This is a repository copy of *Removing seasonality under a changing regime : Filtering new car sales*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/64524/>

Version: Submitted Version

Article:

Thornton, Michael Alan orcid.org/0000-0002-4470-809X (2013) Removing seasonality under a changing regime : Filtering new car sales. *Computational Statistics & Data Analysis*. 4. pp. 4-14. ISSN 0167-9473

<https://doi.org/10.1016/j.csda.2011.06.021>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

NOTICE: this is the authors version of a work that was accepted for publication in Computational Statistics and Data Analysis. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version is forthcoming in the second Special Issue on Signal Extraction and Filtering 10.1016/j.csda.2011.06.021.

Removing Seasonality Under a Changing Regime: Filtering New Car Sales

Michael A Thornton*

*School of Economics
University of Reading
Whiteknights
PO Box 217
RG6 6AH
United Kingdom*

Abstract

The use of filters for the seasonal adjustment of data generated by the U.K. new car market is considered. U.K. new car registrations display very strong seasonality brought about by the system of identifiers in the U.K. registration plate, which has mutated in response to an increase in the frequency with which the identifier changes, while it also displays low frequency volatility that reflects U.K. macroeconomic conditions. Given the periodogram of the data, it is argued that an effective seasonal adjustment can be performed using a Butterworth lowpass filter. The results of this are compared with those based on adjustment using X-12 ARIMA and model-based methods.

Keywords: Seasonal adjustment, Signal Extraction, Frequency Response, Butterworth Lowpass filter.

1. Introduction

The registration system for new cars in the U.K. has long involved a component signifying the period in which the vehicle was first registered. This feature, which means that cars bought and registered at certain points of the year appear to be newer for longer, has made the pattern of new car

*tel: +44 (0) 118 378 5344, fax: + 44 (0) 118 378 4029.

Email address: m.a.thornton@reading.ac.uk (Michael A Thornton)

URL: <http://www.personal.reading.ac.uk/~qb901444/> (Michael A Thornton)

purchase highly seasonal. To illustrate, August is usually a relatively quiet month in all industries in the northern hemisphere but, because an annual identifier incorporated into registrations changed on August 1, it was the month with the highest number of new car registrations every year in the U.K. for 28 years.

Such strong seasonality is, in itself, a challenge for modellers and analysts. A range of techniques have been developed to remove seasonality from data, including empirical procedures such as X-12 ARIMA, see Findley et al (1998), and model-based procedures such as the TRAMO-SEATS procedure of Gomez and Maravall (1996), applied for example in Maravall (2006), or the structural time series approach, see for example Harvey and Todd (1983) or Hindrayanto et al (2010). In this instance, the task is complicated by a seasonal break.

Over time, it became clear that this August effect was dominating the year. In response, the registration system was modified in 1999 so that the identifier changed every six months, in September and March. It was hoped that this would smooth sales to something more like the patterns seen in continental Europe. In reality the seasonality mutated, over an identifiable interval and in an identifiable way, but did not disappear: the standard deviation in monthly sales over a calendar year has fallen a little as a proportion of average monthly sales but remains at around 60 percent. That mutation transformed both the amplitude and frequency of the seasonal cycles while also inducing a modest phase shift: after a period of annual identifiers, there is one seven month identifier before the new regime of six-monthly identifiers. This is not the type of problem for which traditional seasonal adjustment procedures were designed.

As noted in Ghysels and Osborn (2001) there is a long tradition from model-based seasonal adjustment procedures of treating the removal of seasonality as a signal extraction problem. Given the nature of the mutation in seasonal pattern and the relatively low power displayed by the series at non-trend, non-seasonal frequencies, we also make use of a filter that can be defined in terms of its response to frequencies contained within certain bands. This is the Butterworth lowpass filter as discussed by Pollock (2000), Gomez (2001), Harvey and Trimbur (2003), Pollock (2006) and Proietti (2007). In econometrics it has largely been used for identifying business cycle activity, but has been used by Pollock (2000) for removing seasonal components in Swiss unemployment data. The Butterworth filter has two interpretations. The first, which is fully parametric, as the optimal signal ex-

traction filter for an underlying unobserved components model. The second as a semi-parametric technique to split a series into two parts based on two non-overlapping frequency ranges. With a semi-parametric interpretation, it is set up to cope well with mutations in amplitude and phase, enabling the identification of trend-cycle elements across the full span of the series. Other techniques, which are designed around a fixed or slowly evolving seasonal pattern, would be forced into treating the data generated by each of the regimes somewhat differently.

