RESEARCH REPORT SERIES
(Statistics #2008-12)

A Modified Model-based Seasonal Adjustment that Reduces Spectral Troughs and Negative Seasonal Correlation

Tucker McElroy

Statistical Research Division
U.S. Census Bureau
Washington, D.C. 20233

Report Issued: November 21, 2008

A Modified Model-based Seasonal Adjustment that Reduces Spectral Troughs and Negative Seasonal Correlation

Tucker McElroy*

Abstract

The Wiener-Kolmogorov (WK) signal extraction filter, extended to handle nonstationary signal and noise, has minimum Mean Square Error (MSE) among filters that preserve the signal's initial values; however, the stochastic dynamics of the signal estimate typically differ substantially from that of the target. The use of such filters, although widespread, is observed to produce dips in the spectrum of the seasonal adjustments of seasonal time series. These spectral troughs tend to correspond to negative autocorrelations at lags 12 and 24 in practice, a phenomenon that will be called "negative seasonality." So-called "square root" WK filters were introduced by Wecker in the case of stationary signal and noise, to ensure that the signal estimate shared the same stochastic dynamics as the original signal, and thus remove the problem of spectral dips. This represents a different statistical philosophy: not only do we want to closely estimate a target quantity, but we desire that the internal properties and dynamics of our estimate closely resemble those of the target. The MSE criterion ignores this aspect of the signal extraction problem, whereas the square root WK filters account for this issue at the cost of accruing additional MSE. This paper provides empirical documentation of negative seasonality, and provides matrix formulas for square root WK filters that are appropriate for finite samples of nonstationary time series. We apply these filters to produce seasonal adjustments without inappropriate spectral troughs.

Keywords. ARIMA, Seasonality, Signal Extraction, Wiener-Kolmogorov.

Disclaimer This paper is released to inform interested parties of ongoing research and to encourage discussion of work in progress. The views expressed on statistical, methodological, technical, and operational issues are those of the author and not necessarily those of the U.S. Census Bureau.

*Statistical Research Division, U.S. Census Bureau, 4600 Silver Hill Road, Washington, D.C. 20233-9100, tucker.s.mcelroy@census.gov

1 Introduction

The principal task of seasonal adjustment methodology is to remove seasonality. First, seasonality must be defined; given a precise mathematical definition, one can then devise statistical methods to remove the seasonality – see Bell and Hillmer (1984) for a related discussion. This paper defines seasonality as follows: a time series has seasonality if its sample autocorrelation function (ACF) has significant values at multiples of the seasonal period (e.g., for monthly data at lags 12, 24, etc.), *or* if a spectral estimate has noticeable spectral peaks/troughs at seasonal frequencies. More will be said about this definition below, but here we note our main results: typical seasonal adjustment (SA) filters (i.e., the Mean Square Error (MSE) optimal model-based filters) actually leave some seasonality – negative seasonality – behind, and this is by design. Secondly, it is possible to design model-based filters that completely remove negative seasonality, with little extra cost in terms of MSE. These so-called "Dynamic-Matching filters" (DM) are nearly as easy to construct and implement as MSE optimal filters, and are the main result of this work.

Our definition of seasonality is designed to be practical and empirical. One can easily check the ACF plot for a time series, and the confidence bands can be interpreted as a test against a null hypothesis of white noise, with significant ACF values indicating a significant departure from zero correlation (one can also consider more complicated null hypotheses, with corresponding confidence bands determined by Barlett's formula – see Section 7.2 of Brockwell and Davis (1991)). Consider the U.S. Retail Sales of Shoe Stores series, with regression effects such as trading day removed, from 1984 to 1998, which we refer to as the Shoe series. This monthly series is plotted in Figure 1, together with its SA component (estimated using MSE optimal filters obtained from a fitted Box-Jenkins airline model – Box and Jenkins, 1976) and the corresponding ACF plot of the twice-differenced SA component. Note the significant (negative) ACF values at lags 12 and 24. We also display the AR spectrum estimate of the SA component and twice-differenced SA component in Figure 2; the spectral troughs at the seasonal frequencies $\pi j/6$ ($j = 1, 2, 3, 4, 5$) are more noticeable in the right-hand panel. The ACF and spectral plots together can be taken as evidence of residual seasonality.

