

An Empirical Study of Break Point Detection for Seasonal Change in an External Migration Series

Paper presented at Econometric Society of Australasia Meeting,
at Wellington, New Zealand,
July 2008

Guan Yu Zheng^{1}, Richard Penny¹*

Marco Reale², Easaw Chacko²

¹ Statistics New Zealand, NZ, ² University of Canterbury, NZ

* Statistics New Zealand
Private Bag 4741
Christchurch, New Zealand
info@stats.govt.nz
www.stats.govt.nz

Liability statement: Statistics New Zealand gives no warranty that the information or data supplied in this paper is error free. All care and diligence has been used, however, in processing, analysing and extracting information. Statistics New Zealand will not be liable for any loss or damage suffered by customers consequent upon the use directly, or indirectly, of the information in this paper.

Reproduction of material: Any table or other material published in this paper may be reproduced and published without further licence, provided that it does not purport to be published under government authority and that acknowledgement is made of this source.

Abstract

Statistics New Zealand is interested in detection of structural breaks in its time series as these may indicate major changes in the data generating process which could affect the consistency of its outputs over time. As part of its time series outputs Statistics New Zealand produces the original time series but also often produces a seasonally adjusted series and trend estimate from the original series. To produce these outputs Statistics New Zealand uses the U.S. Bureau of the Census product *X-12-ARIMA* (X-12) to decompose the original series into a set of unobserved components; trend-cycle, seasonal, and irregular. Each individual component also provides useful information for the interpretation in the behaviour of the series so it is important that each component is estimated correctly.

There is a large literature on the identification of breaks in trends or levels. However a component of major interest to Statistics New Zealand is the seasonal pattern, so Statistics New Zealand needs to ensure that structural breaks in the seasonal component are identified properly and efficiently. For example, the school term changed to a four term year from three term year in 1996, thus altering the dates of school holidays. This could be reflected as an abrupt change in the seasonal pattern of the travel behaviour of New Zealand residents holidaying overseas. At present, Statistics New Zealand identifies these breaks by visual inspection of the individual monthly or quarterly SI charts.

We have investigated the use of two methods to identify the structural breaks, that of Bai and Perron (BP), and Atheoretical Regression Trees (ART). BP produces optimal solutions but at the cost of lengthy computation for long series and is also conservative in its identification of a break. ART uses regression trees which involves a non-parametric method for fitting piece-wise constant functions to a time series. ART, in contrast to BP, tends to overidentify the number of possible breaks. For Statistics New Zealand this is less of an issue, as not identifying a break is more an issue than identifying spurious breaks.

In our work we applied BP and ART to the SI series for each month produced by X-12 from the NZ short-term departures series. The SI series consists of a combination of the seasonal and irregular components for that particular month. We have found that BP and ART produce broadly similar results. Both methods identified a break in the August SI series in 1996 that could reflect the change of school holidays. They also showed range of possible breaks in the other months. We have found that the minimum allowed length of the segment between breaks is crucial to where breaks are identified. Statistics New Zealand would prefer to find a break in the seasonal pattern within 3 years of its occurrence, so we compared the results specifying minimum segment lengths of 3 and 10 and find that the results are similar.

Key words: Official statistics; Seasonal adjustment; X-12-ARIMA; Structural break; Seasonality; Non-parametric; Parametric; Regression tree; Bai and Perron.

1. Introduction

Statistics New Zealand is interested in detection of structural breaks in its time series as these may indicate major changes in the data generating process. This could be caused by many factors, such as major changes in the business or social environment, or changes in the legal and administrative systems that produce the data. Any structural break could affect the consistency of Statistics New Zealand's time series. If the structural breaks can be identified soon after they occur some recognition of their effects can be made in Statistics New Zealand outputs.

As part of its time series outputs Statistics New Zealand provides the original time series but often a seasonally adjusted series and trend estimate from the original series. To produce these outputs Statistics New Zealand uses the U.S. Bureau of the Census product *X-12-ARIMA* (X-12) to decompose the original series into a set of unobserved components: trend-cycle; seasonal; and irregular. Each individual component also provides useful information for interpreting the behaviour of the series, so it is important that each component is estimated correctly.

