Innovations in Measuring the Australian Labour Market

Sean Crick and Paul Schubert

Introduction

The ABS produces a comprehensive range of data on the Australian labour market. Statistics include core datasets such as the monthly Labour Force Survey, which measures employment, unemployment and labour force participation, business focused data such as Average weekly earnings and Job vacancies, as well as detailed data made available through the quarterly Labour accounts, focused topic releases and an increasingly wide range of microdata products.

Data has traditionally been obtained through surveys of households and businesses. While these surveys remain an integral part of measuring the labour market, they are increasingly difficult and costly to run, requiring new and innovative approaches to conducting surveys.

At the same time, large administrative datasets, such as the Single Touch Payroll dataset obtained from the Australian Taxation Office, offer opportunities for new and complementary statistics, and new techniques to improve traditional data collection methods.

This paper discusses three innovations the ABS has made to its Labour Statistics program in recent years:

- Modernising the design of the Labour Force Survey to improve response rates, respondent experience, and reduce costs
- Improvements to measuring regional labour market outcomes using administrative datasets
- Linking administrative data to survey frames to enhance sample selection, sample enumeration and estimation.

Corresponding author:
Sean Crick, Australian Bureau of Statistics, Locked Bag 10, Belconnen ACT 2616, Australia. Email: sean.crick@abs.gov.au

The authors have no financial or other conflict of interest with respect to the topics discussed in this paper. Both authors are employed by the Australian Bureau of Statistics.

Modernising the Australian Labour Force Survey

The Labour Force Survey (LFS) is the official source of employment statistics in Australia, providing estimates of employment, unemployment, participation and hours worked.

Like many national statistical agencies, the ABS is facing growing challenges in maintaining survey response rates, and managing the cost of data collection. Households are generally less willing to participate in surveys than in the past, and some more cautious about sharing information with government agencies. At the same time, many people now expect online forms to be simple, mobile-friendly and quick to complete, or to only need to provide information to government once.

LFS response rates have declined from consistently above 90% prior to the COVID period, to below 80% in recent months. At the same time, the proportion of households responding online has remained steady at just under 40%, with the reminder enumerated through telephone or face to face interview, a considerably more costly method of enumeration.

Chart 1 shows the decline of response rates over the past 11 years.

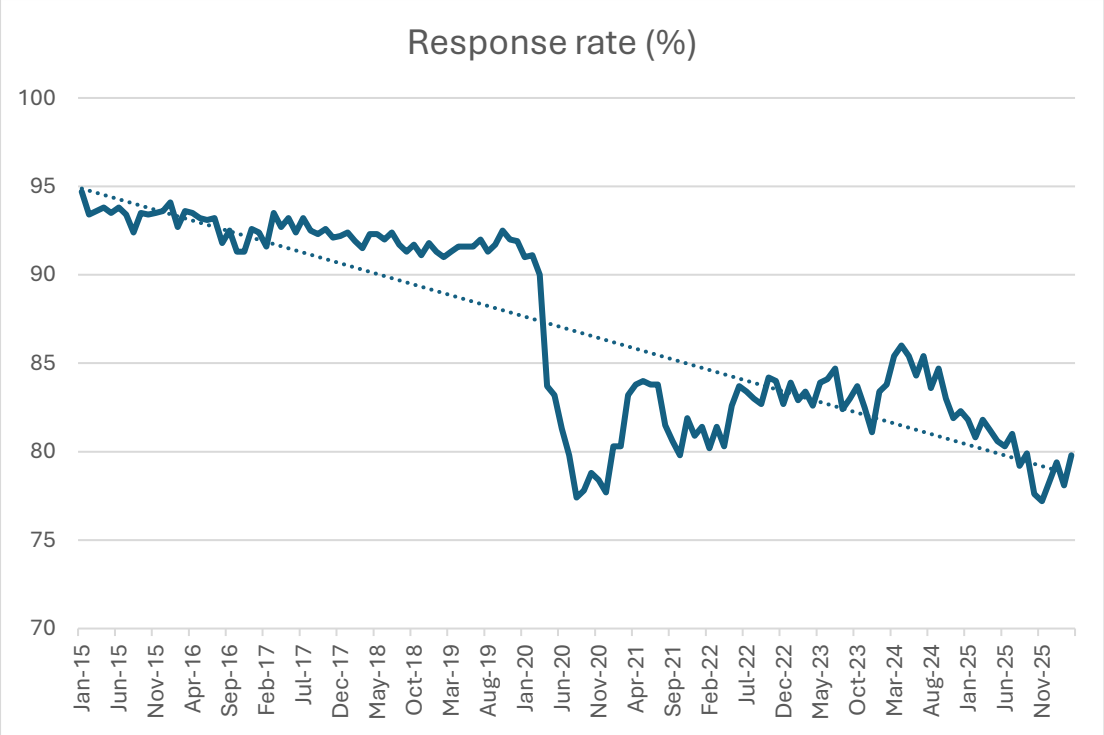


Chart 1 – Labour Force Survey Response Rates 2015 - 2026

Modernising the Labour Force Survey

In April 2026, the ABS introduced a new streamlined model for the LFS design (Australian Bureau of Statistics 2025). The aim is to:

- Ensure the ABS can continue to produce high quality, trusted labour force statistics
- Arrest the decline in response rates, and increase the proportion of households responding online
- Deliver an improved experience for survey respondents, and
- Improve efficiency.

