

Investigation on the impact of the 2016 redevelopment on the Household Labour Force time series

Vinayak Anand-Kumar

Richard Penny

Mark Gordon

Abstract

The Household Labour Force Survey (HLFS) is a sample survey that provides indicators of labour supply. It provides New Zealand's official employment and unemployment statistics, and forms part of Stats NZ's quarterly labour market statistics release. Recently, the HLFS was redeveloped, introducing a number of changes to maintain its relevance as an indicator of labour supply. This paper will explore the work carried out to manage the data quality of the HLFS estimates as these changes were being applied.

One of the more pertinent changes from the introduction of the new HLFS was to the estimates of people employed in New Zealand, with evidence that there was a level shift in the employment series driven more by survey changes as opposed to real changes in the labour market. As part of our investigations on the possible effects of these changes, we fitted time series models with an exogenous variable corresponding to the changes, using both ARIMA and a State Space Modelling for the time series component. Both approaches gave similar results, and we conclude by reporting on how these analyses are used in our decision making.

Vinayak Anand-Kumar

Email: vinayak.anand-kumar@stats.govt.nz

Richard Penny

Email: Richard.penny@stats.govt.nz

1 Introduction

This paper presents the work carried out to investigate potential level shifts in the Household Labour Force Survey time series at June 2016. The introduction is structured so that it presents:

1. The Household Labour Force Survey (HLFS) and the changes that were applied to it in June 2016
2. An overview of time series and intervention analysis
3. The application of intervention analysis to the HLFS time series

1.1 Household Labour Force Survey

The Household Labour Force Survey (HLFS) is a quarterly sample survey that provides an indication of labour supply (i.e. the supply of people into the labour market) in New Zealand. Some of its key measures include how many people are currently working and how many people are willing to work. It provides New Zealand's official employment and unemployment statistics, and is part of Stats NZ's quarterly labour market statistics release. Information from the HLFS is used to develop and monitor labour market and social policy, support research, and help inform on the health and general well-being of New Zealand's economy.

The use of the HLFS in decision making highlights the need for Stats NZ to manage the product's data quality. Aspects of data quality that are considered by the organization include (Eurostat, 2003):

- Relevance: Ensuring statistical concepts meet current and potential users' needs
- Timeliness: It is made available to customers, so that it might be used to inform decisions
- Consistency: Data are comparable across time and between other related data sources
- Accessibility: Data and information pertaining to the product are easily accessible
- Accuracy: It is the degree to which source data and statistical techniques used portray the reality that they were designed to represent
- Interpretability: Information about how the data are produced are available to customers

Most recently, the HLFS applied a number of changes to primarily improve its relevance as an indicator of labour supply, which were collectively referred to as the 2016 redevelopment. The new content provided more information about the nature of individual's employment conditions and work arrangements. As well as including more detail from respondents who work in more than one job during a reference week. Further details about the redevelopment are provided in the paper, Household Labour Force Survey - summary of 2016 redevelopment (Stats NZ, 2016). Other changes included:

- A new questionnaire, which more effectively identified
 - Those that were employed, in particular the self employed
 - The underutilized
- Removal of short form for respondents aged 65-74. Replacing it with the new questionnaire, which was better at picking up different types of employment, may have contributed to an increase in the number employed for this demographic
- Including armed forces residing in private dwellings into the survey
- Incorporating new/ modifying existing data processing steps into the survey that included donor imputation of usual and actual hours worked

Interventions intended to improve one of the six dimensions, can often have impacts on the other five. Although the 2016 redevelopment improved the product's relevance, and in some cases the accuracy of outputs, it had potentially negative consequences for the comparability of the data prior to and after June 2016. The redevelopment changed the way in which the data were collected and processed, and as a result the estimates along a time series were not derived using the same methods.

