A methodological note on leap year adjustments

This note explains how leap years might affect ONS time series and the methods used to adjust for them as part of seasonal adjustment.

Overview

What is a leap year?

A leap year adds an additional day to February generally every four years. It happens in years divisible by four, unless the year is a multiple of 100 but not 400, so 2000 was a leap year but 1900 was not.

How could a leap year affect ONS time series?

The occurrence of a leap year may affect movements in some ONS time series. For example consider a time series measuring output in a particular industry is used. If daily output in that industry is fairly constant, say 1 unit per day, then in a February without a leap year, say 2015, we would expect 28 units to be produced while in a February with a leap year, say 2016, we would expect 29 units to be produced.

This means that production in February has increase by 3.6% between February 2015 and February 2016 just because of the additional day, when in fact daily production has remained constant.

Are all ONS time series affected by leap years?

Not all ONS time series will be affected by the leap year, as it depends on what is being measured and the way in which it is measured. For example, leap year effects are found in Index of Production and Index of Services time series, which are estimates based on data collected from the Monthly Business Survey. However, a series such as the Retail Sales Index will not have leap year effects as the survey period is a four or five week period and not a calendar month. Moreover stock series, such as the Claimant Count are also unlikely to be effected by leap years due to the way in which the data are collected.

How are leap years accounted for?

Leap year effects, if present, are likely to be identifiable in monthly time series, and are routinely estimated and removed as part of the process of seasonal adjustment, where such an effect is statistically significant. Seasonal adjustment aims to remove effects due to the time of the year and arrangement of the calendar. It is difficult to robustly estimate a leap year effect in quarterly and annual time series and therefore these series are not adjusted for leap year effects.

Other months have different lengths, how is this different?

Apart from February other months contain either 30 or 31 days. Unlike February, the number of days in each month remains the same every year. With constant production of 1 unit per day we would expect 30 units to be produced in June and 31 in July so July would have higher total production even though the daily rate is the same. This is a seasonal effect as July would be expected to be higher than June every year. This is accounted for in seasonal adjustment.

Does the day of the week that February 29th falls on matter?

If daily output is 1 unit on a weekday and 0 on a weekend then in a February without a leap year there are 20 weekdays, so 20 units would be produced. In a February with a leap year, if February 29th falls on a weekday then 21 units would be produced, whereas if February 29th falls on a weekend then 20 units would be produced. This means that leap year effects can be related to trading day effects (effects due to the different number of different days of the week in a time period). Trading day effects are accounted for in seasonal adjustment and are closely related to leap year effects.

Details

This part of the note provides details on how leap year effects are accounted for. Section 1 provides some details on seasonal adjustment, including relevant sections of international guidance on seasonal adjustment regarding leap years. Section 2 discusses how different types of time series may be affected by leap years, with regard to what is being measured, at what frequency and the methods used for data collection and estimation. Section 3 provides examples of analysing time series with leap year effects.

1. Seasonal adjustment with leap year effects

Seasonal adjustment is the process of estimating and removing movements in time series associated with the time of year and other calendar related effects. ONS uses the X-11 method of seasonal adjustment as implemented in the software X-13ARIMA-SEATS. There are two main steps in the process of seasonal adjustment.

- RegARIMA, for calendar related adjustment.
- X-11, for seasonal adjustment.

RegARIMA modelling estimates and removes calendar related effects and outliers. The X-11 method decomposes this calendar and outlier adjusted series into trend, seasonal and irregular components. The seasonally adjusted series is then defined as the original series with calendar and seasonal effects removed.

Calendar effects, including leap year effects, are generally estimated in a regARIMA model such that

$$y_t = \sum_i \beta_i x_{it} + z_t$$

Where y_t is the observed time series, x_{it} are regression variables, β_i are the coefficients to be estimated in the regARIMA model and z_t follows an ARIMA model. For more information on the model and estimation of parameters see USCB (2013).

Leap year effects are related to length of month and trading day effects. ONS tests monthly series for trading effects using a regARIMA model. Where trading day effects are found to be statistically significant, an adjustment is usually made for leap year effects too. A series would not be adjusted for leap year effects if a trading day regressor without a leap year effect was found to be preferred.

A leap year effect is adjusted for in slightly different ways depending on whether it is believed that unobserved components of the time series combine multiplicatively or additively.

In the case of a multiplicative combination, a natural logarithm transformation is applied to the observed series and then modelled additively. In this instance a leap year regressor is not included explicitly as February values are adjusted prior to modelling using a fixed proportion. On average over a four year cycle the length of February is 28.25 days long. In non-leap years, output will therefore be about 0.9 percent lower than the average February, while in leap years output will be approximately 2.7 percent higher than the average February. These proportions are used to adjust February values onto an average February of 28.25 days.