The updated model simplifies both the survey design and how it is delivered. This includes:

- An improved digital experience, with participants able to complete the survey for the first time on a smartphone
- A streamlined and redesigned survey form. Households will answer only a minimum set of core labour force questions for their first seven months in the survey, with supplementary labour force topics asked only in their final eighth month of the survey. Duplicated and low value content has been removed.
- All parts of the survey will be delivered using modern, secure digital systems, while retaining telephone and other options for people who cannot respond online.

Appendix 1 shows the new and old survey design.

By making the survey easier and more convenient to complete, less confusing for respondents, and easier to operationalise, we expect to see improvements in response rates, increases in the proportion of households completing online, and reduced cost of managing enumeration.

The changes will lower the total time taken to fill in the survey, as households will no longer face multiple supplementary surveys spread across different months. For respondents, the survey will feel more predictable and straightforward.

Results to date

At time of writing, there are only limited results available from the new model. Only the incoming rotation group (one eighth of the sample, or around 4,000 households) have been enumerated on the new model and systems, and telephone and in-person follow up is still underway. However, there are some key trends coming through in the data.

Online take up

For the incoming rotation group on the new model, we have seen 46% of the sample respond via an online channel. This contrasts with an average of 35% for the incoming rotation group over the past 12 months.

Chart 2 shows the distribution of device used by those households who responded online. Just over one third used a smartphone or tablet, options that previously were not available to survey respondents.

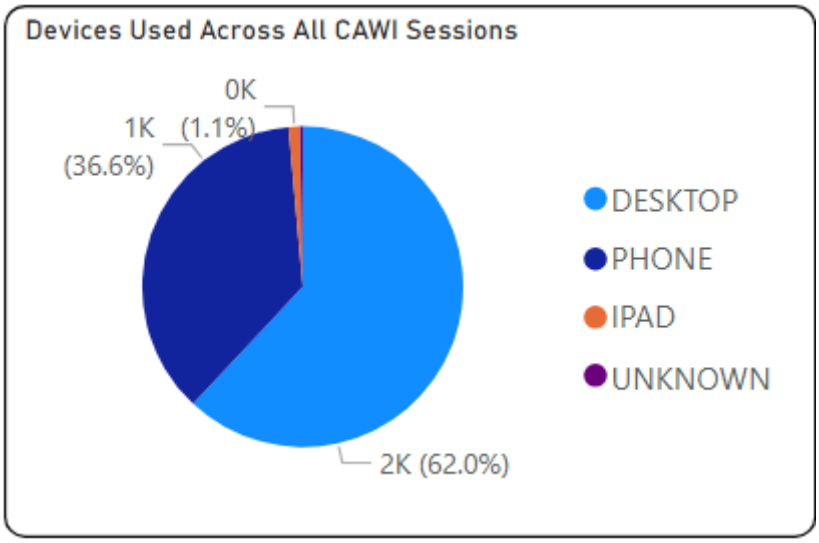


Chart 2 – Distribution of device used to respond online to the LFS, April 2026 incoming rotation group

These initial results demonstrate the importance and value of a simple and convenient survey platform in driving self-response through online channels. Further work, through improved communications and engagement strategies, are planned to keep those respondents engaged and reporting online into their subsequent months in the survey.

Survey length

While data is not yet available on the length of time that households will spend completing the LFS, a field test of the new model undertaken in October 2025 provided guidance on likely timing.

Under the existing model, the average time to complete the survey for a month with only core labour force content was 7.8 minutes, with the average over the 8 months in the survey 95.5 minutes.

Based on the 2025 test, this timing reduced to 5.9 minutes for the survey in months 1-7, and 77 minutes overall (including induction and supplementary content in the final month).

Improvements to measuring regional labour market outcomes

The Labour Force Survey is designed to provide robust national estimates of employment and unemployment, with the secondary design objective of producing states and territory estimates. While the LFS is not designed to produce regional estimates, these have been possible as a by-product of the large sample size required to produce high quality national estimates.

Below state (of which there are 8 states and territories in Australia), the ABS defines the next finest level of geography as the Statistical Area Level 4 (SA4). There are 89 SA4s that cover the whole of Australia without gaps or overlaps, and these are designed primarily to represent regional labour markets.