In section 2 we discuss the following relationship: seasonal peaks correspond to positive seasonal autocorrelations (i.e., values of the ACF at lags 12, 24, etc.), whereas seasonal troughs correspond to negative seasonal autocorrelations. This latter feature of a time series will be called "negative seasonality," and it can be seen in Figure 2. Negative seasonality is common in MSE optimal model-based seasonal adjustment, which is an assertion that we document in section 2. Note that Nerlove (1964) and Grether and Nerlove (1970) contained the first published observation of the spectral trough phenomena in seasonal adjustments; at the time, this was thought to be a defect of MSE optimal filtering. It is now well-known that MSE optimal model-based filters (referred to as WK filters) typically induce spectral troughs – see Bell and Hillmer (1984) for a historical discussion. In the seasonal adjustment community, this is not viewed as a problem, because the SA estimate is MSE optimal for its target given the correctness of the model. However, this optimality is considered time point by time point, and the inherent dynamics of the estimated SA component *considered as a whole* is not accounted for in the criterion function of the MSE approach. Thus, although the signal estimate is MSE optimal, its internal dynamics (which can be assessed through spectral estimates) do not correspond to those of the target signal – and this is by design. A "dynamic-matching" approach, such as that proposed by Wecker (1979), sacrifices MSE optimality for producing an estimate whose dynamics match those of the target process. (Although, as noted in Findley and Martin (2006), when signal and noise have pseudo-spectra with disjoint support, the WK filters will be dynamic-matching.)

We therefore argue the following: (1) seasonality typically remains in MSE optimal SA estimates, in the form of negative seasonality; and (2) such negative seasonality should be removed in order to provide adequate seasonal adjustments. Note that the failure of MSE optimal SA filters to completely remove seasonality is not their fault *per se* – they do what they are designed to do. They are just misapplied: we suggest that MSE alone is not a proper criterion for signal extraction *if we are interested in estimating an entire component* rather than just one time point; we should be interested in an estimate whose statistical properties match those of the target process. Thus, estimated trends should be smooth to the same degree as the target trend processes themselves (which are described through a model), and estimated irregulars should be truly uncorrelated rather than having strange autocorrelation patterns, etc. Figure 3 illustrates the result of using the DM filters: the SA component is somewhat less smooth than the WK estimate, but the negative seasonality is reduced according to the ACF plot.

The DM filters of this paper can be used to estimate a large array of signals – at least as wide as the MSE optimal theory allows. Thus we can treat trends, seasonals, irregulars, SAs, and cycles. It is shown in section 3 that DM filters produce estimates with approximately the same ACF structure as the target signal; this approximation is assessed in practice through looking at a large suite of seasonal time series. We also assess the loss in terms of additional MSE accrued through use of these alternative filters, as well as through phase delay. We argue that DM filters form a tractable, appealing, and easily implementable alternative to traditional WK filters.

The paper is organized as follows: section 2 contains a discussion of the concept of negative seasonality, which is empirically illustrated through a suite of Foreign Trade series from the U.S. Census Bureau. In section 3 we present the mathematical treatment of DM filters, and in section 4 these methods are implemented and applied to the same Foreign Trade series, where comparisons between the WK and DM methods can be readily made. Section 5 concludes and some technical results are contained in the Appendix.