The underlying trends in U.K. new car data are of interest beyond presenting a methodological challenge. The motor industry is widely seen both as a leading indicator, and as a wider linchpin, of the manufacturing sector, particularly during the recent economic downturn. Following the lead in continental Europe, from May 2009 to March 2010, the U.K. ran a car scrappage incentive scheme, whereby purchasers of new cars received £2,000 toward the cost of a new car if they scrap a vehicle that is more than 10 years old. In total 400,000 purchases benefited from this additional funding. Such schemes and their potential to create further cycles in activity have been analysed from a theoretical viewpoint by, among others, Adda and Cooper (2000).

The paper is split into 5 sections. Section 2 discusses the nature of seasonality and popular methods of seasonal adjustment, making the case for the use of a frequency-specific filter as a part of that tool-kit. Section 3 presents the Butterworth lowpass filter. Section 4 explains the changes that took place in the series and applies the filter to the full sample and to two sub-samples representing the different regimes, comparing the results with those using X-12 ARIMA and TRAMO-SEATS, and section 5 concludes.

2. Seasonal Adjustment

Here we follow the definition in Nerlove (1964, pp 264):

In the more general case, then, we may define seasonality as that characteristic which gives rise to spectral peaks at seasonal frequencies.

Seasonal factors tend to fall into three broad categories, as noted by Hylleberg (1992): climate; convention (including the timing of religious festivals); and repetitive practices (such as tax years, accountancy periods and the U.K. registration plate system). Some of these will remain fixed, some may vary

but are always predictable, and some may offer departures around a roughly regular pattern, causing spectral power to concentrate around rather than appearing directly on the seasonal frequencies.

2.1. Seasonal Adjustment Procedures

Procedures for removing seasonal factors from data fall into two broad camps: empirical and model-based techniques. Empirical techniques use statistical smoothing methods without presupposing that the data are generated by an underlying model. The procedure developed at the U.S. Census Bureau, Variant X-11, see Shishkin et al (1967), is still the major part of adjustment procedures used in most statistical agencies. An excellent treatment of the way the program is assembled can be found in Hylleberg (1992), Ghysels and Perron (1993) and Wallis (1982). The more recent X-12 ARIMA, see Findley et al (1998), builds on X-11, improving diagnostics, the treatment of outliers and enabling the use of ARIMA-generated out of sample values in the smoothing. Such techniques, however, have a number of drawbacks for econometric work. Firstly, the use of moving averages with long lags and leads means that a definitive figure for the adjusted series will not be available for a number of years. Secondly, the procedure provides no insight into what seasonality might be and no framework in which we can examine its relationship with the trend component. Thirdly, these procedures smooth data in connection with the rest of the sample leaving a series of data-points which are no longer independent realisations. The number of degrees of freedom lost will depend on a number of issues including the choice of procedure for smoothing outliers and trading day effects.

Model-based procedures, on the other hand, suppose that the series to be treated can be decomposed into unobserved components, perhaps after taking logarithms. A typical form splits a series (X_t) into trend-cycle (TC_t), seasonal (S_t) and irregular (I_t) components,

$$\begin{aligned} X_t &= TC_t + S_t + I_t, \\ TC_t &= \frac{\gamma(L)}{\delta(L)} \epsilon_t, \\ S_t &= \frac{\psi(L)}{s(L)} v_t \\ I_t &= \frac{\beta(L)}{\alpha(L)} \nu_t, \end{aligned} \tag{1}$$

where L is the lag operator, $LX_t \equiv X_{t-1}$, ν_t , v_t and ϵ_t are mutually independent white noise processes. The autoregressive filters are of the form $\delta(z) = (1 - z)^d$ and $s(z) = 1 + z + z^2 + \dots + z^{s-1}$, where s is the number of data periods per seasonal cycle, while $\alpha(z)$ is of order a and has no roots on or inside the unit circle. The moving average filters, $\gamma(z)$, $\psi(z)$ and $\beta(z)$ have roots outside the unit circle and are of order less than d , $s - 1$ and a respectively. Estimating these components becomes a problem of signal extraction following the work of Wiener and Kolmogoroff and can be implemented by the TRAMO-SEATS programme of Gomez and Maravall (1996) or through the STAMP programme of Koopman et al (2007).

As the individual components cannot be observed, however, the above is not sufficient to identify a unique model. The empirical methods discussed above develop their procedures for seasonal adjustment on differing additional assumptions. Burridge and Wallis (1988) investigate some unobserved component models for which X-11 provides linear least square estimates of the seasonal component. The approach taken in a formative article by Hillmer and Tiao (1982) and developed in papers by Bell and Hillmer (1984) and Maravall (1987) is to make additional assumptions about the nature of the variances as well as the lag polynomials. Their so-called ‘canonical decomposition’ comes from choosing the form of equation (1) that maximises the variance of the process driving the irregular, ν_t , given the assumed forms of the lag polynomials and the constraint that the spectral density of each component is non-zero in the region $[0, \pi]$. In other words, as much of the ‘jerkiness’ in the series as possible is allocated to the irregular.