As these regional LFS estimates are based on smaller sample sizes in each region, they are of a lower level of statistical quality compared to those produced at the national and state and territory levels. Smaller sample size means greater sampling variability.

Over time, fluctuations occurred across most of the regional labour force data, particularly in regions with smaller populations. These fluctuations are generally caused by sampling variability, though sometimes they are caused by actual local events, such as natural disasters.

The regional labour force estimates derived from the LFS have large relative standard errors (RSEs), generally in excess of the RSE cautionary threshold used by the ABS. The regional estimates with the highest volatility, and hence the largest RSEs, tended to be in areas with lower populations, such as outback portions of many states.

A new model for timely regional labour market data

In April 2024, the ABS introduced an enhanced, modelled set of regional labour market estimates for each SA4, combining LFS data and administrative datasets. For each SA4, the following estimates are modelled:

- employed persons
- unemployed persons
- persons not in the labour force (NILF)
- associated rates (employment-to-population ratio, unemployment rate and participation rate).

The ABS uses de-identified administrative data to model regional labour market estimates. The two key data sources used are:

- Single Touch Payroll (STP) data from the Australian Tax Office, and

- JobSeeker and Youth Allowance recipients data from the Department of Social Services (DSS).

Single Touch Payroll is an Australian Taxation Office reporting system that requires employers to report payroll information for every payroll transaction that takes place, at the time that transaction takes place.

These administrative data sources contain granular geographic information and are regularly updated. The model produces more stable and reliable estimates, through leveraging the strong relationship between the administrative data and the survey data.

The model is based on a Rao-Yu (1994) area level time series model (Australian Bureau of Statistics 2023b). Estimates of the number of employed and unemployed people by SA4 are modelled separately:

- The unemployment model relies primarily on a time series relationship with the JobSeeker and Youth Allowance recipients data from DSS.
- The employment model relies primarily on a relationship with the STP data from the ATO, although it also makes some use of the JobSeeker and Youth Allowance recipients data from DSS.

Appendix B fully describes the model used.

To ensure coherence with LFS data, modelled estimates of the number of employed and unemployed people by SA4 are calibrated to the state and territory level estimates from the LFS.

Estimates of persons not in the labour force (NILF) are then derived as the residual of the SA4 population that is neither employed nor unemployed. Rates (unemployment, participation and employment to population) are then able to be calculated.

The models assume a roughly constant relationship between LFS variables and the auxiliary data over time. When there are government policy changes and/or unusual events such as COVID-19, these relationships may change significantly. Time series corrections have been applied to account for this, particularly COVID-19 effects on unemployment.

The modelling method explicitly considers the effects of seasonality and survey sampling:

- Residual seasonality was included in the unemployment model because of the different seasonal patterns that occur in unemployment relative to DSS JobSeeker recipients, particularly around the summer holidays when labour market activity is reduced. For employment, no residual seasonality was detected or required.

- The models also include a sampling error term, which is modelled such that the temporal correlations (and variances) induced by the rotating panel design of the LFS are captured. By modelling the direct survey estimates and their associated sampling error covariance structure the complex sample design of the LFS is properly accounted for. The models also include random effect terms for both areas and time points. These terms help to account for differences between areas and time points that the auxiliary data and seasonal terms could not capture.

Results of the model

Chart 3 presents a comparison of the direct survey estimates and modelled estimates of the unemployment rate for the Shepparton SA4, a regional area in the state of Victoria.

The direct survey estimates of the unemployment rate for Shepparton exhibit substantial month-to-month volatility, with some examples of an unemployment rate of 0% where no unemployed persons were observed in the labour force sample for a given month. In contrast, the modelled estimates produce a smoother and more reliable series, reducing the influence of short-term sampling variability. This same improvement is observed in the majority of SA4s.

