Seasonality and structural breaks: NZ visitor arrivals and 9/11

John Haywood* John Randal†

12 December 2006

Abstract

We analyse monthly short term visitor arrival time series for New Zealand, to assess the effect of the 11 September 2001 terrorist attacks. We demonstrate that while some historical events have had a marked structural effect on trends in those arrivals, 9/11 was not one of these. Our conclusions are drawn on the basis of an initial nonparametric analysis, followed by a new iterative approach to fitting parametric structural break models, motivated by iterative methods for seasonal data. We suggest that iterative estimation methods are a simple but crucially important feature of this approach to seasonal modelling. Consequently, our new approach has potential applications for model-based seasonal adjustment, in addition to the dating of structural changes.

Key words: Break dates; Endogenous dating of structural changes; Iterative fitting; Multiple breaks; Seasonal adjustment; Trend extraction.

1 Introduction

There seems little doubt that the terrorist attacks of 11 September 2001 have had a pronounced influence on world events since that time. For example, see US Department of State (2004), for a summary of 100 editorial opinions from media in 57 countries around the world, commenting on the three years following September 2001. Those terrorist events and their subsequent effects have been used to explain apparent movements in many time series, and in this paper we focus on a particular example: the number of short term visitor arrivals to New Zealand.

The economic importance of tourism to New Zealand has recently increased considerably. As Pearce (2001) notes in his review article, international visitor arrivals increased by 65% over the period 1990 to 1999 while foreign exchange earnings increased by 120% (in current terms). More recently, for the year ended March 2004 tourism expenditure was $17.2 billion (Statistics New Zealand, 2005). In that year, the tourism industry

*School of Mathematics, Statistics and Computer Science, Victoria University of Wellington, PO Box 600, Wellington, New Zealand. email: john.haywood@vuw.ac.nz.

†School of Economics and Finance, Victoria University of Wellington.

made a value added contribution to GDP of 9.4%, split between direct (4.9%) and indirect (4.5%) contributions, and 102,700 people (full time equivalent) had work that was directly engaged in tourism: 5.9% of the total employed workforce. Also in that year, tourism's 18.5% contribution to exports was greater than that of dairy products (14.3%), which in turn was greater than the contributions from meat and meat products, wood and wood products, and seafood. The same ranking of industries was also seen in 2003.

The monthly short term New Zealand visitor arrivals series is one direct and easily recorded measurement of the international tourist contribution to the New Zealand economy. An essential precursor to development of tourism policy or business strategy is an understanding of the dynamic behaviour of these seasonal data. A classical time series decomposition includes unobserved components representing an evolving trend, a seasonal encapsulating regular deviation from the trend on a within-year basis, and an irregular, which is the residual or unexplained variation in the data. There are various ways to estimate these components, using both parametric and nonparametric approaches; see for example Harvey (1989), Hamilton (1994), Franses (1998) and Makridakis, Wheelwright & Hyndman (1998). Such a decomposition then allows an interpretation of the dynamic behaviour of visitor arrivals in terms of the estimated components.

Our focus is to detect any longer term, or structural, changes in trend or seasonal components of the arrivals as a result of the 9/11 events. We also wish to compare the magnitude of any 9/11 effects with those due to other causes. Consequently we do not wish to specify the dates of any structural changes, but rather estimate the number and position of these endogenously. To achieve this we use Bai & Perron's (1998, 2003) procedures for estimating multiple structural changes in a linear model. Their approach permits periods of stable dynamic behaviour between relatively infrequent but significant changes to the parameters of the model. There is clearly no empirical requirement that changes in the trend and seasonal components occur simultaneously. As we demonstrate, changes typically occur more frequently in the trend.

It has been appreciated within the seasonal adjustment literature for many decades that it is usually necessary to iterate between estimation of trend and seasonal components; see Bell & Hillmer (1984) for a review of the historical development of seasonal adjustment. In addition, iteration provides a natural way to separately estimate the structural changes in the two components. In contrast, direct application of Bai & Perron's (1998, 2003) methodology fits components simultaneously and yields a relatively poor decomposition, as we demonstrate via simulation and analysis of the visitor arrivals data. We propose a new iterative fitting procedure which gives much improved performance.

In Section 2 we present an exploratory data analysis (EDA) of New Zealand visitor arrivals and a discussion of some apparent sources of variability in the data. Section 3 motivates and presents a parametric model that allows separate structural changes in the trend and seasonal components. In Section 4 we use our iterative approach to model the arrivals data and in Section 5 we give some concluding comments. We find there is actually little to suggest that the September 11 incidents had much effect on New Zealand visitor arrivals, when viewed in the context of 'normal' historically observed movements. In contrast, we identify some other historical events which do appear to have affected visitor arrivals to New Zealand quite markedly.

2 EDA of short term visitor arrivals to New Zealand

We consider 25 complete years of monthly short term visitor arrival series from January 1980 to December 2004. The arrivals are from the seven most important countries of origin, ranked by current proportion of the total: Australia, UK, USA, Japan, Korea, China, Germany, as well as a residual series from 'Other' origins. We analyse these series individually along with their aggregate, denoted 'Total'.

As seen in Figure 1 a 'U'-shaped seasonal pattern is common, with visitor numbers reaching a local maximum in the summer months December to February, and a local minimum in the winter months June and July. Australian and UK arrivals appear to be growing at a relatively steady rate. In contrast, a large downturn in arrivals from the USA is evident in the late 1980s, a period which immediately followed the stock market crash of October 1987. The trend in Japanese arrivals levels off over the last 15 years. The effect of the Asian financial crisis of 1997 is evident especially in the Korean data, with visitor numbers dramatically reduced just after this event. Arrivals from China contain perhaps the most visible short term effect in these series, which is due to the SARS epidemic that virtually eliminated international travel by Chinese nationals during May and June 2003. German arrivals show a clear change from exponential growth prior to the early 1990s to a more stable pattern in recent times. The Other arrivals show a SARS effect much less prominent than that seen in the Chinese arrivals, as do some further series including Total arrivals. One of the more obvious shifts in the aggregate Total series appears to be linked to the Korean downturn, which can be attributed to the Asian financial crisis.