In contrast the structural time series model of Harvey and Todd (1983), sets about identifying the model by making additional assumptions about the processes in equation (1), including that $\psi(z) = 1$ and the irregular component is white noise ($\alpha(z) = \beta(z) = 1$). This has the disadvantage of consigning any spectral peaks outside the seasonal frequencies to the trend/cycle, but it is easily translated into state-space and estimated using the Kalman filter.

Each of these techniques relies upon assumptions about the relationship between the three components of the series. It should be noted, however, that these techniques are not the only way to achieve the aim of removing spectral power at seasonal frequencies. Indeed if we are faced with a series with relatively little spectral power away from the trend/cycle and seasonal frequencies, so that the irregular is a relatively minor component, then a lowpass filter defined to exclude the seasonal frequency and its harmonics is

an attractive method of adjustment and it is to these we turn in the next section.

3. The Butterworth Lowpass Filter

A lowpass filter, $\psi(\omega; \omega_c) : [0, \pi] \times [0, \pi] \rightarrow [0, 1]$, is one intended to allow through frequencies below ω_c and to exclude those above it. The ideal filter would have the properties:

A) $\psi(\omega; \omega_c) = 1$, $\omega < \omega_c$; and,

B) $\psi(\omega; \omega_c) = 0$, $\omega > \omega_c$.

A promising candidate, given the properties of the tangent function over the range $[0, \pi]$, is

$$\psi(\omega; \omega_c) = \frac{1}{1 + \left[\frac{\tan(\omega/2)}{\tan(\omega_c/2)} \right]},$$

which has

A') $\psi(\omega; \omega_c) > \frac{1}{2}$, $\omega < \omega_c$; and,

B') $\psi(\omega; \omega_c) < \frac{1}{2}$, $\omega > \omega_c$,

but only attains A) when $\omega = 0$ and B) when $\omega = \pi$. It is possible to sharpen the transition between stopband and passband, however. Since the term $[\cdot] < 1$ when $\omega < \omega_c$ and > 1 when $\omega > \omega_c$, an effective way is by taking powers of that term in denominator

$$\psi(\omega; \omega_c) = \frac{1}{1 + \left[\frac{\tan(\omega/2)}{\tan(\omega_c/2)} \right]^{2n}}. \quad (2)$$

Moreover, if $\omega_c \in (0, \pi/2)$, as is highly likely in econometric applications, lowering the cut-off frequency also improves the sharpness of the filter. Consider the frequency response some θ radians after the cut-off frequency, $\psi(\omega_c + \theta; \omega_c)$. Using the identities $\tan(a) = \sin(a)/\cos(a)$, $2\sin(a)\cos(b) = \sin(a+b) + \sin(a-b)$ and $\sin(-a) = -\sin(a)$, it can be shown that,

$$\frac{\tan(\omega_c/2 + \theta/2)}{\tan(\omega_c/2)} = \frac{\sin(\omega_c + \theta/2) + \sin(\theta/2)}{\sin(\omega_c + \theta/2) - \sin(\theta/2)},$$

which is a decreasing function in $\omega_c \in (0, \pi/2)$. It then follows that $\psi(\omega_c + \theta; \omega_c)$ is increasing in ω_c .

Now consider the time domain representation of this filter. Denoting $\lambda = [\tan(\omega_c/2)]^{-2n}$ and remembering that $\tan(\omega/2) = i \frac{e^{-i\omega/2} - e^{i\omega/2}}{e^{-i\omega/2} + e^{i\omega/2}}$, where $i = \sqrt{-1}$, then equation (2) can be written equivalently as

$$\begin{aligned}\psi(\omega; \omega_c) &= \frac{1}{1 + \lambda \frac{(1 - e^{-i\omega})^n (1 - e^{i\omega})^n}{(1 + e^{-i\omega})^n (1 + e^{i\omega})^n}} \\ &= \frac{(1 + e^{-i\omega})^n (1 + e^{i\omega})^n}{(1 + e^{-i\omega})^n (1 + e^{i\omega})^n + \lambda (1 - e^{-i\omega})^n (1 - e^{i\omega})^n},\end{aligned}$$

which is the frequency response function of a bi-directional filter, $\phi(L; \lambda)$ of the form

$$\phi(L; \lambda) = \frac{(1 + L)^n (1 + L^{-1})^n}{(1 + L)^n (1 + L^{-1})^n + \lambda (1 - L)^n (1 - L^{-1})^n}. \quad (3)$$

Such a filter can be recognised as one of the Butterworth family of filters, discussed by Pollock (2000), Gomez (2001) and Harvey and Trimbur (2003). As noted by Pollock (2000), there is a trade-off between the sharpness of the filter and the stability of its output, which should be borne in mind when choosing values for the parameters. The parameter λ determines the sharpness of the filter, increasing as ω_c approaches 0 or as n becomes large. As λ becomes large, however, the filter increasingly resembles a bi-directional integrating filter, which is highly sensitive to initial conditions and numerical rounding. This is equivalent to the poles of the filter heading toward the unit circle at an angle of $\pm\omega_c$ to the real line (see appendix), causing the filter to amplify the cycle at ω_c , leading to numerical instability unit modulus and violating the bounded input bounded output (BIBO) conditions in the limit.