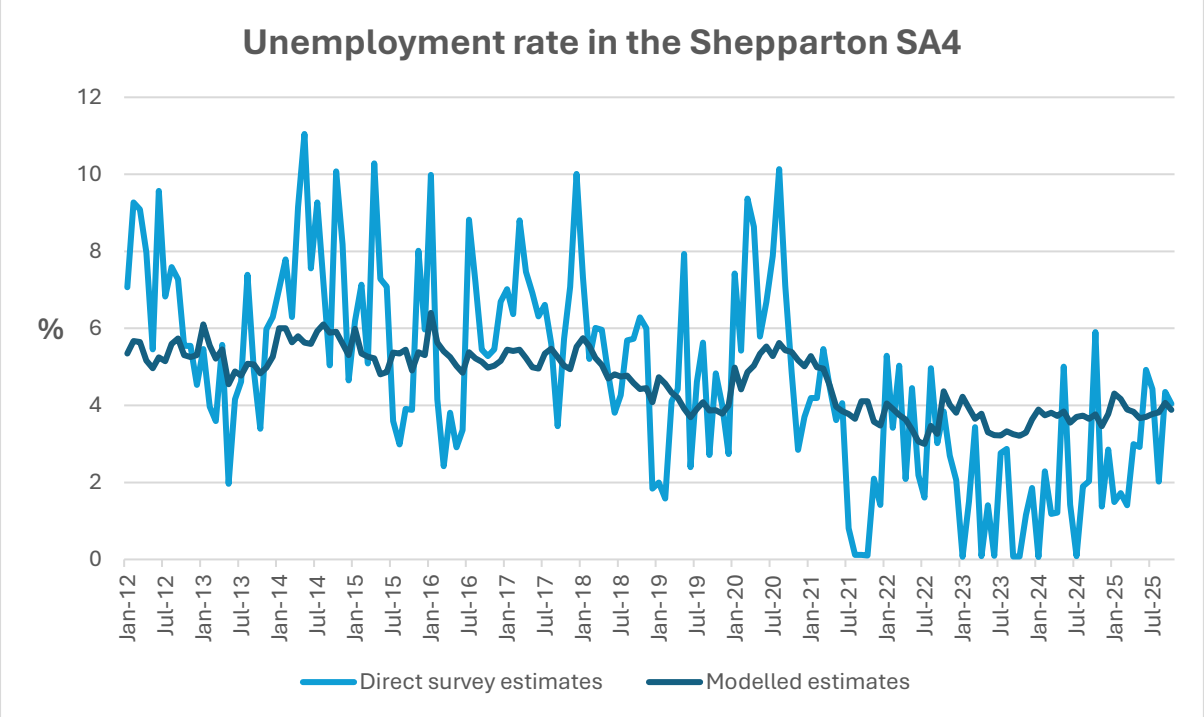


Chart 3 – direct vs modelled estimates of the unemployment rate, Shepparton SA4

The two charts below compare the relative root mean squared error (RRMSE) of the modelled estimates with the relative standard error (RSE) of the direct survey estimates for employment and unemployment levels of all SA4s. These graphs show that the RRMSEs of the modelled estimates are almost always smaller than their direct RSE counterparts. The improvement is particularly pronounced for unemployment, where direct survey RSEs are high for many SA4s due to the small number of unemployed persons captured in the survey sample.