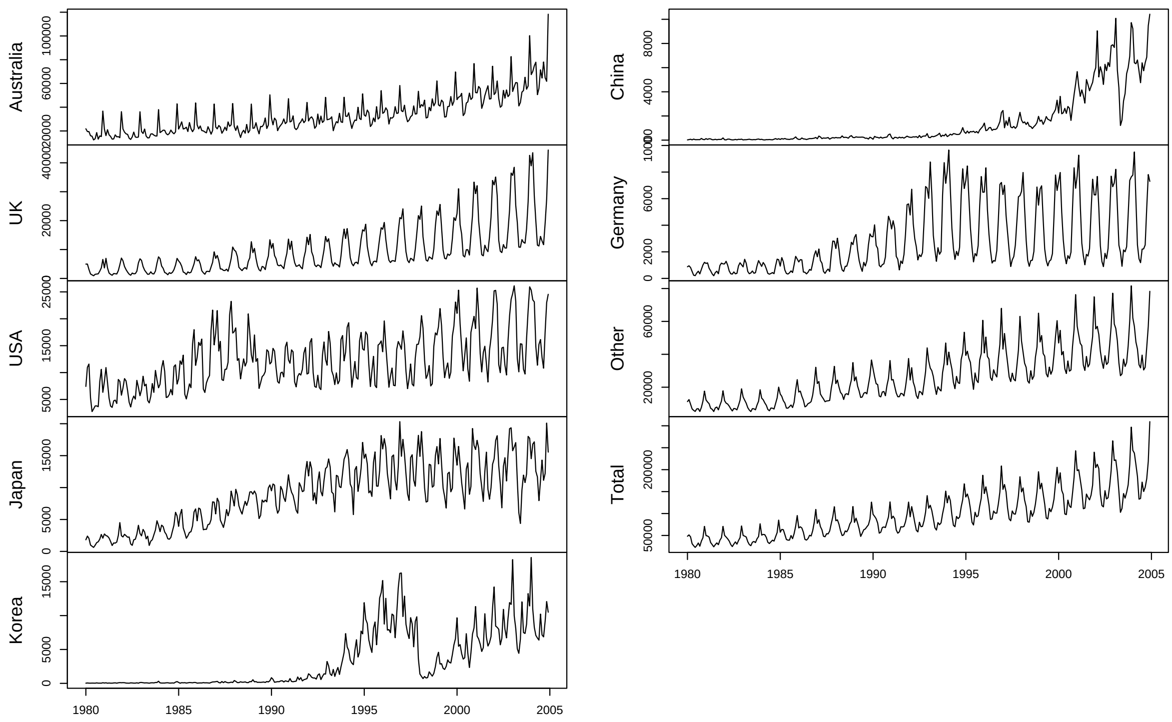

Figure 1: Monthly short term visitor arrivals to New Zealand, by origin, from January 1980 to December 2004. The vertical scales are not equal.

The Asian financial crisis of 1997-1998 markedly affected stock markets and exchange rates in nine East Asian countries: Hong Kong, Indonesia, Japan, Korea, Malaysia, Philippines, Singapore, Taiwan and Thailand. See Kaminsky & Schmukler (1999) for a chronology of the crisis, from the official onset marked by the devaluation of the Thai baht on 2 July 1997 up to the resignation of Indonesian President Suharto in May 1998. Kaminsky & Schmukler (1999) suggest the presence of important contagion effects in those markets, based on an analysis of identified market jitters. More recent analysis by Dungey *et al.* (2004) suggests, however, that increased exchange rate volatility observed in Australia and New Zealand around that time was not due to contagion from Asian countries, or unanticipated factors, but rather to common (anticipated) world factors such as trade linkages. This is one context in which changes in short term visitor arrivals to New Zealand from Asian countries around 1997-1998 can be viewed, since tourism has become such an important sector of the New Zealand economy, as noted above. In particular, Korea is one of the five source countries with the largest current proportion of visitors to New Zealand (Table 1).

Table 1: Summary statistics for the monthly proportion of visitors to New Zealand, by origin. The final three columns give proportions of the Total for the entire 25 year sample period, and the five-year periods 1980-1984 and 2000-2004, respectively.

	Min	LQ	Median	UQ	Max	80-04	80-84	00-04
Australia	21.8	30.0	35.9	41.6	58.8	33.8	44.9	33.3
UK	3.4	6.3	8.0	10.6	18.3	9.8	7.6	11.8
USA	6.3	10.3	13.0	16.3	29.4	12.4	16.7	10.0
Japan	2.8	7.1	9.1	11.0	17.8	9.2	5.9	7.8
Korea	0.0	0.2	1.0	4.3	10.5	3.4	0.2	4.8
China	0.0	0.2	0.4	1.2	4.7	1.4	0.1	3.1
Germany	0.8	1.5	2.2	3.4	7.5	2.9	1.8	2.6
Other	17.9	23.4	26.1	28.6	34.3	27.0	22.8	26.7

Table 1 shows that Australia is by far the biggest single source of visitors to New Zealand, accounting for almost exactly one-third of visitors in the 2000-2004 five year period and slightly more over the entire data period. The maximum proportion in a month from Australia was 58.8% in June 1985, and the minimum was 21.8% in February 1997. An Australian influence is notable in the Total arrivals, because as the nearest neighbour to a geographically isolated country, arrivals from Australia exhibit variation not seen in the remaining data. As seen in Figure 1, the Australian data has a regular seasonal pattern which is quite different from that of any other country. A closer examination indicates three peaks per year before 1987 and four thereafter; we discuss this further in Section 4.

One way of estimating unobserved trend and seasonal components is to use a robust, nonparametric technique such as STL (Cleveland *et al.*, 1990), e.g., as implemented in R (R Development Core Team, 2006). This procedure consists of an iterated cycle in which the data is detrended, then the seasonal is updated from the resulting detrended seasonal subseries, after which the trend estimate is updated. At each iteration, robustness weights are formed based on the estimated irregular component and these are used to down-weight outlying observations in subsequent calculations.

A typical STL decomposition is shown in Figure 2 for the natural logarithm of the Total arrivals. The log transformation is commonly used to stabilise a seasonal pattern which increases with the level of the series, and effectively transforms a multiplicative decomposition into an additive one. The plot shows an evolving seasonal pattern, an upward trend with several changes in slope, and a relatively small irregular component. A vertical line is added to indicate September 2001. There is no obvious (structural) change in the trend at or about this month, although there is a reduction in the slope of the trend nearer the start of 2001, which we discuss further in Section 5 below. More prominent is a cluster of negative irregulars immediately following 9/11, the largest of which is the third largest negative irregular in the sample period. Jointly though, these irregulars are smaller and less persistent than those occurring at the time of the SARS outbreak in 2003. Our exploratory analysis with STL thus suggests that while the events of 9/11 may have had a moderate short term (irregular) effect, there is nothing to suggest that a longer term (structural) effect occurred. We investigate this hypothesis more formally in Section 4.

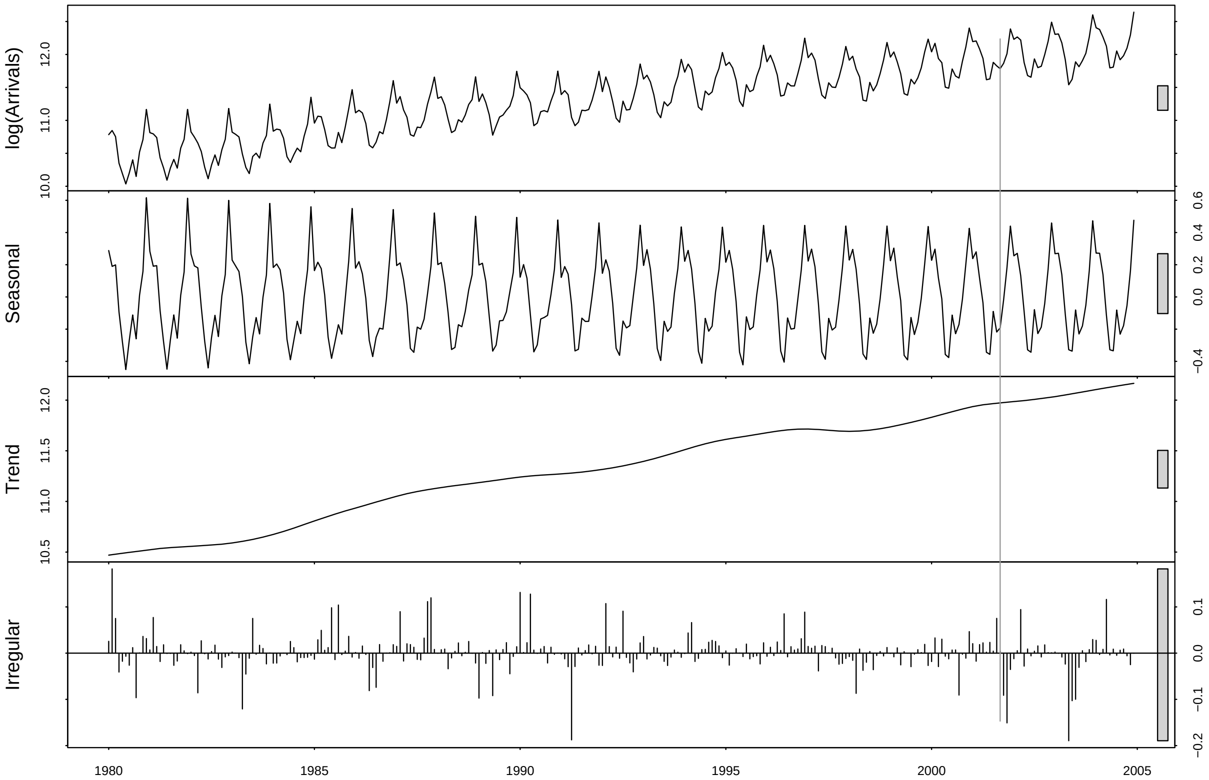

Figure 2: The STL decomposition of the log aggregate monthly visitor arrivals to New Zealand from January 1980 to December 2004. The vertical grey line is at September 2001, and the solid bars on the right hand side of the plot are all the same height, to aid comparisons.