2 Negative Seasonality

Although seasonality, as we have broadly defined it in section 1, is characterized by significant ACF values at lags 12 and 24 (or seasonal autocorrelation), there is a dichotomy in the sign of the ACF. Roughly speaking, positive seasonal autocorrelation can be associated with a seasonal spectral peak, whereas negative seasonal autocorrelation is associated with a seasonal spectral trough. This concept is illustrated through Figure 4, where the weight function $\cos 12\lambda$ is plotted together with two spectral densities f_1 and f_2, with spectral peaks/troughs respectively at the seasonal frequencies. For a stationary process, the lag 12 autocovariance is

$$\gamma(12) = \frac{2}{2\pi} \int_0^\pi f(\lambda) \cos(12\lambda) d\lambda,$$

i.e., a weighted integral of the spectrum f with weights given by $\cos 12\lambda$. As Figure 4 demonstrates, this function gives positive weights in a short band about each seasonal frequency $\pi j/6$, and negative weights to frequencies in-between the seasonals. It follows that if f has a spectral peak at a seasonal frequency – i.e., the values of f are larger in a neighborhood around $\pi j/6$ as compared to further out – then this gives a positive contribution to $\gamma(12)$; conversely, a spectral trough produces a negative contribution to $\gamma(12)$. Adding up over the six seasonals may produce some cancelation, but if they are all peaks or all troughs, then it follows that $\gamma(12)$ will be positive or negative respectively.

The same effect follows for $\gamma(24)$, only the weight function $\cos 24\lambda$ is much more oscillatory, focusing in on a narrower band of frequencies about each $\pi j/6$ (so that broad peaks/troughs get washed out, and only sharp peaks/troughs get picked up). It is in this sense that positive seasonal autocorrelation is associated to seasonal spectral peaks, while negative seasonal autocorrelation is associated to seasonal spectral troughs. Traditionally, only the former is deemed "seasonality," since the typical nonstationary seasonality evident in economic data manifests as a spectral peak of infinite height in estimates of the pseudo-spectrum. However, from the perspective that any significant seasonal autocorrelations are indicative of seasonality broadly defined, we should also be concerned with spectral troughs in raw or seasonally adjusted data. In order to provide some clarity throughout the rest of this section, we propose the following terms *positive seasonality* and *negative seasonality* to distinguish two types of seasonality. Since the concept of a spectral peak or trough can be somewhat nebulous, given the difficulty of obtaining reliable spectral estimates (and the subjectivity inherent in defining how close neighboring frequencies should be), we take as our empirical definition that a significantly positive/negative seasonal autocorrelation (at any of the seasonal lags) indicates the presence of positive/negative seasonality. (The statistical significance level can be determined by choosing a white noise null hypothesis, or more realistically by postulating some moving average model and computing the confidence bands via Bartlett's formula.)

Another feature that is apparent from Figure 4 is that a spectrum with seasonal troughs – like

f_2 – can also be characterized as having peaks at the intra-seasonal frequencies, i.e., $\pi/12$, $3\pi/12$, $5\pi/12$, $7\pi/12$, $9\pi/12$, and $11\pi/12$. Empirically, we have the same finding in many cases: WK SAs induce seasonal spectral troughs, and at the same time induce intra-seasonal spectral peaks (or regions of heightened spectral mass). Hence, an alternative viewpoint to the phenomenon of negative seasonality is that such series have heightened (positive correlation) periodicities at the intra-seasonals. In this case, the two perspectives – namely that we have six seasonal troughs or that we have six intra-seasonal peaks – are literally six of one and a half dozen of the other, in terms of how we explain negative seasonality. Note that the periods associated with intra-seasonals are (in months) 24, 8, 24/5, 24/7, 24/9, and 24/11 respectively, i.e., in 2 years the respective phenomena occur once, thrice, 5 times, 7 times, 9 times, and 11 times. Taking the lowest common denominator, we see that intra-seasonality is characterized by phenomena occurring repeatedly *every two years*; this is in contrast to seasonality, which occurs every year (of course, some components of seasonality and intra-seasonality occur more frequently). Thus, the presence of negative seasonality will be accompanied, in practice, by hidden two-year periodicities identified through significant (typically negative) lag 24 autocorrelations.