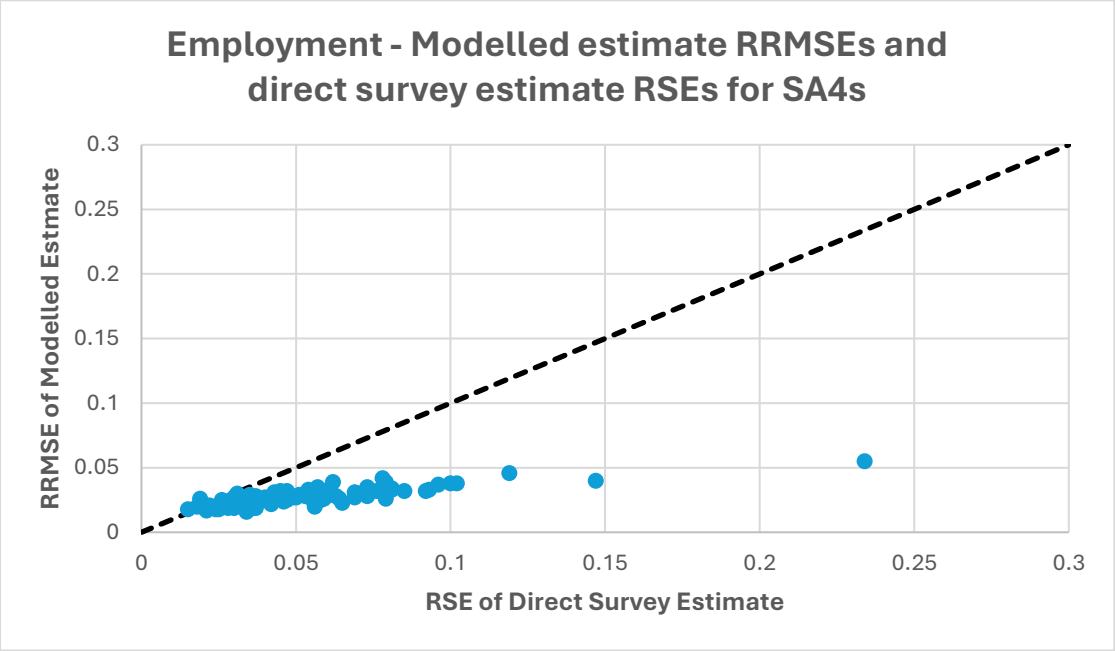


Chart 4 - Comparison of modelled RRMSEs and direct RSEs for employment for SA4s

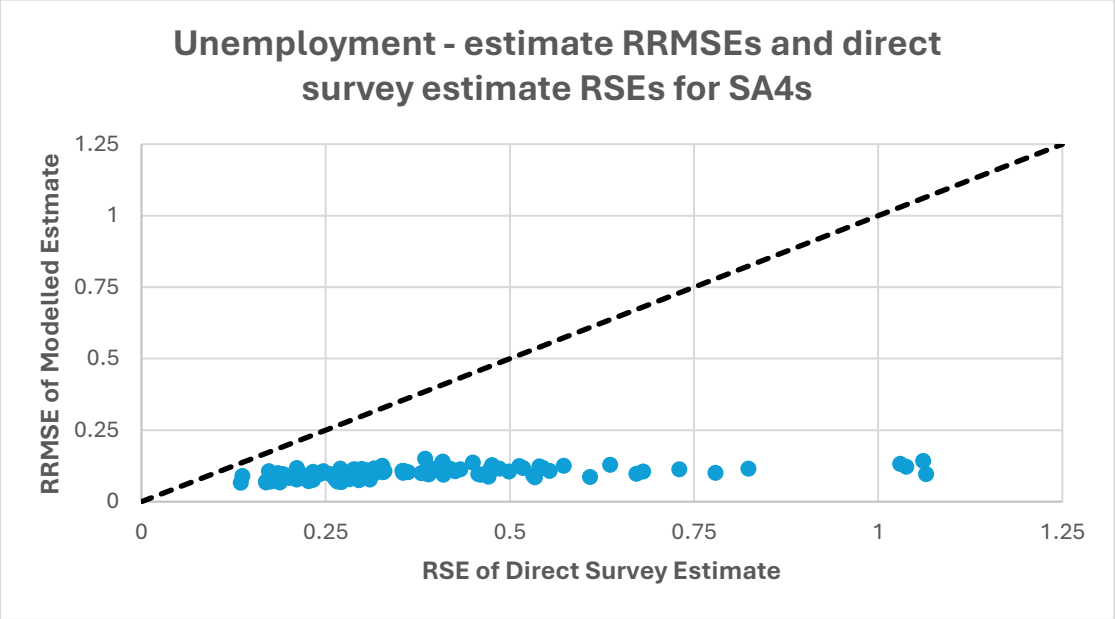


Chart 5 - Comparison of modelled RRMSEs and direct RSEs for unemployment for SA4s

The modelled estimates demonstrate a clear improvement in statistical quality at the SA4 level. This is evident in both the error measures and in the behaviour of the time series with significantly reduced volatility. These improvements are most noticeable in regions with smaller populations where direct estimates are subject to greater degrees of sampling variability.

ABS plans to further refine the model, with a longer-term goal to provide fit-for-purpose estimates at much finer levels of geography, and by age group and gender.

Using linked administrative data to enhance sample selection, sample enumeration and estimation

Australia does not have a population register to assist in conducting surveys of people or households.

The best available source of information from which to draw household survey frames is an Address Register (Australian Bureau of Statistics 2020). The Address Register is itself compiled from a variety of sources, but essentially each record contains only an address, the use of the address (e.g. residential or business) and type of dwelling structure information. Crucially, it does not contain any information about the people who live at the address.

Previous sample designs for household surveys, including the Labour Force Survey, have drawn on information at an area level, most commonly from the Census, to target samples towards areas with higher concentrations of subpopulations of interest. However, when it comes to selecting and surveying individual households, we have no information about the occupants before contact is made with the household (generally through a physical letter in the mail).

Increasing availability of administrative data opens up possibilities for knowing more about the likely occupants of households before sampling, thereby improving the efficiency and effectiveness of household survey estimation. This is important in addressing challenges of declining response across household surveys in Australia and worldwide, and in being able to predict and mitigate any non-response biases that may arise as we transition away from a reliance on face-to-face interviewing to a digital-first approach (eg eforms supplemented by telephone interviews).

The ABS hosts the Person Level Integrated Data Asset (PLIDA), a secure database combining mostly administrative information at a person (and household) level on health, education, government payments, income and taxation, employment, and population demographics (including the Census) over time (Australian Bureau of Statistics 2023). This opens up new possibilities for reducing non-response biases in ABS household survey estimates.

The ABS has developed the Data Improved Frames from the Address Register (DIFFAR) to achieve these goals. The DIFFAR builds on the Address Register by linking administrative data from PLIDA, to create DIFFAR.

<u>Augmenting with auxiliary data to create Data Improved Frames from the Address Register</u>

The administrative data used in DIFFAR is held in PLIDA which contains residential addresses of the people which need to be standardised and geocoded in much the same way as the Address Register so that they can be linked by address.