There has been a large amount of research on the identification of breaks in trends or levels, and its application to a broad range of time series. However a component of major interest to Statistics New Zealand is the seasonal pattern, thus Statistics New Zealand needs to ensure that structural breaks in the seasonal component are identified properly and efficiently. If the seasonal pattern changes abruptly and this change is not modelled properly the seasonally adjusted series, produced by removing this incorrect model of the seasonal component from the actual series, will be misleading. Most likely the seasonally adjusted series will appear more volatile over the years immediately before and after the seasonal break. This is because the seasonal factors for the months affected by the break will not be well estimated over that period. While we may have prior knowledge of a expected break in the seasonal pattern of the series, it is also necessary to identify unexpected break in the seasonal pattern. We will term these breaks in the seasonal pattern *seasonal breaks*.

In the next section we show why identification of seasonal breaks is important and outline the current visual examination method used by Statistics New Zealand for this purpose. The X-12 outputs from a series where we have strong prior evidence of a seasonal break are presented in section 3. The following sections outline the Bai and Perron and Atheoretical Regression Tree approaches to identifying breaks and provide empirical results using these to identify seasonal breaks. Section 6 reports on the application of Atheoretical Regression Trees to two simulated series with seasonal breaks. Our conclusions are given in the final section.

2. Structural breaks

Detection of structural breaks is a topic of considerable interest in applied time series analysis. For time series models stationarity is a key assumption, which implies the mean, variance and trend do not evolve in time. However many time series in economics and finance do not hold to this assumption as the data generating process associated with the time series can change abruptly. That is, a structural break occurs. Hansen (2001) states that an undetected structural break can lead to three major problems in a time series analysis.

1. misinterpretation of time series model;
2. biased estimates;

3. less accurate forecasting.

Statistics New Zealand not only produces the time series as measured, but often provides a seasonal adjusted time series, and sometimes a trend series, to assist users in the interpretation and understanding of the behaviour of the measured series. To produce a seasonally adjusted series the seasonal cycle needs to be estimated. From this seasonal cycle the individual quarterly or monthly seasonal factors are taken, and then removed from the quarterly or monthly measured time series. This requires the seasonal cycle to be well estimated. However it is possible to have a structural break in this seasonal cycle. While there have been many investigations into structural breaks in levels or trends of time series, little work has been done on applying structural break methodology to the seasonal cycle.

To seasonally adjust time series Statistics New Zealand uses the U.S. Bureau of the Census product, *X-12 Variant of Census II Seasonal Adjustment Method*, more commonly referred to as X-12, to decompose the original time series into a set of unobserved components (Ladiray and Quenneville, 2001). An integral part of X-12 is its use of Henderson moving averages to generate some of those components.

A Henderson filter is derived by minimising the sum of squares of the third difference of the moving average series. It is equivalent to minimising the quantity $H = \sum (\nabla^3 q_i)^2$, where ∇ is the first difference operator, and $\{q_i\}$ are the weights or coefficients of the moving average (Henderson 1916, Doherty 2001). The advantage to using Henderson filters is that they provide a smoother output than standard moving average (eg 2 x 12 moving average) and eliminate irregular variations

While X-12 can estimate a range of components, for our exposition we will assume that X-12 decomposes the original time series into 3 unobserved components, the trend-cycle (C), seasonal (S) and irregular (I) components. We also assume that the series is measured monthly. Most work has been on structural breaks in the trends rather than in the seasonal pattern. It should be noted that a seasonal break may not affect the overall level of the series. Instead the annual pattern has suddenly and permanently been altered, thus it may not appear as a structural break in the trend.

X-12 initially estimates the trend and then removes this from the original series. This creates a SI component (ie combination of seasonal and irregular components), from which the individual monthly seasonal factors are estimated. Over a complete annual cycle these factors average to approximately 0. If a break is present in a particular SI series for a particular month, the estimates of the seasonal factors for that month will be distorted before and after the break. The three unobserved components must combine to the values of the original series, so an incorrect estimation of the seasonal component will affect the estimation of at least one of the other components, usually the irregular component. As the SI series for any month are estimates of the seasonal factors for that month for a particular year, the seasonal factors, and thus the seasonally adjusted series, for two to three years before and after the seasonal break will appear more volatile. At present Statistics New Zealand identifies seasonal breaks by visual inspection of the SI plots (figure 3).

3. New Zealand resident short-term departures

When any person arrives or departs from New Zealand they are required to fill in an arrival or departure card. Total arrivals and departures are categorised into three sub-series.