In the project to redevelop the HLFS, Stats NZ did explore how the changes to the survey could be managed. A dual run was considered. However, the added respondent burden, logistics and cost involved meant that this was not a feasible option. Prior to introducing the changes, the organisation elected to monitor the key HLFS time series. In June 2016, Stats NZ did carry out preliminary analysis prior to the release and found no evidence for a large level shift. However, the evidence was not conclusive that there was not a level shift, which could affect time series analysis immediately before and after. Therefore, it was agreed that an assessment would be carried out after several quarters of the redeveloped HLFS data to determine if there had been any structural changes in key time series as a result of the redevelopment. This would allow for the effect to be estimated, which in turn would advise on whether an intervention would be appropriate.

1.2 Time series analysis

Presenting HLFS estimates over time enables customers to interpret movements in the labour market in the context of short and long term patterns. The purpose of analysing time series is to 1) understand the movements that give rise to the observed series and 2) using that understanding to predict future values of a series. Analysts can use a combination of subject matter expertise and statistical tests to learn about the series; identifying the different types of movements that contribute to changes in the data and accounting for these so that customers of the data can make meaningful comparisons between quarters and across years.

The study of changes to the structure of a time series is referred to as intervention analysis. Interventions include any real world events or data derivation practices that cause abrupt changes to the trend of a time series. These exogenous factors can influence the time series for just one observation or have a prolonged impact on the time series. The latter is referred to as a level shift, whereby the mean of the time series is significantly different following a time point compared to before (Cryer & Chan, 2008).

1.3 Intervention analysis of HLFS time series

The change in the total employed time series between March and June 2016 was relatively large. Figure 1 below presents the total employed series, and the changes observed before and after the redevelopment. It is evident that this increase is more pronounced in the seasonally adjusted series. The June 2016 figure was adjusted up to account for the expected March to June seasonal changes in employment. Figure 2 demonstrates how the observed movements were not consistent with expected seasonal patterns; the March to June change was odd in relation to previous March to June movements. In seasonally adjusted terms, the total employed series had a 2.5% increase over the quarter, which was the largest quarterly change reported at the time of publication (see Figure 3).

The inclusion of the armed forces and the ability to more effectively detect those in self-employment were thought to contribute to the relatively large one-off change in the employed series. However an increase in total employed in June is not unknown (Figure 2). At the time, it was not conclusive as to whether the labour market series had a level shift; it was not known how much of the movement could be due to changes in the survey (i.e. level shift) and how much due to prevailing labour market

conditions. The following series were identified by the subject matter experts, as being most likely to have been impacted by the 2016 redevelopment, for the intervention analysis:

- Male employed
- Female employed
- Male Not in the labour force (NILF)
- Female NILF
- Total actual hours worked
- Total usual hours worked
- Employed full-time
- Employed part-time

The purpose of this paper is to present the methods and the results of the level shift investigation, carried out on the key HLFS time series. It was of interest to explore whether the mean of the time series had changed following the 2016 redevelopment. Structural changes of the time series were studied by fitting time series models, to explain the observed movements, and then incorporate an exogenous variable to estimate the shift in the mean function. The time series component was modelled as an ARIMA process (Box et al., 2008), and a State Space Model (Harvey, 1989) to corroborate the former. A common approach to identifying level shifts is that of Bai and Perron (1998) but in this instance, the time point for the level shift was known, and as a result, the intervention analysis approach of Box and Tiao (1975) was used.