Another member of the family, replaces the tangent with the sine function. This filter will, in the case that $\omega_c \in (0, \pi/2)$, be further from the ideal of A) and B). This can easily be seen by writing $\tan(x) = \sin(x)/\cos(x)$ in equation (2) and noting that $\cos(\omega) > \cos(\omega_c) \quad \forall \omega \in (0, \omega_c)$ and $\cos(\omega) < \cos(\omega_c) \quad \forall \omega \in (\omega_c, \pi/2)$.

These filters can be interpreted as examples of the model-based methods discussed in the previous section, albeit in the absence of seasonal factors. As noted by Gomez (2001), the Butterworth sine filter is the minimum mean square error linear filter for removing a white noise component from a random walk of order n , where now λ is the ratio of the variances of the processes driving the random walk and the noise. In these circumstances the filter will deliver the structural decomposition. The tangent filter, meanwhile, is

the equivalent for an integrated moving average (n, n) process, where the moving average takes the form $\gamma(z) = (1 + z)^n$. In these fully parametric interpretations, n reflects the degree of integration of the series. Unlike the sine filter, the tangent filter has $\psi(\pi; \omega_c) = 0$, making its approach generally canonical in the sense that it would not be possible to remove a larger white noise term from the series, although it does not accomplish the canonical decomposition in the general case.

4. Data and Modelling

4.1. *The Identifier*

The registration system for new cars in the U.K. is unusual in featuring a component signifying the period in which the vehicle was first registered. Starting in January 1963, this component was a letter, changed annually, progressing through the alphabet (omitting letters I, Q, U and Z on grounds of identification). Recognising the impact of such a system, the motoring authorities responded to a call to liven up summer sales and in 1967 the day of the annual letter change was moved to August 1.

As the sales-boom in August, with over one fifth of annual registrations, became the dominant feature of the market, policy makers became increasingly aware of the costs it imposed on the U.K. motor industry. Since most U.K. production was aimed at the domestic market, firms were required to hold large inventory stocks in order to meet the surge in demand, preventing them from taking advantage of more efficient, just in time, production techniques. In contrast most overseas manufacturers find that August is very quiet in their domestic markets and are delighted to tailor their output toward the U.K.. Further costs are incurred during the vehicle's lifetime, as a high proportion of vehicles require annual services and road-worthiness tests at the same time of year, testing the capacity of the car maintenance industry.

In order to counter this the system was modified to one which changed every six months. The final annual identifier, the letter 'R' at the beginning of the registration plate, ran from the start of August 1997 to the end of July 1998. The next identifier, 'S', ran from August 1998 to the end of February 1999 (7 months), and from then on the identifier has changed at the start of every March and September. In case the public started to lose track of the implied vintages, from September 2001, the alphabet exhausted, relatively oblique letters were supplanted by more brazen numbering: 02 for March

2002 and 52 for September 2002. The data, obtained with kind permission from the Society of Motor Manufacturers and Traders, are for monthly new car registrations in the U.K. from January 1970 until December 2010, giving 492 observations. Starting the sample 3 years after the annual August identifier was introduced, allows the seasonal pattern underpinned by that system some time to bed in as fleet buyers adjust the timing of their replacements in line with having the most up-to-date number plates. In fact 1970 is the only year in the sample in which the month with the highest number of registrations does not coincide with a change in the identifier. The figure for registrations, rather than actual sales, is used to capture the phenomenon of the ‘nearly new’ car whereby manufacturers, under the pretext of delivery mileage, are able to impose their excess production on the new car market. These cars have only the slightest distinction from new and so are treated as new throughout the data set.

The data, plotted in figure 1, show a high degree of seasonality, which evolves over the sample. There is also a very clear mutation in the seasonal pattern a little under three-quarters of the way through the sample, between August 1998 and March 1999. At first sight, the data also appear to show a trend which may be linked to the business cycle/ economic conditions. For example, total sales were 1,661,624 in 1973 but only 1,194,088 in 1975; total sales reached 2,300,934 in 1989 but fell sharply to 1,592,326 in 1991. Of particular interest is a dramatic fall in sales, following a long period of growth, from May 2008 and the apparent recovery toward the end of the sample.

The periodogram of the data after taking logarithms, in figure 2, shows power in the trend/cycle frequencies, particularly close to the origin with a further smaller spike at a business cycle frequency around $\pi/30$. Away from these lower frequencies, the periodogram shows a striking concentration around the seasonal frequencies of $\pi/6$, $\pi/3$, $\pi/2$, $2\pi/3$, $5\pi/6$ and π , with relatively little power in between the clusters.