1. overseas visitors to New Zealand

2. permanent migration, based on an expected absence from New Zealand of more than 12 months, or staying in New Zealand for more than 12 months
3. New Zealand residents absent short-term from New Zealand, that is, an expected absence of less than 12 months.

The overseas visitor arrivals series was examined by Haywood and Randal (2005), who investigate possible breaks caused by 9/11, but we focus on New Zealand residents' short-term departures (STD). This series will be affected by a range of events and changed circumstances, but generally the effects of these changes are incremental and will not seriously affect Statistics New Zealand's outputs.

While there are many reasons why New Zealand residents travel overseas, it is expected that a major driver of the STD series arises from families holidaying overseas. This leads to an obvious hypothesis that the travel patterns of New Zealand families are heavily influenced by the timing of school holidays. Before 1996, New Zealand had a three-term school year, so school holidays occurred in December–January, May and August–September. In 1996 a four-term school year was introduced, thus holidays are now December–January, March–April, June–July and September–October. It is therefore of interest to compare the seasonal pattern of the STD series before and after 1996.

As noted above, if the seasonal pattern changes abruptly and this is not modelled properly the seasonally adjusted series produced by deseasonalising the actual series using this incorrect seasonal model will be misleading. The seasonally adjusted series will most likely appear more volatile over the years immediately before and after the seasonal breaks as the seasonal factors for the months affected by the break will not be well estimated over that period. For the STD series this can be seen in the increased volatility of the seasonally adjusted series (trend \times irregular) in figure 1. The volatility of adjusted series has become greater in the years around 1996, the year the timing of the school holidays changed.

To deal with this abrupt change in the seasonal pattern Statistics New Zealand splits the series at 1996 and adjusts the two resulting sub-series separately, thus producing a different seasonally adjusted series to that plotted in figure 1.

We have used X-12 to decompose 27 years of monthly STD data, from January 1980 to December 2006, into the three unobserved components and the outputs are shown in figure 2. Again this does not correspond to the decomposition as published by Statistics New Zealand as we have not split the series. The plot shows an increasing trend with potential structural changes and an evolving seasonal pattern. Clearly the amplitude and shape of the seasonal pattern has changed over time, and X-12 has modelled a slowly evolving seasonal pattern. The irregular component is weakly-stationary and homogenous, although some potential outliers appear in the series.

At present, Statistics New Zealand detects seasonal breaks by visually inspecting the SI plots (figure 3). There is an SI plot for each month, and the plot consists of the SI series, the combined seasonal and irregular components, and the seasonal factor estimated by applying a moving average to the SI series. A break in the August factor at 1996 is clearly apparent, but the expected mirror shift in either June, July or September from the change in school holiday dates is not obvious. This could be because the seasonal break is smaller than expected or the change in the seasonal pattern is more complex than initially hypothesised. This demonstrates that the current process is subjective and may be inconsistent, as the identification of a seasonal break depends on the particular analysts' background and knowledge.

Figure 1: Original (black solid) and seasonally adjusted (red, dash) series of NZ resident short term departures between 1994 and 1998. Grey line indicates January 1996. Note that the seasonally adjusted series is not that published by Statistics New Zealand.




Figure 2: X-12 Decomposition. Top graph is actual (black) and seasonally adjusted series (red). Middle graph is seasonal component. Bottom graph is the irregular component. Grey dashed line indicates January 1996. These series differ from those used and published by Statistics New Zealand.