Where series are assumed to be additive no transformation is applied to the observed series and a leap year variable is included in the regARIMA model to estimate the effect of the additional day in a leap year relative to non-leap years. This is the default approach to dealing with leap years in X-13ARIMA-SEATS and is based on research into alternative methods for dealing with length of month adjustments (Bell, 1992).

The leap year regressor (L_t) used in X-13ARIMA-SEATS is defined as

$$L_t = \begin{cases} 0.75 & Leap\ Year\ February\ (calendar\ quarter\ 1) \\ -0.25 & non-Leap\ Year\ February\ (calendar\ quarter\ 1) \\ 0 & for\ all\ other\ months\ (quarters) \end{cases}$$

ONS follows the European Statistical Systems Guidelines on Seasonal Adjustment (ESS, 2015), which provides recommended approaches to dealing with different issues in seasonal adjustment. Leap year effects are mentioned briefly in the ESS Guidelines on Seasonal Adjustment. The main section of relevance is section 2.3 on calendar adjustment, where the main recommendation is to use regARIMA modelling for calendar effects when there is a rationale for such effects in the series.

The use of proportional adjustment for leap years in multiplicative series may appear to go against the guidelines, as proportional adjustment is not recommended. However, as discussed in Bell (1992), there is equivalence in multiplicative series, where the expected coefficient for natural logarithm transformed data with leap year or de-seasonalised length of month regressors is equivalent to a proportional adjustment.

The main point to note is that calendar effects including leap year are tested for, and an adjustment is only made where found to be statistically significant.

2. Types of time series affected by leap years

Different types of time series may or may not be affected by leap years and it is important to consider the nature of the data when analysing a time series for leap year effects. Here we consider three categories

- Type of data
- Time considerations
- Data collection and estimation methods

Type of data

In the example used in section 2, our hypothetical time series was measuring the monthly aggregation of daily output; this is referred to as a flow series and such series may exhibit leap year and trading day effects as an additional day assuming equal activity on each day within the month of February will increase output by a fixed proportion. Examples of such time series in ONS include Index of Production or Index of Services, which are based on estimates of turnover in a month.

Stock series make a measurement at a particular point in time. For example, the Claimant Count is a measure of the number of people claiming particular benefits on the second Thursday of the month. These sorts of series are less likely to be effected by leap year effects, but some types of stock series may be affected. Such effects can be estimated using the regARIMA model but alternative regressors may be required.

Price indices, may have similar issues to stock series. For example, Consumer Price Indices use price data collected on the second or third Tuesday of the month and expenditure information is based on annual data. The extra day in February is unlikely to affect prices on the second or third Tuesday of the month. While there is potential for leap year effects to exist in annual data, they are unlikely to be estimated reliably, as discussed in section 2.5 of the ESS Guidelines on Seasonal Adjustment.

Time averaged stocks, such as unemployment estimates are also unlikely to be affected by the additional day. The Labour Force Survey asks respondents for their employment status over a particular week. It is a continuous survey and the headline figures are rolling quarters using data from a thirteen week period. The design of the survey and the shifting survey periods are discussed further below.

Time Considerations

The length and periodicity of data must also be considered when testing and estimating for leap year effects in a time series. Generally, ONS publishes, monthly, quarterly and annual time series. While leap year effects may in theory be present in all of these periodicities, in practice they are difficult to identify in quarterly and annual time series. For annual time series, the average length of a year is 365.25 days. Therefore in a leap year there are only 0.2 percent more days than the average year length and in non-leap years there are 0.06 percent fewer days than average. As discussed in section 2.5 of the ESS Guidelines on Seasonal Adjustment, estimation of leap year effects in annual data can also be confounded by business cycle effects.

Quarterly series are more ambiguous; there are 0.8 percent more days in the first calendar quarter of a leap year and 0.2 percent fewer days in the first calendar quarter of a leap year relative to the average first quarter length of 90.25 days. Research has suggested that trading day effects in quarterly time series may not be identifiable (see McDonald-Johnson et al, 2009). ONS does not make adjustments in quarterly time series for trading day effects and therefore does not adjust for leap year effects in quarterly time series.

In monthly series leap year effects are more readily identifiable. ONS tests for these effects in series where there is a rationale for such an effect.

The length of a time series is another important consideration in detecting and estimating leap year effects. For short time series (for example, monthly series less than six years in length) ONS does not make adjustments for trading day effects or leap year effects. Many monthly series where we would expect to see trading day and leap year effects, such as the Index of Production, or Index of Services, have relatively long time series. Trading day and leap year effects are assumed to be constant (in proportion or absolute terms) over the whole span of the time series due to the methods used for estimation.