The significance of this can be tested using a basic analysis of variations technique, see Brockwell and Davies (1991). Under the null of no periodicity at a given Fourier frequency, $\omega_j \neq \pi$, the test statistic in a sample of length T is

$$Z = \frac{(T-3)I(\omega_j)}{(SSD - 2I(\omega_j))} \sim F_{2,T-3},$$

where SSD is the sum of squared deviations about the mean and $I(\omega_j) =$

Figure 1: U.K New Car Registrations.

Figure 2: Periodogram of Log U.K New Car Registrations

$\frac{2}{n} |\sum_{t=0}^{T-1} x_t e^{i\omega_j t}|^2$ is the sample periodogram. A similar statistic can be constructed when $\omega_j = \pi$, with only one degree of freedom in the numerator. Table 1 reports the p-value of this statistic for trend, business cycle and seasonal frequencies.

Table 1. Power at Trend/Cycle and Seasonal Frequencies								
Frequency	$\pi/246$	$4\pi/123$	$\pi/6$	$\pi/3$	$\pi/2$	$2\pi/3$	$5\pi/6$	π
Period	492	61.5	12	6	4	3	2.4	2
Z	2.641	0.388	0.507	4.919	1.408	1.173	6.508	0.368
p-value	0.07	0.68	0.60	0.01	0.25	0.31	0.00	0.54

The statistic is in the bottom decile of the distribution for all frequencies outside the trend and seasonal regions, becoming larger at and around all of the seasonal frequencies and highly significant at $\pi/3$ and $5\pi/6$. The relatively low values for some of the larger seasonal frequencies reflect the fact that the peak of the density is not exactly on the harmonic. The peak around frequencies $2\pi/3$ and π actually occur the ordinate before, with values of 3.2313 (p value 0.04) and 2.6519 (0.07) respectively. As figure 2 would suggest, the regions around $\pi/2$ and $\pi/6$ are the only seasonal regions not to have an ordinate that is significant at the 10 per cent level.

4.2. Filtering the Data

We now compare the results of filtering the full sample by X-12 ARIMA, TRAMO-SEATS and using the Butterworth Lowpass filter. A log-additive model is chosen under all of the likelihood tests in X-12 ARIMA and is also chosen by the automatic model selection criteria implemented in TRAMO. Figure 3a shows the adjustment of the data by X-12 ARIMA, including trading day adjustments, while figure 3b shows the irregular.

The irregular is heavily influenced by the change in regime. Including this in the seasonally adjusted series causes it to mimic the new seasonal pattern for a period starting in the late 1990's. There is also some evidence of struggling to adapt to the evolving seasonality in the mid-1970's with a particularly large values positive value in September 1975 being followed by a particularly large values negative value 12 months later, itself followed by large positive values in November and December.

Following pre-adjustment by TRAMO, SEATS selects a seasonal ARIMA (3,1,1)(0,1,1). Figure 3c shows the adjustment of the data by TRAMO-SEATS while figure 3d shows the irregular. The change in seasonal regime is heavily reflected in both irregulars, but X-12 ARIMA also appears to capture

(a) X-12 ARIMA: Seasonally Adjusted

(b) X-12 ARIMA: Irregular

(c) TRAMO-SEATS: Seasonally Adjusted

(d) TRAMO-SEATS: Irregular

Figure 3: Seasonal adjustment of U.K New Car Registrations 1970-2010.

Figure 4: Properties of the Butterworth Lowpass Filter with $n=5$ and $\omega_c = \pi/12$.

some of the seasonal pattern at other times, suggesting it is imposing a more slowly evolving seasonal component. Not surprisingly, this is picked up in X-12 ARIMA's monitoring and quality assessment statistics, with three statistics falling in the failing 1–3 region, those for: the amount of month to month change in the irregular component as compared to the amount of month to month change in the trend/cycle (M3), with a value of 1.570; the amount of autocorrelation in the irregular as described by the average duration of run (M4), with a value of 1.095; and, the test for the number of months it takes the change in the trend/cycle to surpass the amount of change in the irregular (M5), with a value of 2.169.

We now turn to the Butterworth lowpass Filter. We must first decide the order and cut-off of the filter. Both may be informed by an underlying model, with n being the degree of integration and ω_c depending on that and the ratio of the variances of the trend and white noise irregular, but here our treatment is semi-parametric. The order of $n = 5$ gives a reasonably quick transition between stop and pass bands, as seen in figure 4a. Since there is relatively little power in the region between the business cycle frequencies and $\pi/6$, we can choose a cut-off frequency a little below $\pi/6$, at $\pi/12$, 15 degrees, without paying a significant cost in excluding trend frequencies. This delivers a filter that is 'stiffer' in terms of excluding seasonal elements, without the increased instability that a greater n would bring. As a check on stability, the poles and zeros of this filter are drawn in figure 4b.