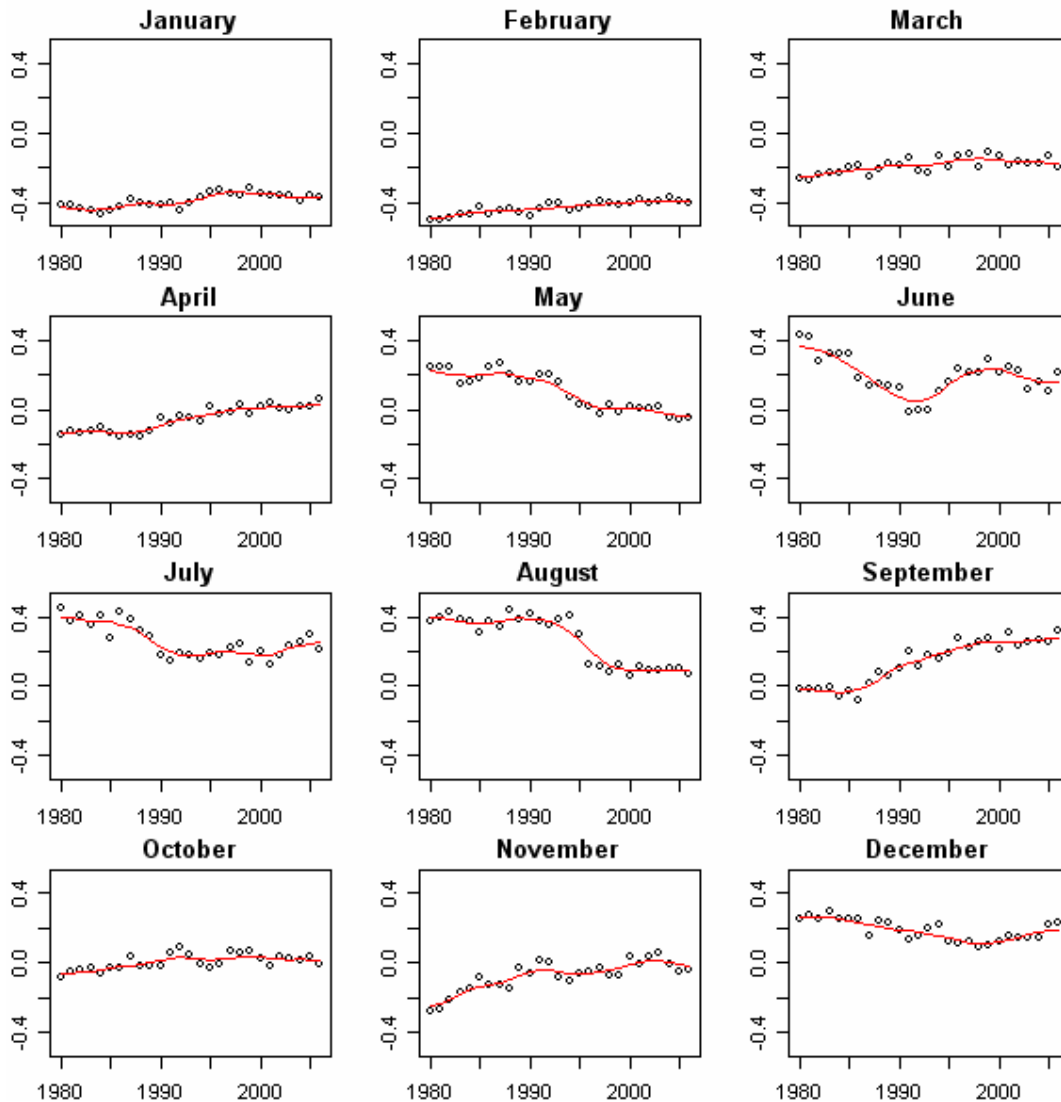


Figure 3: SI charts – SI values (black dots); Final Seasonal Factors (red lines)



4. Structural break methodology

Given the possible problems with visual identification of seasonal breaks we now introduce two methods for identifying structural breaks, that of Bai and Perron (1998, 2003) and Atheoretical Regression Trees (Cappelli and Reale, 2005).

Bai and Perron's method (BP) uses dynamic programming and Fisher's exact optimisation for fitting a regression model with structural breaks. Each serial point is allocated into mutually exclusive and exhaustive sub-series which maximise the sum of squared errors between groups and minimise the sum of squared errors within groups. This guarantees that the differences between adjacent sub-series are maximally distinguished. This methodology has these advantages: it is a global minimiser of the sum of squared errors given a known number of breaks, and it is a parametric procedure that provides statistical inferences for the parameters (eg date of breaks). However it is computationally intensive, an issue when series are long (eg > 1,000 points), and tends to underestimate the number of structural breaks in the time series.

An alternate method for detecting structural breaks is Atheoretical Regression Trees (ART) (Cappelli and Reale, 2005). It is a non-parametric approach, and produces

hierarchical tree structures with break dates. It makes use of a recursive binary division algorithm using Fishers' contiguous partitions method (1958) combined with a Least Square Regression Tree splitting criterion (Breiman et al, 1984). Under this approach, series are separated as far as possible until they reach the minimum size of sub-series/regime and no improvement on the splitting criterion. Usually ART will generate a maximal tree which is considered to be overfitted and is pruned by information criterion, like AIC and BIC. ART is computationally inexpensive and it is easy to interpret its outputs as it is non-parametric and allows inspection of the overall results as tree diagrams. However the splitting process does not guarantee to be optimal, and it tends to overestimate the number of breaks, especially for short series.