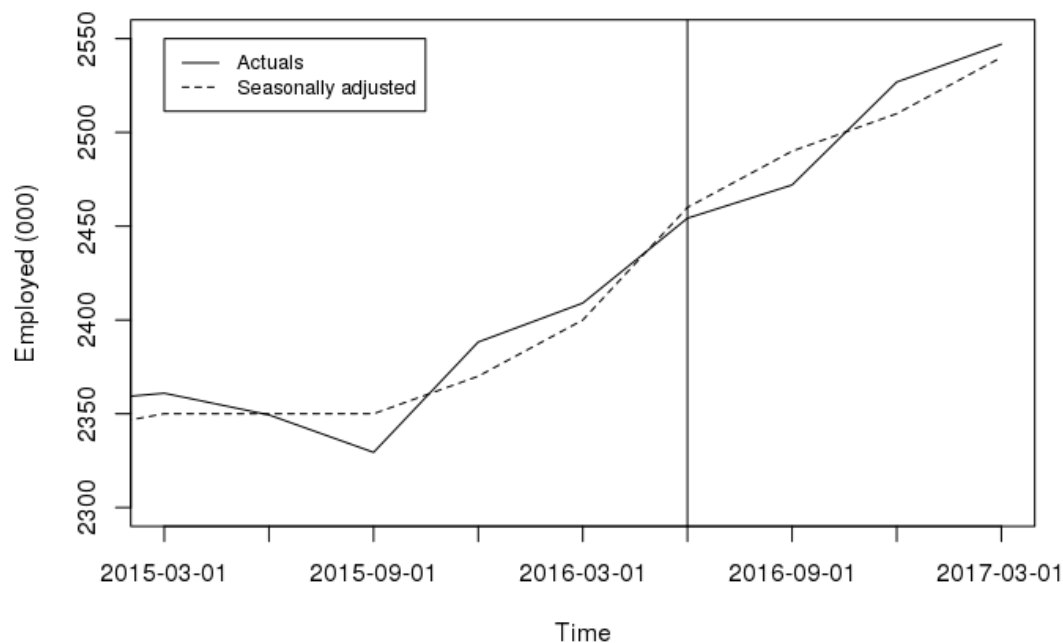


Figure 1. Time series plot presenting actual and seasonally adjusted figures for total employed between March 2015 and March 2017. The vertical line represents the start of the redeveloped HLFS.




Figure 2. March to June percentage changes in total employed between 1986 to 2016.




Figure 3. The distribution of quarterly percentage changes observed in the seasonally adjusted total employed series. The vertical line represents the quarterly change reported for the June 2016 quarter.

2 Methods

The intervention analysis of the HLFS time series involved:

1. Fitting a model to each time series.
2. Estimating the level shift using Regression with ARIMA errors (RegARIMA).
3. Estimates of the level shift were derived using State Space Modelling (SSM), to explore whether the results were similar to those from step (2).

2.1 Fitting a model

Cryer and Chan (2008) present three main steps in fitting a model to a time series:

1. Model specification
2. Model fitting
3. Model validation

2.1.1 Model specification

Model specification involves identifying the type of time series model, which would be appropriate for an observed series. Based on subject matter expertise, it was established that the labour market series would not have a constant mean. That is, it was expected that the long term and seasonal patterns of the series would vary over time. As a result, an ARIMA model was fitted to each time series.

ARIMA models are suitable for time series with stochastic (i.e. variable) trends. A key assumption required to make inferences about the structure of a time series is that there is a constant mean and variance over time. This is referred to as stationarity. The ARIMA model allows stochastic series to achieve a state of stationarity through a process called differencing. Observation of autocorrelation (ACF) plots helped identify the degree of differencing required.

Once the series is stationary, a model for the time series is fitted, which utilises previous figures and the historic error terms to explain the observed movements. These are referred to as autoregressive (AR) and moving average (MA) terms respectively. A seasonal ARIMA model is expressed as $(p,d,q)(P,D,Q)_4$, which reflects the degree of differencing, the order of AR and MA terms used.

2.1.2 Model fitting

The following steps were used to fit an ARIMA model to each time series (Duke University, 2014):

1. Identify the order of differencing:
 - a. The correct amount of differencing is the lowest order of differencing that yields a time series that fluctuates around a well-defined mean value and whose autocorrelation function (ACF) plot decays rapidly to zero.
 - b. If there is a spike at every 4th lag, then apply one order of seasonal differencing
2. Identify the AR and MA terms:
 - a. If the partial autocorrelation function (PACF) of the differenced series displays a sharp cutoff and/or the lag-1 autocorrelation is positive, then consider adding one or more AR terms to the model. The lag beyond which the PACF cuts off is the indicated number of AR terms.
 - b. If the autocorrelation function (ACF) of the differenced series displays a sharp cutoff and/or the lag-1 autocorrelation is negative, then consider adding an MA term to the model. The lag beyond which the ACF cuts off is the indicated number of MA terms
 - c. If the autocorrelation of the appropriately differenced series is positive at lag s , where s is the number of periods in a season, then consider adding a seasonal AR term to the model.