3 Iterative break estimation for seasonal data

Bai & Perron (1998, 2003) present a methodology for fitting a linear model with structural breaks, in which the break points, i.e. the times at which the parameters change, are determined optimally. The optimal positions of m break points are determined by minimising the residual sum of squares, for each positive integer $m \leq m_{\max}$. The optimal number of break points ($0 \leq m^* \leq m_{\max}$) may then be determined by, for example, minimising an information criterion such as BIC. Given a sample of T observations, the selected break points are estimated consistently, with rate T convergence of the estimated break fractions (that is, the proportions of the data between consecutive breaks).

The maximum number of break points, $m_{\max}$, is determined by the number of observations relative to the number of parameters in the model. In general, for a model with m breaks and q parameters, at least q observations are needed between each pair of break points, requiring at least $T \geq (m+1)q$ observations in total. Clearly if the model has many parameters, fewer break points can be estimated from a given sequence of observations.

We consider implementing this approach for a time series of the form

$$Y_t = T_t + S_t + I_t \qquad t = 1, \ldots, T$$

where Y_t are the observed data, T_t is an unobserved trend component, S_t is an unobserved seasonal component with seasonal period s, and I_t is an unobserved irregular component. We assume that between two break points t^*_{j-1} and t^*_j ($j = 1, \ldots, m+1$), the trend T_t is linear,

$$T_t = \alpha_j + \beta_j t \qquad t = t^*_{j-1} + 1, \ldots, t^*_j$$

and the seasonal component is fixed,

$$S_t = \sum_{i=1}^{s-1} \delta_{i,j} D_{i,t} \qquad t = t^*_{j-1} + 1, \ldots, t^*_j$$

where $D_{i,t}$ are seasonal dummies. We use the convention that $t^*_0 = 0$ and $t^*_{m+1} = T$ (Bai & Perron 1998). Under these assumptions, we note that for daily or monthly data (with $s = 7$ and $s = 12$ respectively), and for quarterly data (with $s = 4$) to a lesser extent, the trend component will be parsimonious relative to the seasonal component.

Bai & Perron's (1998, 2003) methodology offers two alternatives for estimating the break points in such a model. The first is that the coefficients of one component are fixed over the entire sample period (a partial structural change model); the second is that parameters in both components should have the same break points (a pure structural change model). We believe that neither of these options is likely to be satisfactory for the type of data we are examining in this paper, i.e. seasonal time series with evolving trends, and large s (in this case, $s = 12$).

When considering trend extraction and assuming that structural breaks will be required, in general we wish to allow break points in the seasonal component, which is inconsistent with a partial structural change model. Conversely, we would not necessarily wish to constrain any seasonal break points to occur at the same places as the trend break points, as required in a pure structural change model. On the face of it, this requirement is not necessarily restrictive, since the parameter estimates of one component are not forced to change from one segment of the data to the next. However, when selecting the optimal number of break points using a penalised likelihood criterion, e.g. BIC, this will compromise our ability to detect break points in the data, i.e. the selected number of breaks may be too low. One example of where these issues may be important is in arrivals from Australia. As noted in Section 2, the Australian arrivals seem to have a seasonal break point in 1987 (changing from three peaks to four), with no apparent change in trend.

To address this concern we estimate the trend and seasonal components separately, using a new iterative approach motivated by the Macaulay cycle seasonal decomposition method (Macaulay, 1931) and the iterative technique of STL. This allows more flexible structural break estimation than fitting both components simultaneously. As above, we assume that the time series can be decomposed into a piecewise linear time trend and a piecewise constant seasonal pattern. Each component is then estimated using the methodology of Bai & Perron (1998, 2003), implemented in R (R Development Core Team, 2006) using the strucchange package of Zeileis *et al.* (2003). In this package the default method of selecting the number of breaks uses BIC.

We estimate the trend of the data Y_t using a piecewise linear model for the seasonally adjusted time series $V_t = Y_t - \hat{S}_t$, i.e.,

$$V_t = \alpha_j + \beta_j t + \epsilon_t \qquad t = t^*_{j-1} + 1, \dots, t^*_j$$

for $j = 1, \dots, m+1$, where ϵ_t is a zero-mean disturbance and t^*_j, $j = 1, \dots, m$, are the unknown trend break points. For the first iteration, we set $\hat{S}_t = 0$ for all t.

Once the trend has been estimated, we estimate the seasonal component of Y_t using a piecewise seasonal dummy model for the detrended data $W_t = Y_t - \hat{T}_t$, i.e.,

$$W_t = \delta_{0,j} + \sum_{i=1}^{s-1} \delta_{i,j} D_{i,t} + \nu_t \qquad t = t'_{j-1} + 1, \dots, t'_j$$

for $j = 1, \dots, m'+1$, where $D_{i,t}$ are the seasonal dummies, ν_t is a zero-mean disturbance and t'_j, $j = 1, \dots, m'$, are the unknown seasonal break points. As before, we take $t'_0 = 0$ and $t'_{m'+1} = T$. The estimates $\hat{\delta}_{i,j}$ are adjusted at the end of each iteration so that they add to zero within each full seasonal cycle (between seasonal breaks), to prevent any change in trend appearing as a result of a seasonal break happening 'mid-year'. That is,

$$\sum_{i=0}^{s-1} \hat{\delta}_{i,j} = 0 \qquad \text{for all } j.$$

This estimation process is then iterated to convergence of the estimated break points. We are thus able to estimate a trend which, due to its parsimonious representation, is able to react to obvious shifts in the general movement of the data. If required, we are able to identify important changes in the seasonal pattern separately.

We now illustrate the importance of this method using simulated data. Consider a time series with piecewise linear trend given by

$$T_t = \begin{cases} 20 + 0.05t & t = 1, \dots, 78 \\ 23.9 & t = 79, \dots, 234 \\ 23.9 + 0.05(t - 234) & t = 235, \dots, 312 \end{cases}$$

and fixed seasonal component with a break point at $t = 156$. The two seasonal cycles are shown in Figure 3, and are identical except for the ordering of the Jan/Feb, Mar/Apr, Jul/Aug and Sep/Oct values. The data are given by

$$Y_t = T_t + S_t + I_t \qquad I_t \sim \text{ i.i.d. } N(0,1)$$

and thus have three break points: two associated solely with the trend, and one associated only with the seasonal component.