We estimate the seasonal factors of the 12 separate monthly SI series derived by X-12 from the STD series (figure 3) by fitting a structural break model to each of the SI series. We use the functions 'strucchange' (Zeileis et al, 2008) and 'tree' (Ripley, 2007) available in R to apply BP and ART respectively to the 12 STD SI series.

5. Results

As we are initially interested in whether BP and ART can identify the break in August 1996 we set the minimum regime length, the minimum length of the time series between breaks, as 10 years. We also asked BP and ART to identify the most likely single break point in the SI series for each of the 12 months (table 1).

Table 1: Date of structural Break identified by ART and BP. Minimum regime length 10.

Month	ART	BP	BP 95% confidence interval
January	1993.5	1994	[1993, 1996]
February	1990.5	1991	[1990, 1993]
March	1993.5	1994	[1991, 1999]
April	1989.5	1990	[1988, 1991]
May	1993.5	1994	[1993, 1995]
June	1989.5	1990	[1983, 2000]
July	1989.5	1990	[1989, 1992]
August	1995.5	1996	[1995, 1997]
September	1990.5	1991	[1990, 1992]
October	1990.5	1991	[1987, 1994]
November	1989.5	1990	[1988, 1994]
December	1989.5	1990	[1987, 1992]

Examination of table 1 shows that BP and ART produce very similar results in terms of identifying the year of a break in the SI series for all months. Both identify the break in 1996 for August; however they also show that the most likely break in the June, July or September factors is around 1990. This was unexpected, because if the effect of the change in school holidays is to shift the August holiday travel earlier to June–July school holiday or later to September this is not reflected in the overall seasonal break pattern identified by BP and ART. It is clear that there was an abrupt change in the August factor after 1995, so Statistics New Zealand is correct in splitting the series for seasonal adjustment, but the expected seasonal break model is not supported by our results. Of interest is the large number of months that have a break in 1989-90. This may be a random construct of the data as we have asked BP and ART to identify at *least* one break point and for each monthly series the size of the seasonal break will vary and may be, for any particular month, small. This possibility is further supported by

inspection of figure 3, as this shows the flatter, and thus more stable an SI series, the wider the confidence interval given by BP.

However it is important to remember that the individual monthly seasonal factors estimated from the SI series lead to an annual seasonal pattern. This seasonal pattern should, over any 12-month cycle, have the 12 monthly seasonal factors average approximately to 0, as most seasonal models assume seasonal patterns change slowly over time. A rapid change in the seasonal pattern could result from a set of small seasonal breaks in a number of months, which may be difficult to recognise by independently examining each month.

As Statistics New Zealand would like to identify a seasonal break as soon as possible after it occurs, a minimum regime length of 10 implies that the seasonal break would not be identified until 10 years have elapsed. We now repeat the work done above, but instead specify a minimum regime length of 3 years (table 2).