- d. If the autocorrelation of the differenced series is negative at lag s , consider adding a seasonal MA term to the model. This is likely to occur if a seasonal difference has been used, which should be done if the data has a stable and logical seasonal pattern.

2.1.3 Model validation

After fitting a model to the time series, it was of interest to explore:

1. How robust the model was at explaining the movements of the entire time series? That is, had there been any changes in the data generating process over the 30 years, which would limit the ability of a single model to explain all the movements. In the case of the HLFS, the data generating process is the underlying phenomena that give rise to the movements observed in the labour market series. These include the motivations for individuals to enter/ leave the labour force, and to seek employment. To evaluate this, the ACF and PACF plots for the first 15 years of the series were compared with the second 15 years of the series
2. Whether the 2016 redevelopment resulted in changes to the data generating process? To explore this, the ARIMA model fitted to the entire series was compared with the ARIMA model fitted to the series excluding observations between March 2015 and 2017. It was expected if the underlying structure of the series had remained unchanged over the previous two years, then the same ARIMA model would be fitted to the two series.

Details of the models fitted for the eight series reported on are available from the corresponding authors (see footnote on title page).

2.2 Estimating the level shift

When estimating the effects of an exogenous variable on an ARIMA time series, often an ARMAX approach is used but for technical and practical reasons, the RegARIMA approach was used.

2.2.1 RegARIMA

Once a model was fitted, the next step was to incorporate the effects of the intervention (or exogenous factors) on the time series. The model error terms were studied, to determine whether these were unusual relative to normal behaviour, at around the period of the redevelopment.

For simplicity, the demonstration below, which illustrates how exogenous factors are accounted for in a time series, are made in relation to non-seasonal ARIMA models. Moreover, the model presented below is assuming the data has already been made stationary. As a result, it is expressed using only AR and MA terms. Otherwise expressed as ARMA(p,q).

Let the time series be denoted by y_1, y_2, \dots, y_n

The time series y_t can be defined as an ARMA(p,q) model with no covariates:

$$y_t = \phi_1 y_{(t-1)} + \dots + \phi_p y_{(t-p)} - \theta_1 z_{(t-1)} - \dots - \theta_q z_{(t-q)} + z_t$$

Auto-regressive terms are denoted by ϕ_p . Moving average terms are denoted by θ_q .

A short hand for expressing a time series model is:

$$y_t = n_t$$

where the ARMA(p,q) model is captured by the expression n_t .

Box and Tiao (1975) provide a framework for studying the impact of exogenous factors on a time series; proposing that the intervention impacts the data generating process by changing the mean function. In the case of a single intervention, the model could be expressed as:

$$y_t = m_t + n_t$$

where m_t is the change in the mean function, and n_t is an ARMA model. In the case of the HLFS time series, it was of interest to explore whether the function m_t was significantly different to zero from June 2016. That is, whether there was a change in the trend following the 2016 redevelopment.

The impact of the intervention on the time series was studied using a regression model with ARIMA errors (RegARIMA), which is defined as follows:

$$y_t = \beta x_t + n_t$$

where x_t is a covariate at time t and β is its coefficient. Prior to June 2016 x_t is assumed to be zero, and after (and including) June 2016 it assumed to be one (see Figure 4 for a graphical representation of a change in the covariates).




Figure 4. Figure demonstrating a permanent abrupt shift in a time series. The figure reflects the change in the covariate x_t after time t (in this instance it suspected to be is June 2016).

The RegARIMA model was fitted using the seasonal adjustment program X-13-ARIMA-SEATS, as this is the software used by Stats NZ for estimating other exogenous effects in time series (e.g. outliers, trading day effects etc).

