…	1.3	-1.0	-0.2	0.1
Unemployed	…	1.9	2.4	0.0	2.8
Not in labour force	…	2.5	-1.6	0.5	-1.1
Unknown	…	0.5	1.1	-0.1	3.1
Total	…	1.7	-0.9	…	…

1. Estimates do not account for undercoverage and relate to the HLFS survey population only.
Symbol: … not applicable
Source: Stats NZ

It is conceivable that, all things being equal, the introduction of the new HLFS sample may yield an increase in both employment and unemployment, and a decrease in the number not in the labour force, post calibration. The differences will be fully realised only after the new sample is completely embedded (September 2016 quarter), and movements throughout the transition could be less predictable.

Appendix figures 1 to 3 consider a few key estimates, by rotation group, to determine whether there are outliers that may influence the survey outputs over the transition.

Appendix figure 1

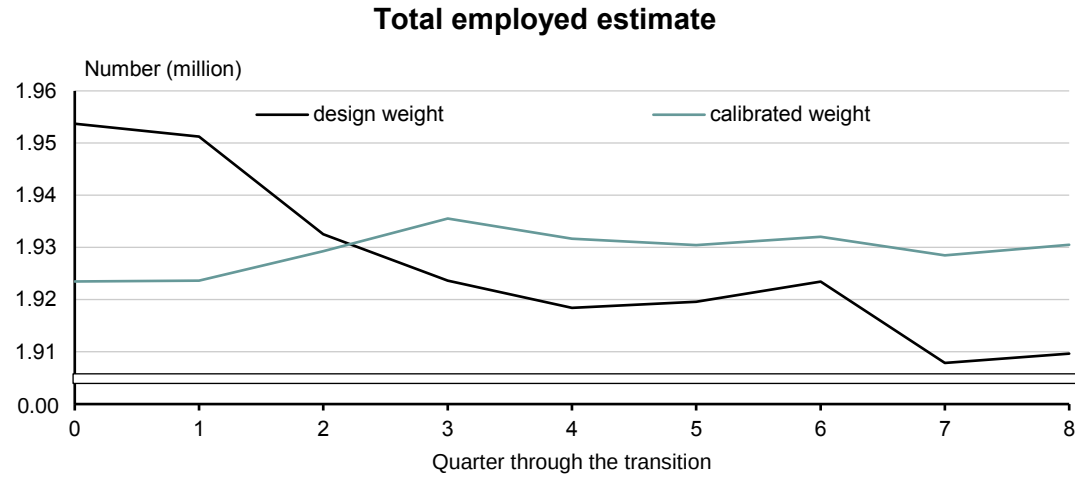

Source: Statistics New Zealand

Appendix figure 2

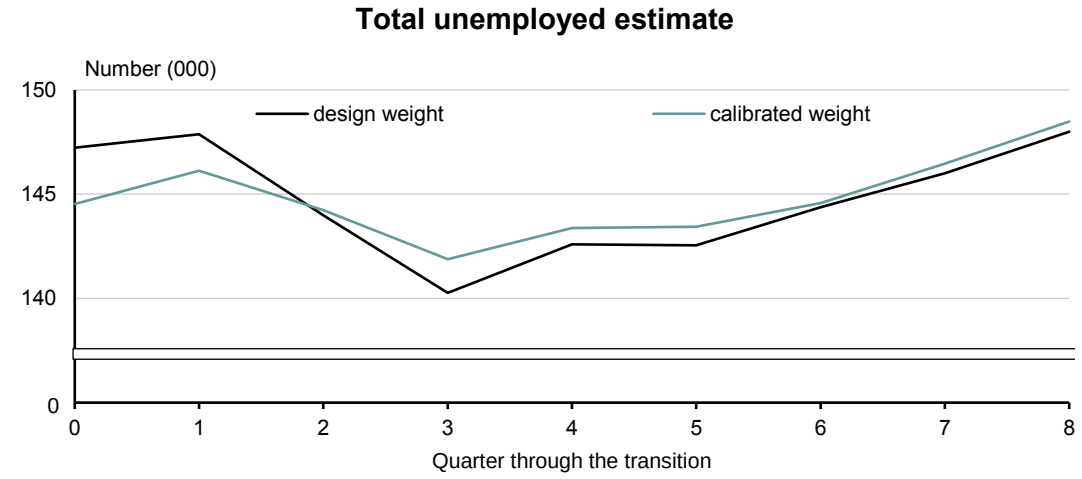

Source: Statistics New Zealand

Appendix figure 3

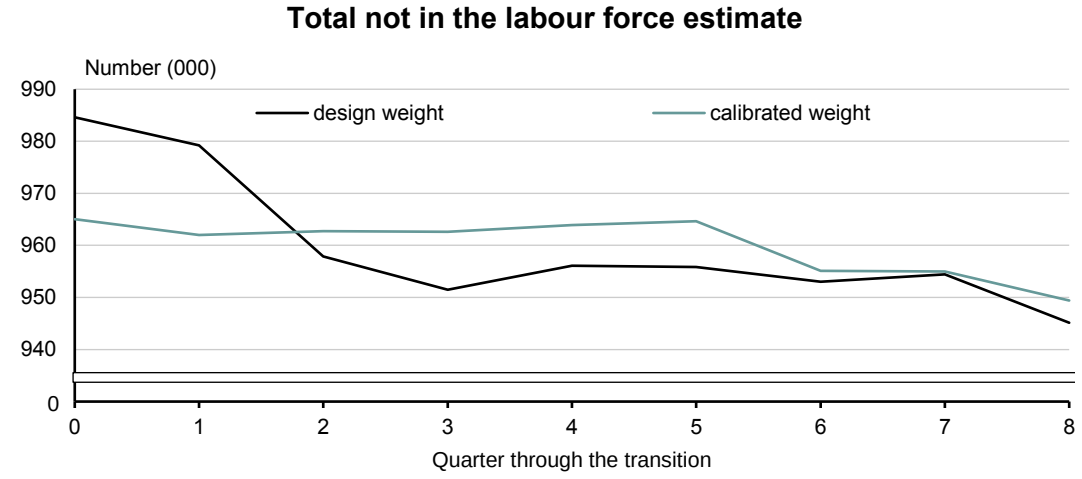

Source: Statistics New Zealand