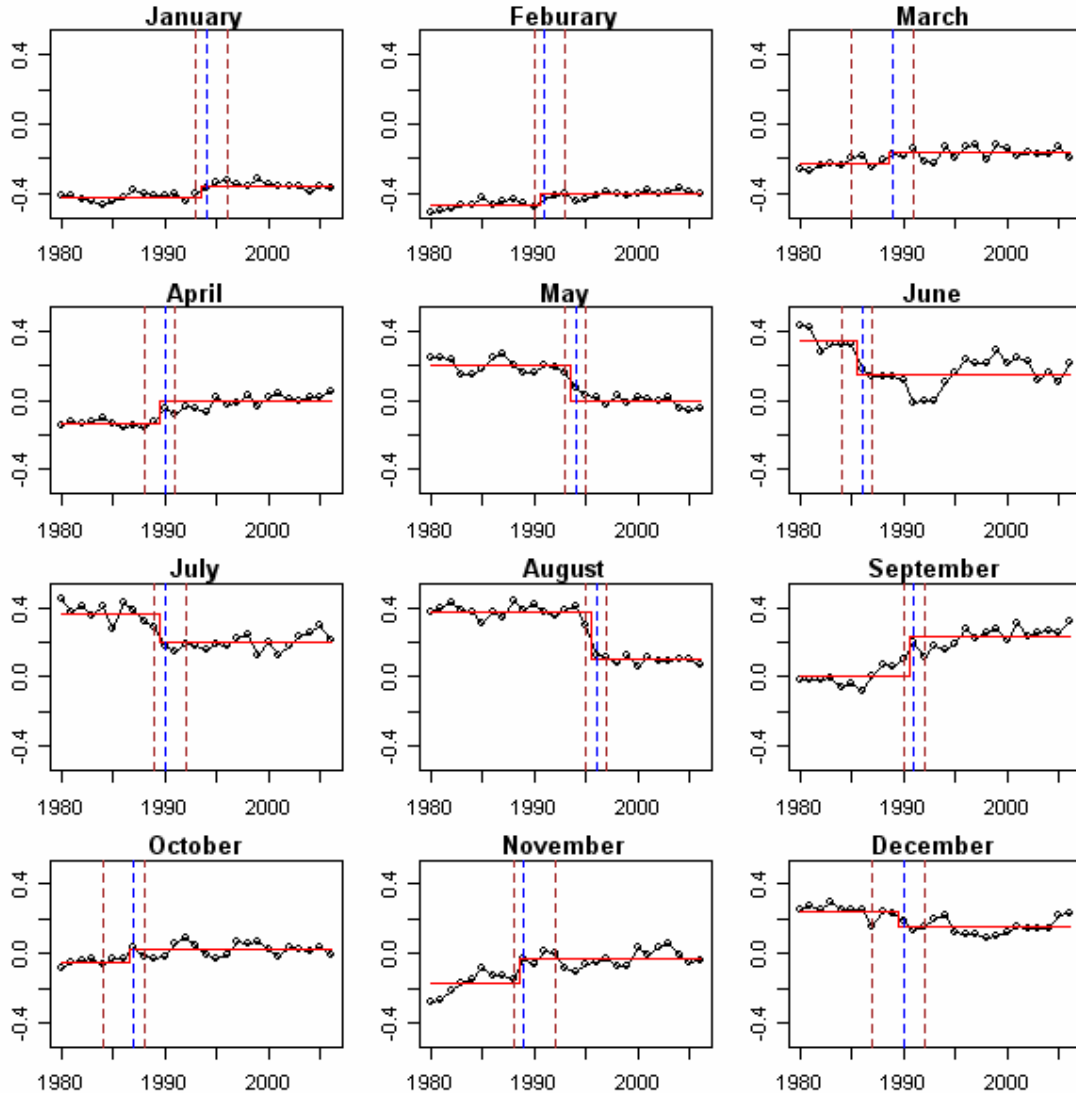
Table 2 Date of structural Break identified by ART and BP. Minimum regime length 3.

Month	ART	BP	BP 95% confidence interval
January	1993.5	1994	[1993, 1996]
February	1990.5	1991	[1990, 1993]
March	1988.5	1989	[1985, 1991]
April	1989.5	1990	[1987, 1991]
May	1993.5	1994	[1993, 1995]
June	1985.5	1986	[1984, 1987]
July	1989.5	1990	[1989, 1992]
August	1995.5	1996	[1995, 1997]
September	1990.5	1991	[1990, 1992]
October	1986.5	1987	[1984, 1988]
November	1989.5	1990	[1987, 1992]
December	1989.5	1990	[1987, 1992]

Overall the results using a regime length of 3 (table 2) are broadly similar to those using a minimum regime length of 10 (table 1). The most likely seasonal break identified by BP and ART changes only for March, June and October. A minimum regime length of 3, rather than 10, provides a much larger range of possible years for a break point to occur, since the first and last seven years of the SI series can now be tested for break points. The similarity of the results using the two different minimum regime length gives us some confidence having a minimum regime length of 3 will not significantly increase the number of spurious seasonal breaks.

We are not confident at this stage that BP and ART can replace inspection of the SI plots to identify seasonal breaks, but feel it is useful to combine the approaches. This is done by overlaying the seasonal break points and regime levels on the SI plots (figure 4). This could lead to more consistency and better interpretation of those plots as this addition information provides input into when a seasonal break may have occurred along with an indication of its significance. It also provides input into the scale of the break across the various months, which is important when trying to define a possible structural change in the seasonal pattern involving a large number of months.