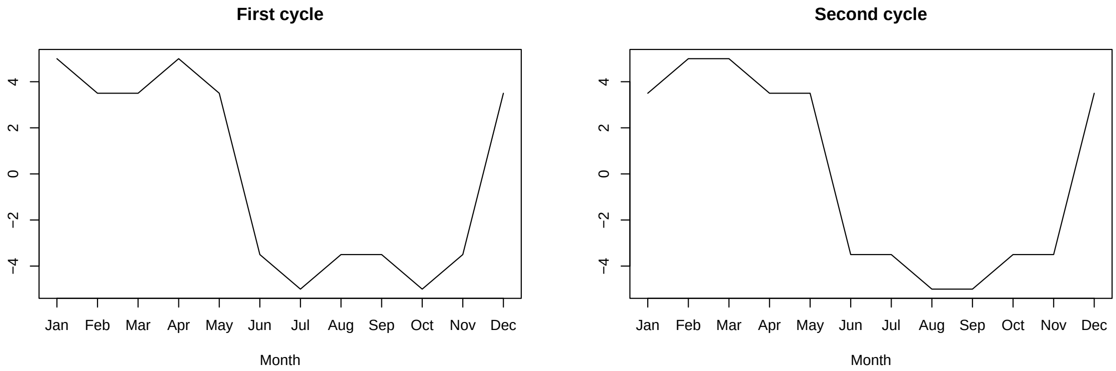

Figure 3: The two seasonal cycles used for the seasonal component in the illustrative simulation.

We simulated 500 independent series as above, and for each of them estimated the trend and seasonal components simultaneously using the Bai & Perron approach with a minimum period between breaks of 36 observations (3 years of data, and roughly 11.5% of the series). In particular, we restricted the parameters in both the trend and seasonal components to change simultaneously; i.e., a pure structural change model. Between breaks, the constant term in the estimated trend was corrected so that the seasonal component added to zero. We also applied the iterated methodology to the same series, fitting the trend and seasonal components separately. Figure 4 shows the estimated break points for the 500 series using the two competing methodologies. In both panels are boxplots of the estimated break points for the series, and these have been grouped depending on how many break points were estimated.

The estimated break points obtained fitting both components simultaneously are in the upper panel of Figure 4. In 343 series (68.6%) only one break point was estimated, and the sample distribution of these is summarised in the extreme left boxplot. The estimates appear to be unbiased for the central (seasonal) break point. In the middle of the upper panel are the sample distributions of the two break points, as estimated in 146 series (29.2%). These appear to be biased estimates of the true trend break points, with each being closer to the seasonal break point than to the nearest end of the series. On the right, the sample distributions of the three break points are displayed, estimated in 11 of the series (2.2%). These appear unbiased, although the sample size is very small. In conclusion, fitting the two components simultaneously has allowed us to correctly identify the true break points in only 2.2% of the series. Despite correctly dating the breaks in these 11 series, this approach does not attribute the change to any one component of the structural model. Clearly though, each true change affects only one of the trend or seasonal components.

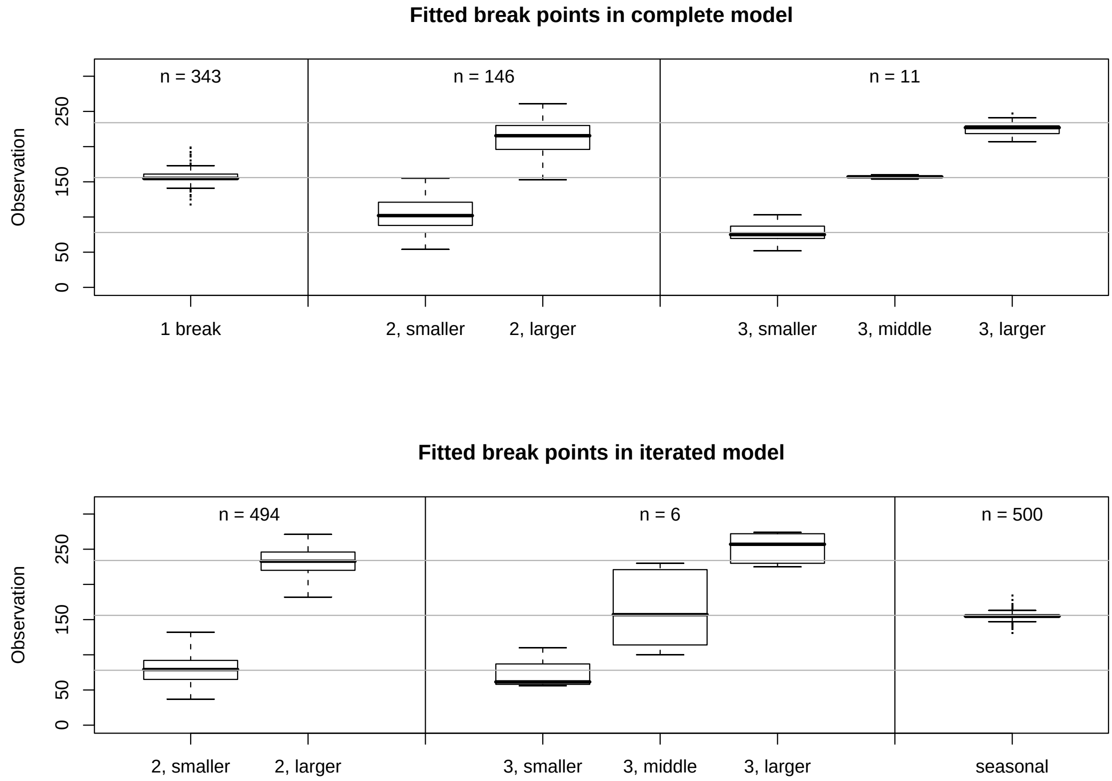

Figure 4: The estimated break points for 500 simulated series. True break dates are at the horizontal grey lines. Results in the upper panel are for the Bai and Perron methodology, with the new iterated methodology in the lower panel. The number of cases is shown above each boxplot, while the total number of estimated breaks is given below. In the iterated panel, trend break points are shown on the left and the (single) seasonal break point on the right.

In the lower panel of Figure 4 are comparable results for the new iterated approach. All series have at least two estimated trend break points and a single seasonal break. Two trend breaks are estimated in 494 series (98.8%) and are displayed on the left; they appear to be unbiased estimates of the true trend break points. In the remaining six series (1.2%), a third trend break is estimated. The estimated seasonal break points for all 500 series are shown on the right. These are clearly unbiased, and estimated relatively precisely. In conclusion, fitting the two components iteratively allowed us to correctly identify and attribute the trend and seasonal break points in 98.8% of the series; a dramatic improvement over the simultaneous approach. Use of sequential F-tests instead of BIC to select the number of breaks in the simultaneous approach does not substantially change the number of series with the correct number of estimated breaks.

Figure 5 shows a single example series that results in typical estimation behaviour, with the true trend and both estimated trends also plotted. For the complete model, BIC selects only one break point at $t = 168$. Applying the iterated methodology to the data, two trend break points are estimated at $t = 83$ and $t = 249$, and a single seasonal break point is estimated at $t = 157$. Note that the complete model induces a 'quadratic' trend in the residuals on either side of the single break, which is expected after viewing the estimated and true trends together as in Figure 5. That (local) trending behaviour is