Next, we investigate the WK SAs of the Foreign Trade series (Imports and Exports) – from January 1989 through November 2001 – checking for negative seasonality. There are 19 Import series (denoted with an "m" prefix) and 20 Export series (with an "x" prefix), with the results displayed in Tables 1 and 2 respectively. The SAs were obtained in the following manner: each log-transformed series was first regression adjusted for trading-day effects and outliers using X-12-ARIMA, and then to each series a Box-Jenkins airline model was fitted, with the component models determined via canonical decomposition (Hillmer and Tiao, 1982). The reason this model was selected was for uniformity of procedure, in order to facilitate comparisons (for 35 out of the 39 series the parameter estimates were adequate, but for four series a seasonal moving average unit root was obtained, indicating perfectly periodic seasonality; such a case, though, is not a problem for the algorithm). Once the component models are determined, it is a simple matter to obtain WK SAs via matrix formulas (McElroy, 2008). Tables 1 and 2 provide the lag 12, 24, and 36 sample ACFs for the twice-differenced SAs of each series, with an asterisk if they are significant at the 5% level (we did not take multiple comparisons – examining the ACF at three separate lags – into account when determining the significance level)[1].

Out of the 40 series considered (the Imports, Exports, and Shoe series), 24 (60 %) exhibit

[1]Significance levels were determined as follows: the "true" seasonally adjusted component is supposed to follow a trend plus irregular model, which by the canonical decomposition method for an airline model is an IMA(2,2). Since the SAs are twice-differenced, it suffices to consider the MA(2) portion of the model; its correlations were computed and plugged into Bartlett's formula for the variance of the sample ACF (formula 7.2.5 of Brockwell and Davis, 1991). Standard errors assume a Gaussian distribution. Hence significant values can be interpreted as a rejection of the hypothesis that the twice-differenced seasonally adjusted estimate follows the given MA(2) model for the trend-irregular.

significant negative seasonality (in the WK SAs). This result is unsurprising, but seems to be worth documenting in order to prove our point. We will return to these same series later in section 4, where we find that negative seasonality can be greatly reduced in SAs by using DM filters.

Why do WK filters produce negative seasonality? If we let f_{SA} denote the pseudo-spectrum of the (true) SA component, as determined by canonical decomposition, then we generally have

$$f_{SA}(\lambda) = g(\lambda)|1 - e^{-i\lambda}|^{-2d},$$

where d is the order of trend integration, and is typically 1 or 2 (e.g., 2 for the Airline model); here g is a bounded function. Let $U(z) = 1 + z + z^2 + \cdots + z^{11}$ be the annual summation operator, so that the seasonal pseudo-spectrum is given by $f_S(\lambda) = k(\lambda)|U(e^{-i\lambda})|^{-2}$, for a bounded function k. Then the data pseudo-spectrum is

$$f_y = f_{SA} + f_S = \frac{f_w}{|U(e^{-i\cdot})|^2 |1 - e^{-i\cdot}|^{2d}}$$

for a bounded function f_w. Then the SA component estimate – assuming a bi-infinite sample and the truth of our component models – has pseudo-spectrum

$$\frac{f_{SA}^2(\lambda)}{f_y(\lambda)} = \frac{g^2(\lambda)|U(e^{-i\lambda})|^2}{|1 - e^{-i\lambda}|^{2d} f_w(\lambda)}.$$

The factor $|1 - e^{-i\lambda}|^{2d}$ in the denominator gives the appropriate poles at the trend frequency of zero, but the factor $|U(e^{-i\lambda})|^2$ in the numerator generates zeroes in the pseudo-spectrum at the seasonal frequencies, i.e., spectral seasonal troughs. Thus the annual summation operator is the culprit that generates negative seasonality. The ratio of the SA estimate pseudo-spectrum to the true SA pseudo-spectrum – a measure of failure of the WK procedure to match dynamics – is $f_{SA}/f_y = g|U(e^{-i\cdot})|^2/f_w$. This function multiplies the target SA pseudo-spectrum, producing troughs at the seasonal frequencies (among other distortions caused by multiplication by g and division by f_w). We can further illustrate the effect by specializing our example below.