Figure 5: Butterworth Filter with $n = 5$ and $\omega_c = \pi/12$. applied to U.K New Car Registrations 1970-2010.

To aid comparison with the log-additive models, the filter is applied to the data after taking logarithms before converting the estimates of the components back for comparison with the original series.

In comparing figure 5a with figure 3a and 3c it is clear that, in this application, the Butterworth filter does not allow short-run effects from the change in regime to dominate the seasonally adjusted series. In both X-12 ARIMA and TRAMO-SEATS, the change to the new seasonal pattern is captured by the irregular, which then feeds into the seasonally adjusted series. A better comparison is between the Butterworth and the trend/cycle components estimated by X-12 ARIMA, shown in figure 6a and by TRAMO-SEATS, shown in figure 6b. The performance of the filters in isolating trends is broadly similar, despite the simpler nature of the Butterworth filter. Its tendency to isolate low frequencies, however, causes it to understate the impact on registrations of the financial crisis from the end of 2007, prior to the introduction of the scrappage scheme, compared to both X-12 ARIMA and TRAMO-SEATS. Moreover, the Butterworth filter ends on an upward trend, perhaps buoyed by memory of the scrappage scheme, while the other filters end with a downward trend.

4.3. *Splitting the Sample*

The poor performance against X-12 ARIMA's quality assessment statistics and the increase in volatility in the irregular as the regime changes, suggests that investigation of the two periods separately may be worthwhile.

Figure 6: Trend/Cycle Components of U.K New Car Registrations 1970-2010.

The series has a very clear break somewhere between August 1998 and March 1999. As the change in regime was widely anticipated it is possible that it began to influence the pattern prior to its implementation. This can be seen in the data where anticipation of the first March identifier depresses the number of registrations in January and February 1999 to 78 per cent and 47 per cent respectively of those in the previous year. Here we regard January 1999 as the break-point. In reality, the transition between regimes will have been smoother, with the prospect of a new regime exerting a growing influence in the months running up to March 2000, but allocating four of the seven months between identifier changes to the old regime is not unreasonable. The periodogram corresponding to the annual regime is in figure 7a, while that corresponding to the six month regime is in figure 7b.

The new regime has a noticeably different pattern, while still displaying a high level of seasonality. The six month cycle is reflected in greatly diminished power at $5\pi/6$ and at other harmonics of $\pi/6$ that are not also harmonics of $\pi/3$. Power remains concentrated on the seasonal frequencies under each regime as over the combined period, reinforcing the assertion that a lowpass filter will perform a passable seasonal adjustment.

Figure 8 shows the results of applying X-12 ARIMA, TRAMO-SEATS and the Butterworth lowpass filter over the two sub-samples separately and splicing the results together.

Again log-additive models were chosen by both methods. Splitting the

Figure 7: Sub-sample Periodograms of Log U.K New Car Registrations.

sample improves the performance against X-12 ARIMA's monitoring and quality assessment statistics markedly. All statistics are met over the first period, although M3 and M5 remain close to 1. Over the second period, the procedure fails M5 very marginally, with a value of 1.034, and comes close on M3, with a value of 0.912. This suggests discomfort with the volatility of the irregular, which is perhaps not surprising given the sudden slump in registrations following the financial crisis, reflected in figure 8b. TRAMO-SEATS chooses a simpler ARIMA(0,1,1)(0,1,1) 'Airline' model for both regimes, which are estimated as unobserved components models as in 1 with $d = 2$ and $\alpha(z) = \beta(z) = 1$. The chosen parameters are similar, but not identical, across the regimes with moving average roots for the seasonally adjusted series of 0.676 and 0.926 over the first regime and 0.712 and 0.900 over the second. Not surprisingly, there is no longer a jump in volatility of the irregular around the time of the regime shift.

All three filters provide reasonably similar de-seasonalised estimates. X-12 ARIMA calculates a large downward irregular for December 1999, which the other methods ignore. TRAMO-SEATS identifies a large negative irregular late in 2007, and so produces larger estimates of the effect of the financial crisis than the other two.

5. Summary and Conclusions

The adjustment of data on U.K. new car registrations, a series characterised both by the strength of seasonal pattern and its mutation part way

(a) X-12 ARIMA: Seasonally Adjusted

(b) X-12 ARIMA: Irregular

(c) TRAMO-SEATS: Seasonally Adjusted

(d) TRAMO-SEATS: Irregular

(e) Butterworth Lowpass

(f) Butterworth Highpass

Figure 8: Seasonal adjustment of U.K New Car Registrations 1970-98, 99-2010.

through the sample, has been considered. A semi-parametric method of seasonal adjustment based on the Butterworth filter was elaborated. This method can be interpreted parametrically as a special case of the canonical decomposition. Taking a less parametric view, however, this filter enables a series with a relatively small irregular error to be split into two components in an intuitive way. The filter was applied to the data, as a whole and after splitting into two sub-samples, and the results were compared with those given by X-12 ARIMA and TRAMO-SEATS.