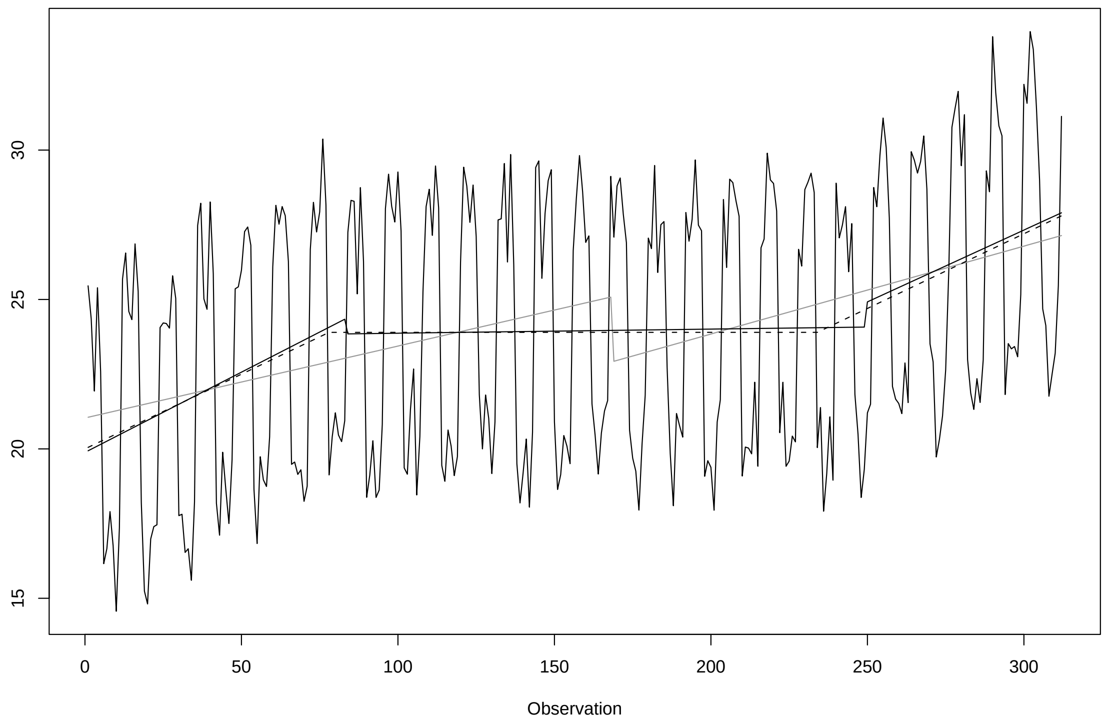

Figure 5: Simulated data (solid) with true piecewise linear trend (dashed). Estimated piecewise linear trends using the Bai and Perron methodology (grey) and the iterated methodology (black).

reflected by significant residual sample autocorrelations at low lags; clear evidence that the model is misspecified. In contrast, the residuals from the iterative approach show no significant autocorrelations at low lags, reflecting the more appropriate modelling of the true trend component.

This simulation illustrates the undesirable consequences of fitting the two components simultaneously in cases when a parameter-rich seasonal component has break points at times other than the break points of a relatively parsimonious trend component. Our new iterated approach to fitting such components appears to address this concern. In the next section we apply this iterated technique to the visitor arrivals data.

4 Modelling the arrivals using an iterated approach

The seasonal variation of the arrivals typically increases with the level of the series (Figure 1). Applying the new iterated approach directly to the untransformed data would certainly require seasonal breaks to account for the changes in amplitude of the seasonal component. This is clearly undesirable, in part because such changes typically evolve smoothly and so should not be modelled with abrupt changes. Consequently a stabilising transformation is needed.

A log transformation is one obvious possibility, but this does not yield an optimal stabilising transformation for all these series and instead we estimate a power transformation, identified using the robust spread-vs-level plots described in Hoaglin *et al.* (1983). For each individual series we calculate the median and interquartile range (IQR) of the monthly arrivals for each of the 25 calendar years, then regress log IQR on log median. The appropriate stabilising transformation is $x^{1-\text{slope}}$, and the transformed series are shown in Figure 6, with the estimated powers. Confidence intervals for the slopes in these regressions support the use of logs only in the case of the UK, USA and Total arrivals (i.e. a power of zero, or a slope of one). In the case of Germany the estimated power is negative, so $-x^{1-\text{slope}}$ is used to preserve order in the transformed arrivals. All further analysis is conducted on the transformed data.

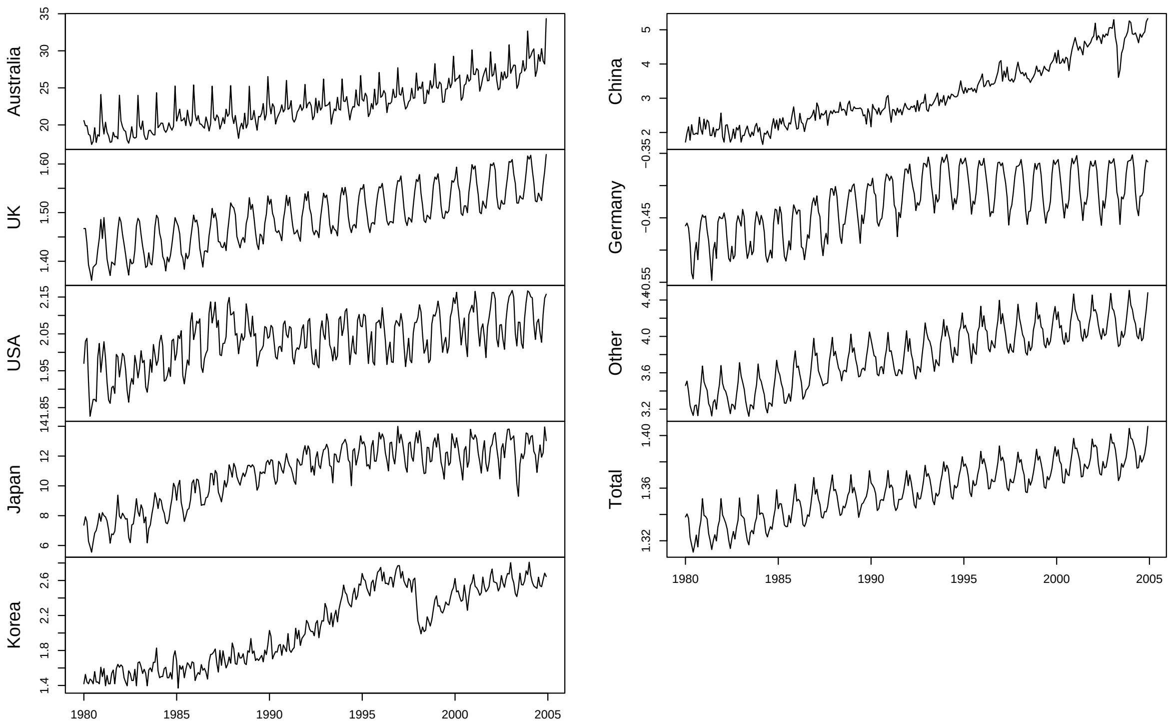

Figure 6: Power transformed monthly short term visitor arrivals to New Zealand, by origin, from January 1980 to December 2004. The power transformations are: Australia 0.3, UK 0.05, USA 0.08, Japan 0.27, Korea 0.11, China 0.18, Germany -0.11, Other 0.13, and Total 0.03.

In the case of the transformed arrivals data, each linear time trend requires two parameters, and each dummy seasonal an additional $s - 1 = 11$. Figure 6 indicates that for most series a linear time trend would need breaks. Further, while the seasonal patterns generally have constant variation over the length of the series due to the power transformations, we do not wish to preclude seasonal changes during the data period. As the simulation study demonstrated, the parameter-rich trend-plus-seasonal (complete) model would severely limit our ability to appropriately fit the data, since the large number of seasonal dummies would reduce the possible number of breaks, especially when selected by BIC.

As with the simulated data, we use a minimum period between breaks of 36 observations (in this case 12% of the data) for estimation of both trend and seasonal components. In fact, when estimating the trend and seasonal components iteratively, there is scope to reduce that minimum period for estimation of the trend component, since it only requires two parameters between breaks. This possibility further increases the flexibility of trends estimated using our new iterated approach. However, to simplify comparisons we have not pursued this option here. For the iterative approach, three iterations were sufficient to ensure convergence of the estimated break points in all cases but Other and Total, which each required four.

The estimated trend break points are shown in Table 2 along with estimated confidence intervals. The confidence intervals have been formed with heteroscedasticity and autocorrelation consistent estimates of the covariance matrix (Andrews, 1991). Note the intervals are not symmetric about the estimated breaks, indicating that the likely range of the true break date is very small in one direction.