Suppose further that $d = 1$ and $g(\lambda) = |1 - \theta e^{-i\lambda}|^2$, so that the model for the differenced SA component corresponds to an MA(1) with unit innovation variance and parameter $\theta > 0$. Then the numerator of f_{SA}^2/f_y corresponds to the spectral density of the MA process $(1 - \theta B)^2 U(B)\epsilon_t$, where ϵ_t is unit variance white noise. Ignoring (for purposes of illustration) the contribution of f_w to the autocovariance of the once-differenced SA estimate, we find that the lag 12 ACF is given by $-2\theta(1 - \theta + \theta^2)$, which is always negative since θ is positive. Of course, this is an overly simplistic illustration: f_w will contribute the ACF, θ may be negative, we more often have $d = 2$, and the pseudo-spectrum of the SA component does not in general have the given form due to finite-sample effects. The real "proof" of the existence of negative seasonality is given by the above empirical results, and the preceding discussion only serves as a heuristic explanation.

3 Mathematical Treatment of Dynamic-Matching Filters

In this section we discuss the DM filter, which is a $n \times n$ matrix that left-multiplies the data vector $Y = \{Y_1, \cdots, Y_n\}'$. Section 3.1 sets out some basic conventions, and section 3.2 gives explicit formulas for DM filter.

3.1 Defining the Component Models

We suppose that the signal and noise processes are ARIMA, with differencing operators $\delta^S(z)$ and $\delta^N(z)$ respectively. These polynomials include only unit roots, and are assumed to be of order d_S and d_N respectively. A crucial assumption is that the two polynomials share no common factors; in practice this is easily accomplished as follows. For the canonical decomposition approach to component modeling, we have in mind an ARIMA model for the data of the form

$$\delta(B)Y_t = \Psi(B)\epsilon_t =: W_t$$

where $\Psi(z)$ is a rational function (with no poles on the unit circle), and ϵ_t is white noise. Since $Y_t = S_t + N_t$, it follows that each factor of $\delta(B)$ must appear as a "left-hand operator" in the ARIMA equation for either S_t or N_t (or both). Making *a priori* allocations of the factors of $\delta(z)$ to either the signal or the noise constitutes part of the definition of the components; we can choose to do this in such a way that no factors are shared. This is sensible too, since the left-hand operators $\delta^S(z)$ and $\delta^N(z)$ serve to define some of the key dynamics of the signal and noise processes, so that making the operators distinct serves to separate the components and assist in making them identifiable (we use this term in a broad sense here). We consider an example below.

Suppose that a time series has the fitted ARIMA model $(2,1,3)(0,1,1)_{12}$ given by

$$(1 - 2\rho\cos\omega B + \rho^2 B^2)(1-B)^2 U(B) Y_t = \Theta(B)\epsilon_t,$$

where ρ and ω control the strength and location respectively of a cycle, $U(B) = 1+B+B^2+\cdots+B^{11}$ is the annual summation operator associated with nonstationary seasonality, $\Theta(z)$ is an order 15 polynomial with zero coefficients at certain particular lags, and ϵ_t is white noise. If we are interested in suppressing seasonality, then we naturally let $\delta^N(B) = U(B)$ – since this is associated with the seasonal frequencies – and $\delta^S(B) = (1-B)^2$, which corresponds to the trend frequencies. If instead we want cycle estimation, then $\delta^N(B) = (1-B)^2 U(B)$ and $\delta^S(B) = 1$; the $(1 - 2\rho\cos\omega B + \rho^2 B^2)$ operator will be an autoregressive operator defining the stationary signal component. If we wish to detrend the series, then $\delta^N(B) = (1-B)^2$ and $\delta^S(B) = U(B)$.