All filters picked out the main macroeconomic trends underpinning the series, although there was some divergence over particular events, such as the size of the effect of the recent financial downturn and the direction of the trend. Further investigation of the relationship between the series and macroeconomic variables would be interesting. X-12 ARIMA and TRAMO-SEATS tended to attribute the shift in seasonal regime to the irregular and performed far better once the sample was split. Splitting the sample, of course, entails a loss of information on trend/cycle components that run across the break. Model-based procedures designed to maintain the trend/cycle while the seasonal changes are worthy of investigation with this series.

Acknowledgements

I am extremely grateful to the Associate editor and two anonymous referees, whose suggestions have improved the paper considerably. I am most grateful to D.S.G. Pollock, for his supervision of the original incarnation of this paper. The procedures outlined were realised using gretl. I am also very grateful to the Society of Manufacturers and Traders for making this time series available. All errors remain my own.

Appendix A. The poles and zeros of the Butterworth filter

The zeros of the filter, which can be identified from equation (3), all lie on the unit circle at the point -1 , as would be expected from condition C).

The poles of the filter are the solutions to the expression

$$\begin{aligned} & (1+z)^n(1+z^{-1})^n + \lambda(1-z)^n(1-z^{-1})^n \\ &= 1 + \lambda \left(i \frac{1-z}{1+z} \right)^{2n} = 0. \end{aligned} \quad (\text{A.1})$$

Defining $s(z) = \left(i \frac{1-z}{1+z} \right)$, it is clear that $\tilde{z}$ is pole if and only if $s(\tilde{z})$ is one of the $2n$ roots of $-1/\lambda$. This leads to $2n$ values of s equally spaced around the circle of radius $\lambda^{-1/2n} = \tan(\omega_c/2)$. Denoting this radius as τ , we have the poles of s at

$$s(z) = \tau \exp \left\{ \frac{i\pi j}{2n} \right\}, \quad j = 1, 3, 5, \dots, 2n-1, 2n+1, \dots, 4n-1.$$

Inverting the function $s = s(z)$ gives

$$z = \frac{i-s}{i+s} = \frac{(i-s)(i+s)^*}{(i+s)(i+s)^*} = \frac{1+i(s+s^*)-ss^*}{1-i(s-s^*)+ss^*},$$

where the symbol $*$ denotes the complex conjugate. Using the Euler equations it follows that at the $2n$ poles

$$\begin{aligned} s + s^* &= \tau \left[\exp \left\{ \frac{i\pi j}{2n} \right\} + \exp \left\{ -\frac{i\pi j}{2n} \right\} \right] = 2\tau \cos \left\{ \frac{\pi j}{2n} \right\}, \\ s - s^* &= \tau \left[\exp \left\{ \frac{i\pi j}{2n} \right\} - \exp \left\{ -\frac{i\pi j}{2n} \right\} \right] = 2i\tau \sin \left\{ \frac{\pi j}{2n} \right\}, \\ ss^* &= \tau^2 \left[\exp \left\{ \frac{i\pi j}{2n} \right\} \exp \left\{ -\frac{i\pi j}{2n} \right\} \right] = \tau^2. \end{aligned}$$

Hence the $2n$ poles are given by

$$z = \frac{1 + 2i \cos \left\{ \frac{\pi j}{2n} \right\} \tau - \tau}{1 + 2 \sin \left\{ \frac{\pi j}{2n} \right\} \tau + \tau^2}, \quad j = 1, 3, 5, \dots, 4n-1. \quad (\text{A.2})$$

The filter is BIBO stable provided the values of z satisfying equation (A.1), which correspond to the first n roots of $s(z)$ (that is for $j = 1, 3, 5, \dots, 2n-1$) lie within the unit circle; correspondingly those of z^{-1} lie outside. From

now on we concentrate on the first set of poles. It follows from (A.2) that when these are expressed in polar form, as $\rho_j \exp(i\theta_j)$, then

$$\rho_j = \frac{\sqrt{(1 - \tau^2)^2 + (2\tau \cos \{\frac{\pi j}{2n}\})^2}}{1 + \tau^2 + 2\tau \sin \{\frac{\pi j}{2n}\}},$$

and

$$\theta_j = \arctan \left\{ \frac{2\tau \cos \{\frac{\pi j}{2n}\}}{1 - \tau^2} \right\}, \quad j = 1, 3, 5, \dots, 2n - 1.$$

The two poles that have the largest modulus are those when $j = 1$ and $j = 2n - 1$, which are reflections of one another in the real line. As n becomes large their respective values for $\cos \{\frac{\pi j}{2n}\}$ head toward 1 and -1 , and for $\sin \{\frac{\pi j}{2n}\}$ head toward zero. This results in these poles having unit modulus at angles of ω_c and $-\omega_c$ with the real line, following application of the tangent half-angle formula. As ω_c and hence τ approaches zero these, and all other, poles converge to a common point on the unit circle at $+1$.