Table 2: Estimated trend break points for the transformed monthly visitor arrivals to New Zealand, by origin, from January 1980 to December 2004. The middle column gives the estimated break points, while the first and third columns give the lower and upper 95% confidence limits respectively, estimated using a HAC estimate of the covariance matrix.

Australia	1984(5)	1985(1)	1985(2)
	1989(3)	1989(4)	1990(5)
	1997(10)	1997(12)	1998(1)
	2001(1)	2001(10)	2001(11)
UK	1985(10)	1986(1)	1986(4)
	1990(7)	1990(8)	1996(5)
USA	1982(12)	1983(3)	1986(12)
	1988(9)	1988(10)	1990(1)
	1998(6)	1998(8)	2001(6)
Japan	1987(3)	1987(6)	1987(8)
	1996(2)	1996(8)	1996(10)
Korea	1982(8)	1983(12)	1984(4)
	1990(9)	1990(10)	1990(11)
	1994(9)	1994(11)	1994(12)
	1997(10)	1997(11)	1997(12)
	2000(10)	2000(11)	2001(1)
China	1984(7)	1984(8)	1985(1)
	1988(10)	1989(7)	1989(9)
	1997(1)	2000(11)	2000(12)
Germany	1986(6)	1986(7)	1986(11)
	1994(5)	1994(6)	1994(7)
	1999(6)	1999(8)	2000(11)
Other	1983(1)	1983(3)	1983(4)
	1985(6)	1986(8)	1986(9)
	1990(7)	1990(10)	1990(12)
	1992(11)	1994(1)	1994(3)
	1997(3)	1997(6)	1997(8)
	2001(4)	2001(7)	2001(9)
Total	1982(12)	1983(1)	1983(3)
	1987(10)	1987(12)	1988(4)
	1989(8)	1990(12)	1991(4)
	1997(1)	1997(3)	1997(6)

The estimated parametric trends and break points (with 95% confidence intervals) are shown in Figure 7, along with nonparametric trends estimated by STL. September 2001 is included in only two confidence intervals, indicating the possibility that the terrorist events of 9/11 may be linked to a structural break in the trend of arrivals for those two origins: Australia and Other. Other is difficult to interpret given its composite nature, although it is plausible that the 9/11 events did have an effect on tourist behaviour in some of these countries. An alternative (or complementary) explanation is discussed in Section 5.

In the case of Australia, a break is estimated in the month following 9/11, which results in an increased trend slope but a decreased intercept. A relevant confounding effect is the collapse of Ansett Australia, which occurred just three days after the terrorist attacks of 9/11; hence it is impossible to separate these two effects with monthly data. The termination of flights by Ansett Australia and Ansett International on 14 September 2001 certainly affected capacity and timing of arrivals to New Zealand. In addition, in the following week, strike action targeted at Air New Zealand occurred at Melbourne and Perth airports (Air New Zealand had acquired control of Ansett Australia during the year preceding its collapse). Those strikes required the cancellation of all Air New Zealand trans-Tasman flights operating from Melbourne and Perth. These physical constraints on passenger numbers are a plausible explanation for a decrease in intercept, while the subsequent increase in the rate of arrivals from New Zealand's nearest neighbour is unlikely to have any causal links from the terrorist events of September 2001.

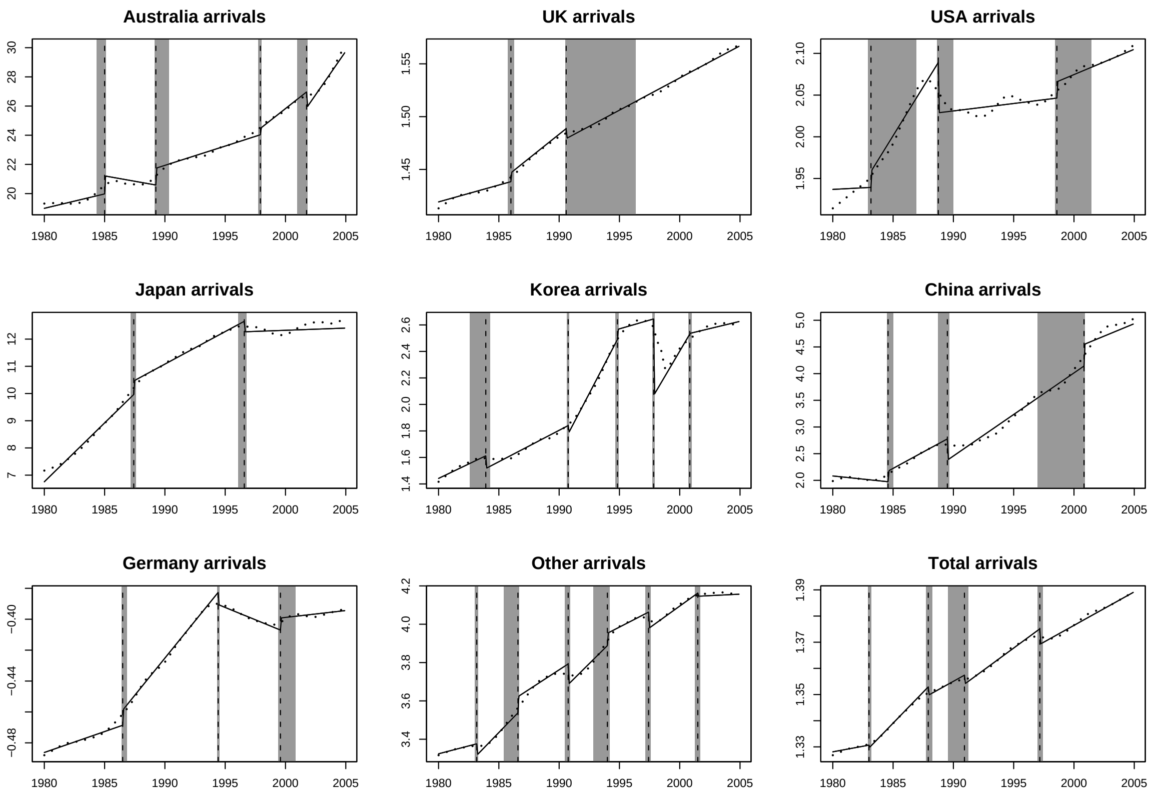

Figure 7: Estimated trends and trend break points for the transformed monthly visitor arrivals to New Zealand, by origin, from January 1980 to December 2004. The solid line is the piecewise linear time trend, while the dotted line is the estimated STL trend. The vertical dashed lines and grey regions respectively indicate the fitted break points and their 95% confidence intervals, estimated using a HAC estimate of the covariance matrix.

Focusing on Figure 7 more generally, we note that it is often difficult to distinguish between the two alternative trend estimates; i.e. those from our iterated approach and STL. In particular, the iterative parametric method achieves similar flexibility in its trend estimate to the nonparametric technique, with the latter essentially fitting linear time trends at each point in the series using only a local window of observations to estimate parameters. The break point technology allows instantaneous changes in the trend however, unlike the STL technique. In effect, STL is requiring an 'innovational outlier' approach to any structural changes in the data, while our parametric procedure models the changes directly and permits an 'additive outlier' approach. (In a series of papers, Perron and coauthors popularised the use of these 'outlier' terms, to describe an approach which is attributed to the intervention analysis work of Box & Tiao (1975).) An obvious contrast between the two approaches is seen in the Korean data at the time of the Asian financial crisis. The parametric break point is dated at November 1997 (with a narrow 95% confidence interval of October to December), which corresponds exactly to the month that the financial crisis first affected Korea (Kaminsky & Schmukler, 1999). However STL spreads the downward impact of the crisis over a number of months, in contrast to the observed behaviour.