2.2.2 State Space Time Series Modelling

When Stats NZ seasonally adjusts its time series it decomposes them into components (trend, seasonal, calendar and irregular) which is an Unobserved Components Model (UCM) or, more generally, a State Space Model (SSM) (Harvey 1989). While the SSM paradigm is not generally used in official statistics agencies, it was applied in this study; a State Space Model with exogenous terms was fitted to the series for comparison purposes.

The basic SSM is

$$y_t = \mu_t + \gamma_t + \epsilon_t$$

where y_t is the labour force series, μ_t is the level component, γ_t the seasonal component and ϵ_t is random error (or more correctly, a sequence of independently, identically distributed (i.i.d.) zero

mean Gaussian random variables $\epsilon_t \sim N(\theta, \sigma_\epsilon^2)$). Harvey (1989) extends the model in the level component by allowing for a stochastic trend

$$y_t = \mu_t + \gamma_t + \epsilon_t$$

$$\mu_t = \mu_{(t-1)} + \beta + \eta_t$$

The slope (the first difference) is β and η_t is again a sequence of i.i.d. zero mean Gaussian random variables, though with variance σ_η^2 , rather than σ_ϵ^2 .

More generally it is possible, and more realistic, to allow for not only the level but also the slope to change over time,

$$y_t = \mu_t + \gamma_t + \epsilon_t$$

$$\mu_t = \mu_{(t-1)} + \beta_{(t-1)} + \eta_t$$

$$\beta_t = \beta_{(t-1)} + \xi_t$$

which allows the trend level to stochastically change as well as the change in the change of the trend to be stochastic over time. The final terms in both equations are estimates of uncertainty, though with different variance estimates. It is easy to extend the second equation to allow for exogenous terms, both regression and autoregressive.

$$y_t = \mu_t + \gamma_t + \sum_{i=1}^m \beta_i x_{it} + \epsilon_t$$

The exogenous regression term is the same structure as used in the RegARIMA approach of the previous section. The SSM was fitted using the UCM procedure in SAS.

3 Results

The results for each series are provided below, including: 1) The model fitted to the series 2) The estimates of the level shift with the corresponding confidence interval

3.1 Models fitted

As can be seen from Table 1 the best ARIMA model for all the series examined is the same, except for "Female, not in the labour force" (FNLF). The similarity of the models is not unexpected given the relationship between the series. The reason for the FNILF ARIMA model being different is being explored.

Table 1. Table presenting the ARIMA models used for each time series.

Series	Model
Male employed	(0,1,0)(0,1,1)
Female employed	(0,1,0)(0,1,1)
Male not in the labour force	(0,1,0)(0,1,1)
Female not in the labour force	(0,1,0)(1,0,0)
Part time employed	(0,1,0)(0,1,1)
Full time employed	(0,1,0)(0,1,1)
Usual hours worked	(0,1,0)(0,1,1)
Actual hours worked	(0,1,0)(0,1,1)

3.2 Estimating the level shift

The RegARIMA model and the SSM produced very similar estimates and confidence intervals for the level shift at June 2016. This provides support to our conclusion that the results are robust. For efficiency of exposition we only discuss the RegARIMA results, but information on the SSM outputs are available by contacting the corresponding author (see footnote on title page).

The findings were consistent with subject matter expectations. Earlier analyses that suggested the redevelopment was more effective at detecting those in employment, particularly self-employment, were reflected by the evidence of a level shift in the male, female and full time employed series, as well as usual hours worked. The argument that this was partly due to individuals previously being captured as "not in the labour force" (NILF) gained some support with evidence of a shift down in the male NILF series over the March to June 2016 quarter (Stats NZ, 2016; Stats NZ, 2016).

Table 2. Table presents the level shift estimates in 000s along with their confidence interval at 95%. The series marked with an asterisk presented with significant evidence of a level shift.