Adda, J., Cooper, R., 2000. Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies. *J. Politi. Econom.*, 108, 778–806.

Bell, W. R., Hillmer, S. C., 1984. Issues involved with the Seasonal Adjustment of Economic Time Series. *J. Business Econom. Statist.*, 2, 291–320.

Brockwell, P. J., Davies, R. A., 1991. *Time Series: Theory and Methods* 2nd edition. Springer Verlag, New York. pp. 335.

Burridge, P., Wallis, K. F., 1988. Unobserved Components Models for Seasonal Adjustment Filters. *J. Business Econom. Statist.*, 2, 350–9.

Findley, D. F., Monsell, B. C., Bell, W. R., Otto, M. C., Chen, B., 1998. New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program. *J. Business Econom. Statist.*, 16, 127–177.

Franses, P. H., 1996. *Periodicity and Stochastic Trends in Economic Time Series*. Oxford University Press. Oxford.

Gomez, V., Maravall, A., 1996. Programs TRAMO and SEATS, Instructions for the User. Bank of Spain Working Paper No 9628.

- Gomez, V., 2001. The Use of Butterworth Filters for Trend and Cycle Estimation in Economic Time Series, *J. Business Econom. Statist.*, 19, 365–373.
- Ghysels, E., Perron, P., 1993. The Effect of Seasonal Adjustment Filters on Tests for Unit Root. *J. Econom.*, 55, 57–98.
- Ghysels, E., Osborn, D. R., 2001. *The Econometric Analysis of Seasonal Time Series*. Cambridge University Press, Cambridge, pp. 108.
- Harvey, A. C., Todd, P. H., 1983. Forecasting economic Time Series with Structural and Box-Jenkins models: A Case Study. *J. Business Econom. Forecasting*, 1, 299–307.
- Harvey, A. C., Trimbur, T. M., 2003. General Model-Based Filters for Extracting Cycles and Trends in Economic Time Series. *Review of Econom. Statist.* 85, 244–255.
- Hilmer, S. C., Tiao, G. C., 1982. An ARIMA Model-Based Approach to Seasonal Adjustment. *J. Amer. Statist. Assoc.*, 77, 63–70.
- Hindrayanto, I., Koopman, S.J., Ooms, M., 2010. Exact maximum likelihood estimation for non-stationary periodic time series models. *Computational Statistics and Data Analysis*, 54, 2641–2654.
- Hylleberg, S. (ed), 1992. *Modelling Seasonality*. Oxford University Press. New York.
- Koopman, S. J., Harvey A. C., Doornick J. A., Shephard N., 2007. *STAMP 8.0: Structural Time Series Analyser Modeller and Predictor: The Manual*. Timberlake Consultants Press, London.
- Maravall, A., Pierce, D. A., 1987. A Prototypical Seasonal Adjustment Model. *J. Time Ser Anal.*, 8, 177–193.
- Maravall, A., 2006. An application of the TRAMO-SEATS automatic procedure; direct versus indirect adjustment. *Computational Statistics and Data Analysis*, 50, 2167–2190.
- Nerlove, M., 1964. Spectral Analysis of Seasonal Adjustment Procedures. *Econometrica*, 32, 241–86.

- Pollock, D. S. G., 1998. Time Series Analysis, Signal Processing and Dynamics, Academic Press, London.
- Pollock, D. S. G., 2000. Trend Estimation and Detrending via Rational Square Wave Filters, *J. Econom.*, 99, 317–334.
- Pollock, D. S. G., 2006. Econometric methods of signal extraction. *Computational Statistics and Data Analysis*, 50, 2268–2292.
- Pollock, D. S. G., 2007. Wiener–Kolmogorov Filtering, Frequency-Selective Filtering and Polynomial Regression. *Econometric Theory*, 23, 71–88.
- Proietti, T., 2007. Signal extraction and filtering by linear semiparametric methods. *Computational Statistics and Data Analysis*, 52, 935–958 .
- Shishkin, J., Young, A. H., Musgrave, J. C., 1967. The X-11 Variant of the Census Method II Seasonal Adjustment Program. Technical paper No. 15. Bureau of the Census. Washington DC: Department of Commerce.
- Wallis, K. F., 1982. Seasonal Adjustment and Revisions of Current Data. *J. Royal Statist. Soc.*, 145, 74–85.