Table 3 gives the estimated seasonal break points for the transformed arrivals, estimated using detrended data; the estimated seasonal components are shown in Figure 8. Korea, China, Germany and Other have no estimated seasonal break points. As the power transformations have effectively stabilised the seasonal variation, any changes in the seasonal patterns more likely reflect behavioural changes in the time of year when visitors arrive. For example, in Australia's seasonal pattern the 'middle' peak has moved and one extra peak has been added, reflecting a shift from a three-term school year to a four-term year in New South Wales in 1987 (NSW Department of Education, 1985). The placement of the seasonal break point coincides exactly with the final month under the old three term system, with the first holiday in the new sequence occurring in July 1987. The UK data show a shift in arrivals from the second half of the year to the first and a shift in the peak arrivals from December to February. The USA and Japanese arrivals have had relatively complex changes, while the Total series has seen most change in the winter months.

Table 3: Estimated seasonal break points for the transformed monthly visitor arrivals to New Zealand, by origin, from January 1980 to December 2004. The middle column gives the estimated break points, while the first and third columns give the lower and upper 95% confidence limits respectively, estimated using a HAC estimate of the covariance matrix. Korea, China, Germany and Other have no estimated seasonal break points.

Origin	Point estimate and 95% CI		
Australia	1987(1)	1987(6)	1987(9)
UK	1985(11)	1986(6)	1987(9)
USA	1995(1)	1995(4)	1995(12)
Japan	1987(10)	1988(6)	1988(12)
Total	1987(3)	1987(7)	1988(1)

To conclude this section, we compare the trend estimates obtained from our new iterated approach to the trends obtained fitting a complete structural break model (with 13 parameters between breaks), and using STL. In Figure 9 we present trends for the Korean arrivals and those from Other origins. We also show sample autocorrelation functions for the three sets of residuals from each series. The trends are all similar, but the agreement is closest for the iterated approach and STL. Some differences are evident particularly at the end of the series though, which would be important for prediction. For Other arrivals, the number of parameters required for the complete model clearly restricts

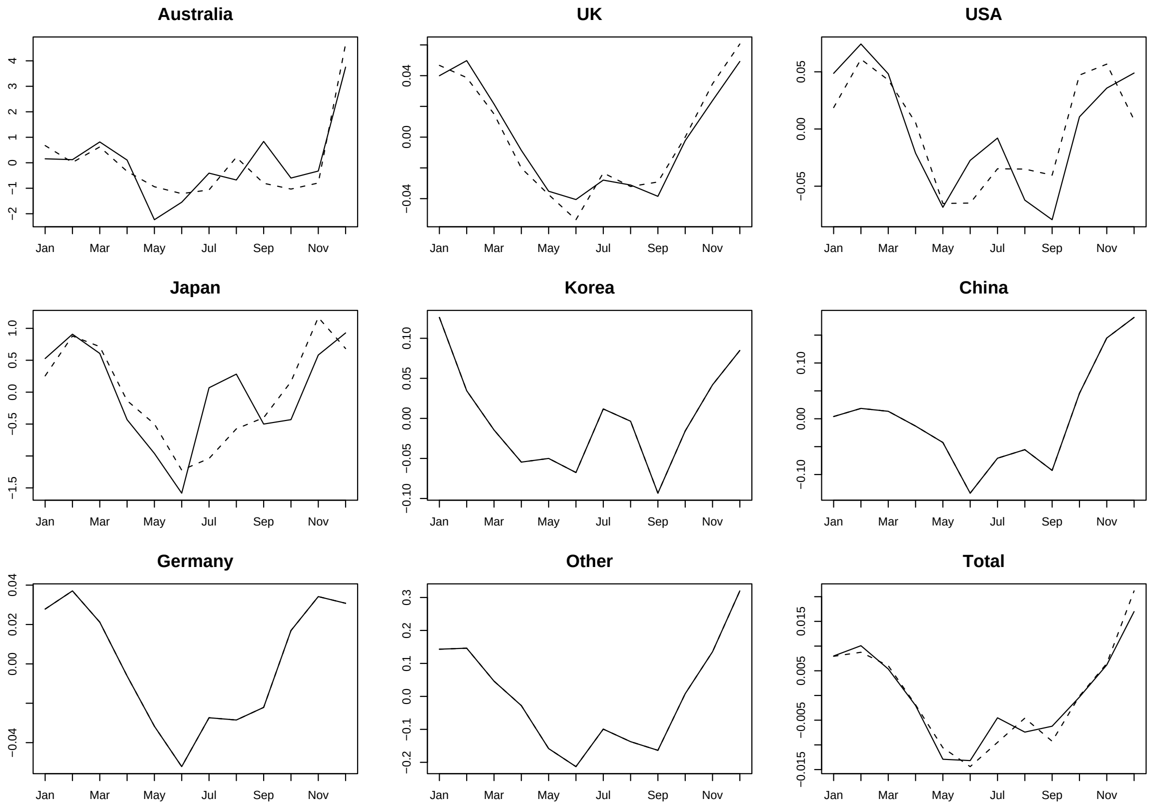

Figure 8: The estimated seasonal components for visitors to New Zealand by origin. The solid line is the final estimated seasonal component; it is the only estimate in four of the nine cases, where no seasonal breaks were detected. The five dashed lines are the seasonal components prior to the seasonal break points listed in Table 3.

the estimated number of breaks, leading to greater departures from the STL trend than achieved by iteration. The irregular components also favour the iterated approach over STL and the complete model, as the residuals for the latter are highly autocorrelated, especially at low lags. In contrast, the residuals of the iterated method exhibit far less autocorrelation, indicating a better overall decomposition.

5 Discussion

The growth in the number of visitor arrivals to New Zealand was lower than expected in late 2001 (e.g., by the New Zealand Ministry of Tourism, as noted in Haywood & Randal, 2005), yet there is no conclusive evidence to attribute this forecast error solely to the terrorist events of 9/11. The termination of flights by Ansett Australia on 14 September 2001 certainly affected capacity and timing of arrivals from Australia to New Zealand, and that would have affected Total arrivals in September 2001 somewhat as well. Indeed Australia is the only (individual) country of origin with a structural change in trend identified relatively close to 9/11. The subsequent rate of Australian arrivals to New Zealand in fact shows an *increase*, following an initial drop which is plausibly explained by the Ansett effect; see Table 2 and Figure 7.