We calculate various propensities based on models which have been developed using linked historical survey data and PLIDA. One of the most important propensities is the propensity of a household to respond via an online form, but other propensities relate to various characteristics of the household:

- propensity to receive a government pension or benefit
- propensity that the household includes at least one person with a disability
- propensity that it is a single-adult household
- propensity that it is a couple-only household
- propensity that the household includes at least one child aged 0-14
- propensity to be in the lowest decile when it comes to income of the household
- propensity to be in the lowest decile of a summary indicator of socio-economic status.

The propensities are then converted into deciles which are attached to each person (via the household) in the PLIDA dataset. Deciles are used for privacy purposes. There may be some loss in utility in using deciles rather than raw PLIDA data items, but this is expected to be small.

The PLIDA with propensities is then linked by address to the Address Register and the resulting Address Register enhanced with the propensities becomes the DIFFARR. Frames produced from the Address Register therefore include propensities that can be used to enhance sample selection, enumeration and estimation.

Using DIFFARR for sample selection

Many of the subpopulations captured by the propensities (eg low income households, households containing a disabled person) are known to have lower response rates to our surveys than the general population, particularly low online (ie self enumerated rather than interviewer assisted) response rates.

We can use DIFFARR in two ways in sample selection:

- if our survey aims to target particular subpopulations, we can give greater probabilities of selection to those households with higher propensities to contain people within those subpopulations
- we can oversample households with lower propensity to respond online so that they are not under-represented in the final responding sample.

In this latter case, the sample selection weights are adjusted by a factor which is the ratio of the target fully responding sample in stratum h to the expected number of responses to be achieved in stratum (estimated from the modelled response propensities) ie the initial selection weights become

$$w_{i,h} = \frac{n_h^t}{n_h^r} \times \frac{N}{n^t}$$

where

n_h^t is the number of fully responding persons (FRP) target from stratum h

n_h^r is the number of FRP achieved from stratum h

N is the Estimated Resident Population of in-scope persons

n^t is the FRP target across all strata.

Using DIFFARR for sample enumeration

Historically, the only dimension of response rates that have been monitored for ABS Household Surveys is geographic eg State/Territory, or metropolitan / ex-metropolitan. System limitations made it hard to see and monitor the data in real time, only in retrospect which is too late for use in enumeration.

Having the PLIDA data matched to the survey frame and sample helps us to monitor response by characteristics of household and people in it – this enables us to prioritise non-response follow-up towards groups where response is lagging. We have recently implemented this in our household survey enumeration monitoring.

In the future we could go further and use an adaptive design approach by changing our strategy for enumeration in advance based on household characteristics and likelihood of enumeration success. For instance, if we know a household is unlikely to respond online from the outset, we may go straight to a different enumeration mode such as a phone call or visit for this household. Alternatively, we could change the tone in our reminder correspondence accordingly – for example, push harder and more firmly for responses from households we know are more likely to be able to respond via an online form.

Using DIFFARR for estimation

Both the bias and variance of survey estimates can increase if we have unequal response rates over different portions of the sample.

Previously, we were very limited in the benchmarks we could use to adjust for representation of the population in the responding sample – typically, just demographics were available.

Some data items of interest would not necessarily be strongly correlated with these benchmarks, so if response rates dropped we would be exposed to risk of bias in survey estimates.

PLIDA can provide a much richer set of benchmarks to adjust for under-representation of various segments of the population eg low income or socio-economic households; government pension or benefit recipients.

Examples of use of DIFFARR

Three examples of the use of DIFFAR are presented, primarily using other ABS household surveys (as this is where DIFFAR has been applied to date) but with comments indicating relevance to and future use in the LFS.

6.1 Sample selection

The DIFFARR approach was used to boost sample selections in lower-response deciles in the 2024 Time Use Survey. An aim of the design was to achieve approximately 1000 diary days in each of the response deciles (ranging from low propensity to respond deciles, to high propensity to respond).

Response decile	Target number of diary days	Using boosting strategy		Without boosting strategy		Relative weights	
		% of diary day target achieved	Total diary days collected	% of diary day target achieved	Total diary days collected	Weight with boosting	Weight without boosting
All / Total	**10,038**	**179%**	**18,010**	**177%**	**17,809**	**1.00**	**1.00**
1	1,009	222%	2,244	116%	1,174	0.81	1.52
2	1,008	201%	2,028	138%	1,393	0.89	1.28
3	989	178%	1,763	145%	1,434	1.01	1.22
4	995	177%	1,765	163%	1,624	1.01	1.09
5	1,003	141%	1,414	139%	1,395	1.27	1.28
6	988	133%	1,314	137%	1,349	1.35	1.30
7	976	170%	1,658	181%	1,769	1.06	0.98
8	1,014	199%	2,020	232%	2,350	0.90	0.77
9	1,011	176%	1,783	230%	2,324	1.02	0.77
10	1,045	193%	2,021	287%	2,998	0.93	0.62
Range		**89pp**	**930**	**171pp**	**1,824**	**0.54**	**0.90**

Table 1: Effect of sample boosting in lower-response deciles in the 2024 Time Use Survey

Table 1 shows that we went slightly too far with the low-response decile boosting – we achieved most response in the lowest two response deciles, deciles 9 and 10 (which would have had the lowest response without boosting) and relatively not quite enough response in the middle response deciles. However, the response with boosting is much

more evenly distributed than it would have been without boosting, as reflected in the range of weights in the last two columns.