Data collection and estimation methods

Some time series, for example from the Retail Sales Index, might be expected to exhibit leap year effects, as there may be an extra shopping day in February during a leap year. However, it is important to consider the methods used in collection and estimation of the time series. The Retail Sales Index is a monthly time series, but the survey periods over which data are collected are not strict calendar months. Respondents are requested to return data for either a four week or five week period and not the calendar month. Therefore with the estimation methods used, the monthly time series will not exhibit a leap year effect. A similar argument can be made for data collected from the Labour Force Survey, where an LFS monthly survey period is either a four or five week period rather than a strict calendar month.

3. Analysis of time series

In this section X-13ARIMA-SEATS is used to analyse two time series from the Index of Production. The first series discussed is the Manufacture of Petrochemicals (£millions, chained volume measure). Figure 1 shows two versions of the seasonally adjusted series for the Manufacture of Petrochemicals, one where leap year has been adjusted for and one where it has not. The vertical blue bars are positioned at leap year Februaries. As can be seen in the chart there are noticeable differences at leap year Februaries between the series where leap years have also been adjusted for (dark blue) and the series where they are not adjusted for (orange). There are smaller differences at non-leap year Februaries and even some very small differences in other months. The reason for the differences in months other than February is due to small differences in the estimated coefficients for trading day effects.

Figure 1 **Manufacture of Petrochemicals (£millions, chained volume measure) seasonally adjusted (dark blue) and seasonally adjusted with no leap year adjustment (orange)**

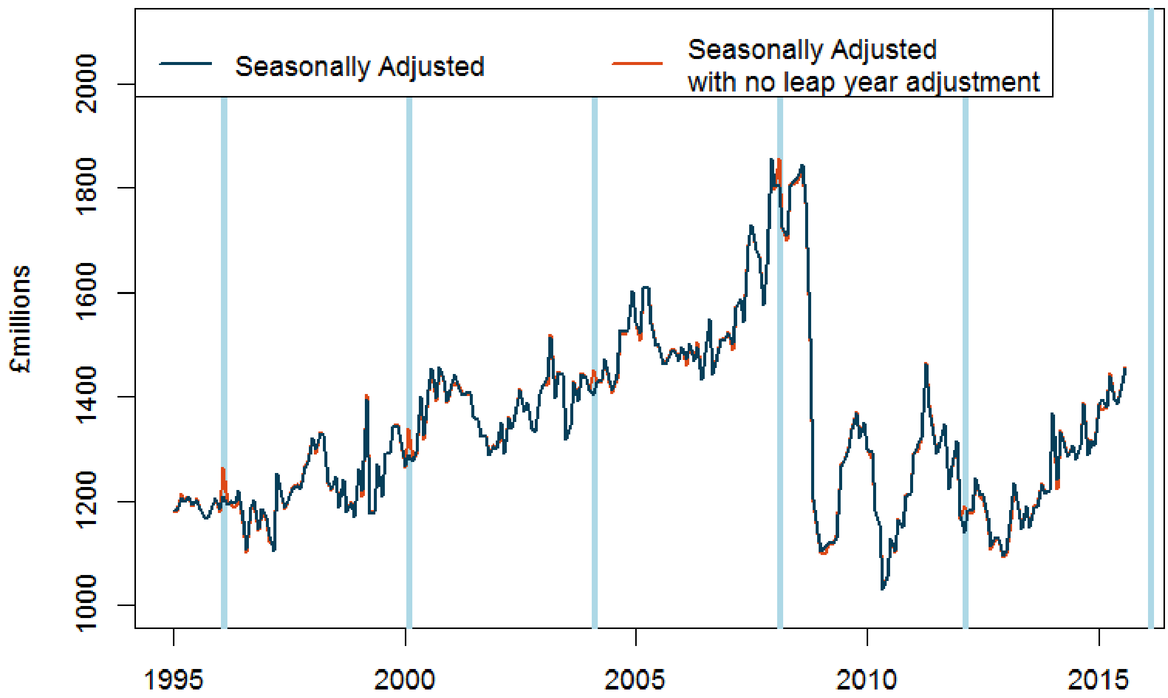

Note that there are a number of other effects accounted for in the model for this series, including a regressor to deal with the rapid change in level of the series from September to November 2008. The series is modelled without a transformation and therefore there is a leap year regressor explicitly included in the model. The estimated coefficient for the leap year regressor is 70.8 and is significant at the 5% level with a t-value of 3.2. This means that on average a leap year is expected to be about £53 million greater than the average level for February (everything else being equal) while a non-leap year will be about £18 million less than the average level for February. Note that there are also trading day effects and therefore the actual adjustment for leap year Februaries will depend on which day of the week is added to February in a particular leap year. Note also that all non-leap year Februaries will be adjusted by £18 million as there are an equal number of each day of the week in non-leap years therefore a zero trading day adjustment for these periods.