Figure 4: SI plots with break points (vertical blue line) and their 95% confidence intervals (red vertical lines) and regime levels (horizontal red line). Minimum regime length is 3.



6. Monte Carlo simulations

As the seasonal break pattern we expected from the change in the timing of school holidays does not seem to occur we have done some preliminary investigations using synthetic time series with 2 modelled seasonal breaks. For each seasonal break 500 series are generated from the model which consists of a mean trend of zero, a seasonal pattern that is a sine function and Normally distributed noise (eg white noise)

$$y_t = s_t + e_t = 10 * \sin\left(\frac{p}{6} t\right) + e_t,$$

Based on expected types of seasonal breaks we modelled 2 types of seasonal breaks.

1. A general shift of one month in the monthly pattern. This resembles a seasonal change that occurred some 20 years ago when the central government financial year changed from a March year to a June year. We term this a 'seasonal pattern shift'.

2. A seasonal break as expected for STD. That is, two seasonal factors swapped after the seasonal break point. We term this a 'seasonal factor exchange'.

The length of each simulated series is 240 points (ie a monthly series of 20 years). In each set of simulations we had the seasonal break occurring either in the fifth, tenth or fifteenth cycle for each of 500 replications. For simplicity we modelled one seasonal break. We have passed the series through X-12 to produce the 12 SI series. The minimum regime length is 3, and we used ART to identify the most likely seasonal break.

Seasonal pattern shift

In a seasonal pattern shift we move the sine function to the right by one month at the break point. In our model before the seasonal break the peak is in April and in May afterwards, but the shape of the seasonal pattern remains the same (figure 5).


Detecting a seasonal break in a particular seasonal factor will depend on the change in the factor before the break compared to its value after the break. The greater the change in the value of the seasonal factor the greater the likelihood of its detection.

Examining figure 5 and looking at the vertical separation of the points (ie the change in value of the seasonal factors) before and after the shift one would expect seasonal breaks to be more readily identified for those months not near the peak or trough. The results in table 3 support this: the months at the peaks and troughs (April and May, October and November respectively) are not identified as breaks by ART, while for the other months ART does identify a seasonal break. This shows that different types of seasonal pattern shifts, in terms of when the shift occurs in the seasonal cycle along with the number of months the pattern is shifted, will produce different patterns of seasonal breaks for individual months.

Table 3: Seasonal Pattern Shift – Percentage of breaks identified by ART in a cycle

Month	Break in the 5th cycle (%)		Break in the 10th cycle (%)		Break in the 15th cycle (%)	
	5 th	6 th	10 th	11 th	15 th	16 th
January	8.4	75.2	8.0	76.4	8.6	76.4
February	10.0	75.6	9.2	75.0	7.8	75.8
March	13.6	52.2	12.6	57.2	11.2	57.8
April	12.2	16.4	9.0	14.2	9.0	19.8
May	16.4	8.0	16.0	10.8	16.0	9.0
June	55.4	11.6	59.4	10.2	56.0	10.8
July	73.2	10.6	72.0	8.8	78.8	8.0
August	70.0	9.8	74.2	10.0	75.2	9.0
September	60.0	9.0	50.6	13.2	54.2	13.2
October	16.4	11.2	16.2	8.8	14.0	10.0
November	14.8	8.2	17.6	8.8	16.0	10.6
December	54.4	9.6	55.8	13.4	57.0	9.0

Figure 5: Seasonal Pattern Shift – Seasonal pattern prior to break (solid line); seasonal pattern after break (dashed line).



6.2 Seasonal Factor Exchange

In this seasonal break model we have swapped the March and April factors (figure 6). We can see from figure 6 that, as we have swapped the factors at the peak of the seasonal cycle, a seasonal break may be difficult to identify as the seasonal factors have not changed markedly. The results from applying ART (table 4) support this. However it is worth noting that ART has increased the probability of a break at the seasonal break point (ie March and April) compared to the probabilities for the other months. In any case it is clear, as for seasonal pattern shifts in the previous section, that where the structural break occurs in the seasonal cycle has a major effect on the identification of seasonal breaks in the individual SI series


This leads us to conclude that it is not the individual SI series alone that are relevant, but rather that the distribution of seasonal breaks in all the SI series needs to be considered. To sum up, it would be useful to have an understanding of possible abrupt changes in the seasonal cycle and how that change will affect the individual monthly seasonal factors. We believe ART with its tree structure for its outputs may be better suited than BP for this task. However a more work is required investigating a range of

possible changes seasonal patterns before it would be possible to examine the break point pattern across months and identify possible changes in the seasonal pattern.