As an aside, this survey which only had online and telephone response options (no face-to face interviewing) achieved a uniformly higher level of response than we anticipated, partly due to the success of the online diary design.

DIFFAR has the potential to be improve the efficiency of a future sample redesign of the LFS in a similar way, for example by targeting the sample towards households that are more likely to contain people in labour force transition (people under 30, people in their 60s, people on unemployment benefits).

6.2 Sample enumeration

The 2025 General Social Survey also was conducted using Computer Assisted Web Interviewing (CAWI) and Computer Assisted Telephone Interviewing (CATI) only. It had a short enumeration period of only one month for the CAWI plus one month for CATI follow-up.

Midway through the first month of enumeration (May) it was clear that the households under most financial stress (decile 10), as modelled with the PLIDA data, were disproportionally non responders – by quite a lot compared to the remainder of the sample. At mid-May, only 57.1% of our target of 1,275 responding households had been achieved, a much lower rate than seen in all other deciles.

GSS Response by Financial Stress Decile						
		Mid-May		Final		
Financial stress decile	Target	Target Rate	Relative weight	Target Rate	Relative weight	Improvement
ALL	10,002	84.20%	1.00	132.7%	1.00	-
1	787	102.40%	0.82	151.6%	0.88	0.05
2	792	101.50%	0.83	149.2%	0.89	0.06
3	761	99.60%	0.85	150.4%	0.88	0.04
4	796	93.10%	0.90	142.7%	0.93	0.03
5	774	95.60%	0.88	145.2%	0.91	0.03
6	757	93.00%	0.91	145.8%	0.91	0.00
7	815	89.20%	0.94	141.2%	0.94	0.00
8	899	80.60%	1.04	132.2%	1.00	0.04
9	1,179	79.10%	1.06	129.8%	1.02	0.04
10	1,275	57.10%	1.47	101.4%	1.31	0.17
99	1,166	65.20%	1.29	105.1%	1.26	0.03

Table 2: Response in the General Social Survey by Financial Stress Decile[1]

[1] Decile '99' contains those households with unknown financial stress (no auxiliary data).

We were able to target these households in our follow-up to reduce the extent of under representation in the responding sample – see decile 10 in the table in the graph, which compares mid-enumeration status with end of enumeration status. The improvement from mid-enumeration was greatest for this decile and, while it was still the lowest response decile, it achieved the target, reaching a total of 1,293 responses. (whereas the other deciles exceeded it).

It would be challenging to adopt a similar approach for the Australian LFS, given that it is monthly and so has a very short enumeration period. However, it may be more feasible to employ DIFFAR-type information during enumeration in countries with a quarterly LFS that have a longer enumeration window.

6.3 Estimation

As we transition the sample onto the new redesigned LFS survey form, DIFFAR can be used to help detect and manage any resulting statistical impacts.

Characteristics from DIFFAR of the responding sample using the new survey form can be compared to those from the old form, to assess whether there are any significant differences. If so, they can be used to adjust estimates in order to ensure that the time series of data remain on a comparable basis eg by 'calibrating' the new form sample to be the same in characteristics distribution as the old form sample.

References

Australian Bureau of Statistics (2020) ABS Address Register, Users' Guide | Australian Bureau of Statistics, ABS Website, accessed 31 August 2025.