Series	Level shift estimates			Estimate at June 2016
	Lower	Estimate	Upper	
Male employed*	7	29	50	1296
Female employed*	0	21	42	1158
Male not in the labour force*	-30	-14	0	454
Female not in the labour force	-28	-10	7	680
Part time employed	-35	-10	15	527
Full time employed*	24	57	58	1927
Actual hours worked	-370	2550	5470	85005
Usual hours worked*	1880	3330	4781	92509

4 Discussion

In 2015, Stats NZ anticipated that the 2016 Household Labour Force Survey (HLFS) redevelopment may lead to structural changes in the key labour market series. For those series where data driven back casting was not an option, it was proposed that the redevelopment changes be managed using time series analyses. Furthermore, it was highlighted that any intervention of the time series be considered after sufficient data was available following the redevelopment. At least three data points were required (post intervention) in order to gauge whether there was a sustained level shift following the improvements to the collection and processing of HLFS data.

With four quarters of data generated using the redeveloped HLFS, an intervention analysis was carried out to estimate the size of the level shift on the key labour market series. The observed series were studied by fitting a time series model (i.e. the underlying data generating process) with a variable to estimate the possible change in the level of the series. Two different methods were used, which presented similar results.

The results of the intervention analysis revealed the following about the impact of the 2016 redevelopment on the labour market series:

- Male employed: there was evidence of a level shift increase
- Female employed: there was evidence of a level shift increase
- Male not in the labour force: there was evidence of a level shift decrease
- Female not in the labour force: there was no evidence of a level shift
- Part time employed: there was no evidence of a level shift
- Full time employed: there was evidence of a level shift increase
- Total actual hours: there was no evidence of a level shift
- Total usual hours: there was evidence of a level shift increase

Adjusting for the level shifts in the published labour market series has to be made in consideration for the impact any changes would have on data quality. The wide use of the HLFS in decision making by a range of customer groups highlights the need for any revisions to be associated with a degree of certainty that it is providing customers with higher quality data. Here quality does not merely reflect the accuracy of the estimate, but also includes its relevance, timeliness, consistency, accessibility and interpretability.

With consideration to the quality of the series, and the impact on customers, the analysts leading the investigation recommend that no adjustments be made to the key HLFs time series. The intervention analysis found that there was a large degree of uncertainty in the level shift estimates, and in turn, the uncertainty in whether applying an adjustment would definitively improve the accuracy and comparability of the time series. The uncertainty associated with the level shift estimates presents a challenge in making inferences about how much of the changes between the March and June 2016 quarters can be attributed to the re-development. Moreover, applying any adjustments to the time series prior to June 2016 would require further investigation into how the changes from the re-development may have impacted the labour market series at different time points. The time and resources required for this work, along with the risk appetite for implications on the accuracy and comparability of the time series should be weighed against the potential value added by this intervention.

References

Bai, J., & Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1), 47-78.

Box, G.E.P, Jenkins, G.M. and Reinsel, G.C. (2008), *Time Series Analysis: Forecasting and Control* (4th ed.), Wiley, New York

Box, G.E.P., & Tiao, G.C. (1975). Intervention analysis with applications to economic and environmental problems. *Journal of the American Statistical Association*, 70(349), 70-79.

Cryer, J., & Chan, K-S. (2008). *Time series analysis: With applications in R* (2nd ed.). New York: Springer.

Duke University (2014). ARIMA models for time series forecasting. Retrieved from <https://people.duke.edu/~rnau/arimrule.htm>

Eurostat. (2003). *Assessment of quality in statistics*. Retrieved 30 January, 2016, from <http://ec.europa.eu/eurostat/documents/64157/4373735/02-ESS-quality-definition.pdf/f0fdc8d8-6a9b-48e8-a636-9a34d073410f>

Harvey, A.C. (1989). *Forecasting, structural time series models and the Kalman filter*. London School of Economics and Political Science.

Stats NZ (2016). Household Labour Force Survey: Summary of 2016 redevelopment. Retrieved from http://www.stats.govt.nz/browse_for_stats/income-and-work/employment_and_unemployment/improving-labour-market-statistics/hlfs-summary-of-changes-2016.aspx

Stats NZ (2016). Employment status. Retrieved from http://www.stats.govt.nz/browse_for_stats/income-and-work/employment_and_unemployment/improving-labour-market-statistics/employment-status.aspx