Table 4: Seasonal Factor Exchange – Percentage of breaks identified by ART in a cycle.

Month	Break in the 5th cycle (%)		Break in the 10th cycle (%)		Break in the 15th cycle (%)	
	5 th	6 th	10 th	11 th	15 th	16 th
February	6.8	6.6	5.4	6.6	7.2	7.2
March	11.8	4.8	13.0	10.6	16.8	11.6
April	18.4	8.2	14.6	9.0	16.0	10.8
May	6.4	6.6	5.0	5.2	8.4	6.6

Figure 6: Seasonal Factor Exchange – Seasonal pattern prior to break (solid line); seasonal pattern after break (dashed line).



7. Conclusions

Statistics New Zealand produces a range of time series beyond those that it directly measures. By providing seasonally adjusted series and trend estimates Statistics New Zealand assists users in their interpretation and understanding of the time series.

However it is important that any major changes in the data generating process associated with a time series is effectively identified as early as possible. Changes not identified and adequately modelled may lead to volatile and less useful seasonally adjusted series and trend estimates. This requires not only identifying breaks in trends, but also breaks in seasonal patterns.

We have investigated the use of structural break methods, normally used to identify trend breaks, to breaks in seasonal patterns. Our work focused on the individual seasonal factors in the seasonal pattern by applying the structural break methods to the SI series produced by X-12. Using an output from X-12 in this way has been revealed to have potential for identifying seasonal breaks. We have also shown that using a short regime length does not appear to affect the quality and stability of the seasonal break identification, thus it may be feasible for Statistics New Zealand to identify seasonal breaks in a timely manner.

It is clear that examining the seasonal factors for individual months is only a first step to identifying, interpreting, and understanding the break in the seasonal pattern. This is to be expected because the seasonal pattern over any annual cycle should average out to the trend. It is therefore necessary to do further simulations with a range of changes in seasonal patterns, as well undertake some theoretical research, possibly using some research into co-breaking (Hendry & Massmann, 2005).

References

- Bai, J., and Perron, P. (1998), "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica*, **66**, 47–78.
- Bai, J., and Perron, P. (2003), "Computation and analysis of multiple structural change models," *Journal of Applied Econometrics*, **18**, 1–22.
- Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. (1984), *Classification and Regression Trees*, Monterey CA: Wadsworth & Brooks.
- Cappelli, C., and Reale, M. (2005), "Detecting changes in mean levels with Atheoretical Regression Trees". *Research report UCMSD 2005/2*, Department of Mathematics and Statistics, University of Canterbury, NZ.
- Doherty, M. (2001). The Surrogate Henderson Filters in X-11. *Australian and New Zealand Journal of Statistics*, **43**, 385–392.
- Fisher, W.D. (1958), "On grouping for maximum homogeneity," *Journal of the American Statistical Association*, **53**, 789–798.
- Hansen, B. (2001), "The new econometrics of structural change: Dating breaks in U.S. labour productivity," *Journal of Economic Perspectives*, **15**, 117–128.
- Haywood, J., and Randal, J. (2006), "Seasonality and Structural Breaks: NZ visitor arrivals and 9/11," *MSCS Research Report 06-10*, School of Mathematics, Statistics and Computer Science, Victoria University of Wellington, New Zealand.
- Henderson, R. (1916), "Note on Graduation by Adjusted Average". *Transactions of the American Society of Actuaries*, **17**, 43-48. Note: a reproduction of this paper is available in Appendix 2 of Australian Bureau of Statistics (2003).
- Hendry, F.D., and Massmann, M. (2005), "Co-breaking: Recent advances and a synopsis of the literature," *Journal of Business and Economic Statistics*, **25**, 33-51.
- Ladiray, D., and Quenneville, B (2001), *Seasonal Adjustment with the X-11 Method*, Berlin:Springer.

Ripley, B. (2007), tree, version 1.0-26. CRAN:<http://cran.r-project.org/web/packages/tree/index.html>

Zeileis, A., Leisch, F., Hansen, B., Hornik, K., and Kleiber, C (2008), strucchange, version 1.3-3. CRAN:<http://cran.r-project.org/web/packages/strucchange/index.html>