The second example presents seasonally adjusted series from Other Manufacturing. Figure 2 compares two versions of seasonally adjusted data, both of which have been adjusted for leap year effects but using different methods. As the modelled series is assumed to be multiplicative a natural logarithm transformation has been applied to the series. The method used to adjust for leap years is a fixed proportional adjustment as described in sections 2 and 3.

Figure 2 **Other Manufacturing (£millions, chained volume measure). Seasonally adjusted series with fixed proportional adjustment for leap years (dark blue) and with estimated leap year adjustment (orange).**

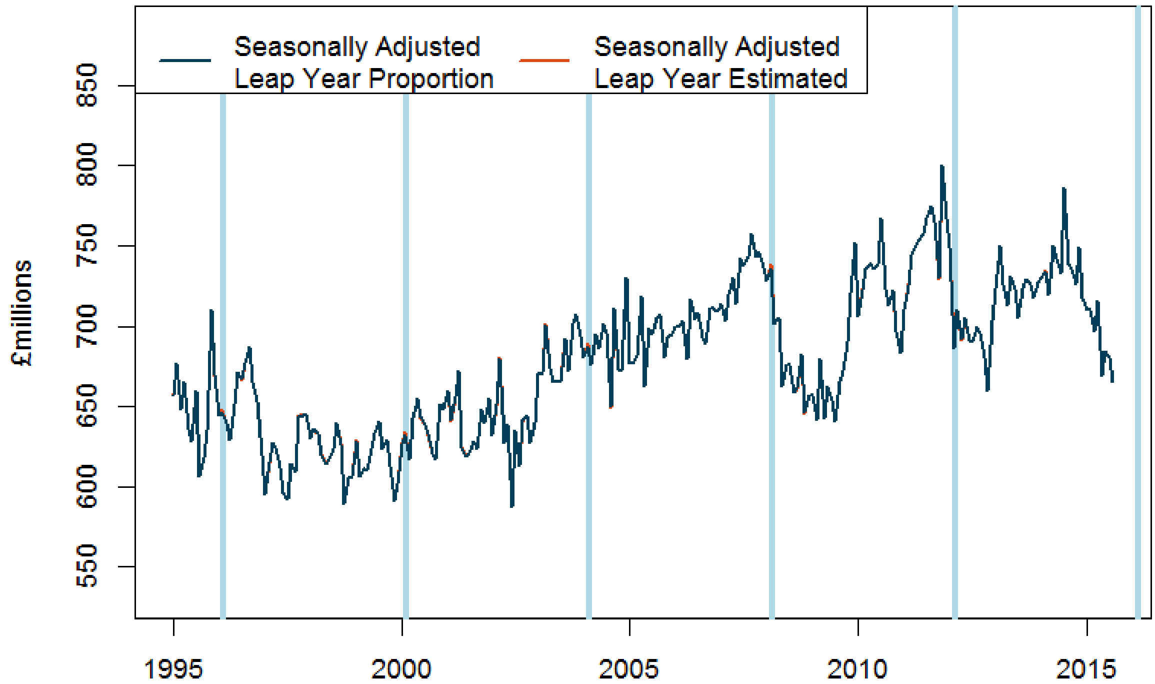

For the purpose of comparison, the second method estimates the coefficient for a leap year regressor. As can be seen there are negligible differences between the two methods used for adjustment. The reason for these negligible differences is due to the estimated coefficient of 0.031 (statistically significant at the 5% level with a t-value of 2.14). This is equivalent to expecting a 2.4 percent increase in leap year Februaries, whereas the fixed proportion adjusts for a 2.7 percent increase, which would be equivalent to a coefficient of approximately 0.035.

Further Information

For further information on points raised in this note please contact the Time Series Analysis Branch

tsab@ons.gov.uk

References

Bell, W. R. (1992) ‘Alternative Approaches to Length of Month Adjustment’ Research Report Series (Statistics #1992-17), Statistical Research Division, U.S. Census Bureau

ESS (2015), *ESS guidelines on seasonal adjustment*, 2015 edition, Eurostat, European Union 2015

McDonald-Johnson, K.M., Findley, D.F., Cepietz, E. (2009) ‘Investigating Quarterly Trading Day Effects’, U.S. Census Bureau https://www.census.gov/ts/papers/jsm09kmj.pdf

USCB (2013) *X-13ARIMA-SEATS Reference Manual*, Version 1.1, Time Series Research Staff, U.S Census Bureau, September 10 2013.