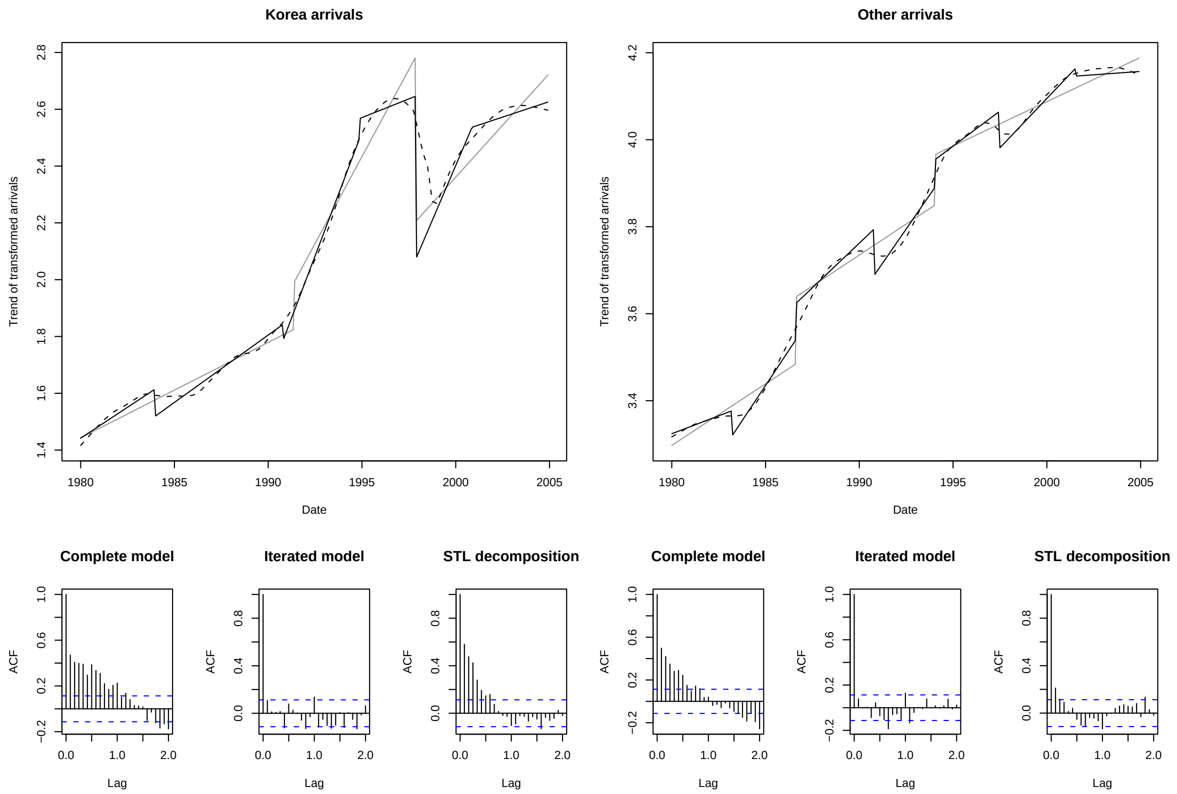

Figure 9: Trend estimates for the Korean arrivals and those from Other origins. The trend estimates are based on the complete model (grey), the new iterated approach (black) and STL (dashed). Also shown are sample autocorrelation functions for the residuals from the three methods.

A further plausible cause for the lower than forecast number of visitors is the US recession dated March 2001 (Hall *et al.*, 2001), along with the world wide flow-on effects from a slow down in the US economy. The recession predates 9/11 by six months but that is consistent with observed features of the data. In particular, March 2001 corresponds exactly to the minimum in the second difference of an STL trend of Total monthly (log) arrivals, indicating a maximum decrease in the slope at that time. It is possible that the slow down seen in the Other (composite) arrivals series, dated July 2001, may be due in part to the flow-on effects from this US recession.

It seems quite clear that the events of 9/11 did not have much influence on the longer term numbers of visitors to New Zealand, and especially not a negative influence. In contrast our analysis identifies other events which have had marked structural effects on the trends in these data, especially from certain countries of origin. In particular, the stock market crash of October 1987 preceded a dramatic decline in arrivals from the USA, followed by a sustained period of only moderate growth. In turn, both the intercept and slope of Total arrivals decreased in December 1987. Similarly, the Asian financial crisis of 1997-1998 precipitated a massive drop in arrivals from Korea, with the intercept and slope of Total arrivals again both decreasing in 1997. The SARS epidemic affected arrivals from China in a different way, with a very short-lived but large reduction, which we class as temporary and not structural. The overall effects of 9/11 might also be seen as temporary and negative, but of a smaller magnitude than those associated with SARS.

Estimation of structural breaks was facilitated by a new implementation of Bai & Perron's (1998, 2003) work. Use of an iterative approach to estimate the trend and seasonal components separately enabled us to locate structural breaks in the data, and to attribute these to either changes in the trend or the seasonal pattern. Estimating these components simultaneously did not achieve the same flexibility in the estimated components, nor in the location of the break points. The agreement between the estimated parametric trends from the iterated approach and the nonparametric STL trends is especially pleasing, as is the lack of residual structure around those parametric trends when compared to other trend estimates.

Finally, we note that our approach may be fruitful for model-based seasonal adjustment; something that is still extremely rare in the production of official statistics around the world. One major reason for this somewhat surprising rarity is a lack of appropriate flexibility in the evolution of model-based trends and seasonal components; see Bruce & Jurke (1996). Our parametric trends have similar flexibility to those of STL and our power transformation approach ensures that seasonal patterns do not change simply due to amplitude variation. Hence there is a real possibility that a seasonal-trend decomposition obtained from this approach may be a viable competitor to that achieved using X-12-ARIMA, but with the obvious advantages of model-based standard errors and forecast functions obtained with appropriate prediction intervals. A more detailed comparison to support this suggestion is the subject of a further paper.

Acknowledgments

Statistics New Zealand kindly supplied the data. We thank those who commented on presentations at Statistics New Zealand, the Reserve Bank of New Zealand, Victoria Management School, the ASC/NZSA 2006 Conference, and the TSEFAR 2006 Conference. We also thank Peter Thomson for some helpful suggestions that improved the paper.

References

Andrews, D.W.K. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica* **59(3),** 817–858.

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

Bai, J. & Perron., P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics* **18,** 1–22.

Bell, W.R. & Hillmer, S.C. (1984). Issues involved with the seasonal adjustment of economic time series. *Journal of Business and Economic Statistics* **2(4),** 291–320.

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

Bruce, A.G. & Jurke, S.R. (1996). Non-Gaussian seasonal adjustment: X-12-ARIMA versus robust structural models. *Journal of Forecasting* **15,** 305–328.

Cleveland, R.B., Cleveland, W.S., McRae, J.E. & Terpenning, I. (1990). STL: A seasonal-trend decomposition procedure based on loess. *Journal of Official Statistics* **6,** 3–73.

Dungey, M., Fry, R. & Martin, V.L. (2004). Currency market contagion in the Asia-Pacific region. *Australian Economic Papers* **43(4),** 379–395.

Franses, P.H. (1998). *Time Series Models for Business and Economic Forecasting.* Cambridge: Cambridge University Press.

Hall, R., Feldstein, M., Bernanke, B., Frankel, J., Gordon, R. & Zarnowitz, V. (2001). The business-cycle peak of March 2001. Technical Report, Business Cycle Dating Committee, National Bureau of Economic Research, USA. Available at http://www.nber.org/cycles/november2001/

Hamilton, J.D. (1994). *Time Series Analysis.* Princeton: Princeton University Press.

Harvey, A.C. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter.* Cambridge: Cambridge University Press.

Haywood, J. & Randal, J. (2005). Structural change and New Zealand visitor arrivals: the effects of 9/11. In *Proceedings of the 40th ORSNZ Conference*, pp. 326–335, Auckland: Operational Research Society of New Zealand (Inc.).

Hoaglin, D.C., Mosteller, F. & Tukey, J.W. (Eds.) (1983). *Understanding Robust and Exploratory Data Analysis.* New York: Wiley.

Kaminsky, G.L. & Schmukler, S.L. (1999). What triggers market jitters? A chronicle of the Asian crisis. *Journal of International Money and Finance* **18,** 537–560.

Macaulay, F.R. (1931). *The Smoothing of Time Series.* New York: National Bureau of Economic Research.

Makridakis, S., Wheelwright, S.C., & Hyndman, R.J. (1998). *Forecasting: Methods and Applications*, 3rd edn. New York: Wiley.

NSW Department of Education (1985). 1987: four term future. *Perspectives: Looking at Education* **8(4)** (April).

Pearce, D. (2001). Tourism. *Asia Pacific Viewpoint* **42,** 75–84.

R Development Core Team (2006). *R: A Language and Environment for Statistical Computing.* Vienna: R Foundation for Statistical Computing.

Statistics New Zealand (2005). *Tourism Satellite Account 2004.* Wellington: Statistics New Zealand.

US Department of State (2004). Three years after 9/11: mixed reviews for war on terror. Available at http://www.globalsecurity.org/security/library/news/2004/09/wwwh40915.htm

Zeileis, A., Kleiber, C., Krämer, W. & Hornik, K. (2003). Testing and dating of structural changes in practice. *Computational Statistics and Data Analysis* **44,** 109–123.