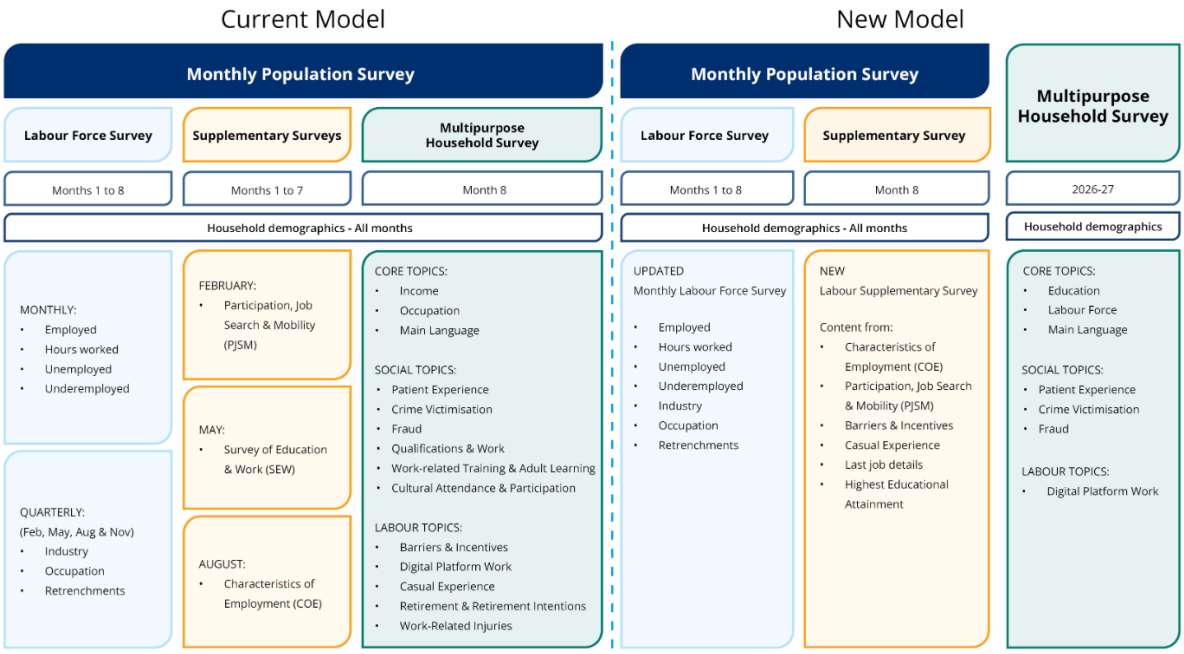
Australian Bureau of Statistics (2023a) Person Level Integrated Data Asset (PLIDA) | Australian Bureau of Statistics, ABS Website, accessed 31 August 2025.

Australian Bureau of Statistics (2023b) A Rao-Yu model for small area estimation of labour force statistics | Australian Bureau of Statistics, ABS Website, accessed 15 April 2026.

Australian Bureau of Statistics (2025) Modernising the Labour Force Survey | Australian Bureau of Statistics, ABS Website, accessed 15 April 2026.

Fraser, B, Schubert, P, Holmberg A, Reducing bias in household survey estimates using enhanced frames, 2026.

Appendix A – Current and New Labour Force Survey Model



Appendix B – Further detail on the model used for estimating employment and unemployment in the regional labour market

A Rao-Yu model has been chosen and our application of it is described in this section.

The key target variables are employment, unemployment and not-in-the-labour-force (NILF) counts along with the unemployment rate. It is reasonable to assume that the original LF series are multiplicative in nature with variances increasing with level. Hence a log transform is first applied for linearity and additivity and with the expectation that the transformed variables will better satisfy model assumptions.

Employment and unemployment counts are modelled independently. NILF has not been modelled directly but is derived by subtraction of the employment and unemployment modelled estimates from the LF total in-scope population. Similarly, the unemployment rate is derived from the modelled employment and unemployment counts. For employment and unemployment the model is described below.

$$\theta^*_{dt} = \mathbf{x'}_{dt}\boldsymbol{\beta} + v_d + u_{dt}$$
$$y^*_{dt} = \mathbf{x'}_{dt}\boldsymbol{\beta} + v_d + u_{dt} + e_{dt}$$

where:

θ^*_{dt} is the log of the population count θ_{dt} where $d = 1, 2, \ldots, D$ refers to the SA4 and $t = 1, 2, \ldots, T$ to the time.

y^*_{dt} is the log of the direct sample estimate of θ_{dt}

$\mathbf{x'}_{dt}$ is a vector of SA4-specific covariates, some of which may vary with t

v_d is a random SA4 effect with $v_d \sim NID(0, \sigma_v^2)$

u_{dt} is a random SA4-by-time effect modelled as a first-order autoregressive process:
$u_{dt} = \rho u_{d,t-1} + \epsilon_{dt},\ |\rho| < 1,\ \epsilon_{dt} \sim NID(0, \sigma_\epsilon^2)$

$\mathbf{e}^*_d = (e^*_{d1}, e^*_{d2}, \ldots, e^*_{dT})$ is a vector of logged sampling errors for SA4 d which is assumed to follow a $N(\mathbf{0}, \boldsymbol{\psi}_d)$ distribution, where the covariance matrix $\boldsymbol{\psi}_d$ is known.

The models were fitted to the following state groups:

- NSW with ACT
- VIC with TAS
- QLD
- SA
- WA
- NT with outback SA4s from other states – 315, 406 and 508

This was due to the small number of SA4s in the smaller states of ACT, NT and TAS. Models were fitted at state group levels in preference to an Australia level, (i) because many of the model parameters are expected to be state-specific, and (ii) because of the computational burden of fitting an Australia-level model.