Once $\delta(z)$ has been partitioned among the signal and noise appropriately, one typically assumes a balanced ARIMA process for each component, so that the "right-hand operator" in each ARIMA equation has order equal to the left-hand, i.e., we have an MA polynomial of order $d_S + p$ or $d_N + p$ for the signal or the noise, respectively, where p is the order of any autoregressive polynomials in the model. Note that because the factors are made distinct by construction, the order of $\delta(z)$ is $d = d_S + d_N$. The MA polynomials for the signal and noise component processes are then determined via partial fractions, as discussed in Hillmer and Tiao (1982).

Hence, we will proceed from the standpoint that ARIMA models have been found for each of the components, with the following notation:

$$\delta(B) Y_t = W_t = \Psi(B) \epsilon_t$$
$$\delta^S(B) S_t = U_t = \Psi^S(B) \epsilon_t^S$$
$$\delta^N(B) N_t = V_t = \Psi^N(B) \epsilon_t^N.$$

This corresponds to the classical signal extraction scenario (Hillmer and Tiao, 1982). For ARIMA models, one typically assumes that the leading coefficients of AR and MA polynomials are unity.

Next, for any polynomial g of order h we define $\Delta(g)$ to be the $(n-h) \times n$ matrix with entries given by $\Delta_{ij} = g_{i-j+h}$ (with the convention that $g_k = 0$ if $k < 0$ or $k > h$). This means that each row of this matrix consists of the coefficients of the polynomial g, horizontally shifted in an appropriate fashion. We are principally interested in $\Delta(\delta)$, $\Delta(\delta^S)$ and $\Delta(\delta^N)$, which we write as $\Delta(Y)$, $\Delta(S)$ and $\Delta(N)$ for short (the dimension of the matrices will be apparent in formulas, being determined by their compatibility with other terms). Thus

$$W = \Delta(Y) Y \qquad U = \Delta(S) S \qquad V = \Delta(N) N$$

where W, U, V, S, and N are defined analogously to Y. To express

$$W_t = \delta^N(B) U_t + \delta^S(B) V_t \tag{1}$$

in matrix form we need the following result (Lemma 1 of McElroy and Sutcliffe (2006)):

$$\Delta(Y) = \Delta(N)\Delta(S) = \Delta(S)\Delta(N). \tag{2}$$

This is an abuse of notation, because the dimensions of these (non-square) matrices differ, but this compact notation will be preferred. Then we can write down the matrix version of (1):

$$W = \Delta(N) U + \Delta(S) V. \tag{3}$$

We also adopt the following notation: Γ_X denotes the covariance matrix of a random vector X, and for any square integrable (possibly complex) function f, $\Gamma(f)$ is the corresponding covariance matrix with jkth entry

$$\Gamma_{jk}(f) = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(\lambda) e^{i\lambda(j-k)} \, d\lambda.$$

Hence if X is stationary with associated spectral density f, then $\Gamma_X = \Gamma(f)$. There are a few other concepts that we need as well. The components in (3) are essentially "over-differenced," and will be referred to as ∂U and ∂V respectively, as a short-hand. Then they have covariance matrices

$$\Gamma_{\partial U} = \Delta(N)\Gamma_U \Delta'(N)$$
$$\Gamma_{\partial V} = \Delta(S)\Gamma_V \Delta'(S).$$

Finally, we will need to consider the matrix square root of a given symmetric positive-definite matrix (Golub and Van Loan, 1996, p.149). (This is not the same thing as the Cholesky decomposition.) Given such a symmetric positive-definite matrix A, its singular value decomposition takes the form $A = QDQ'$ with Q orthogonal and D diagonal with positive entries. Then $A^{1/2} = QD^{1/2}Q'$ by definition, and satisfies $A^{1/2}A^{1/2} = A$. Moreover this square root is symmetric and has inverse $QD^{-1/2}Q'$, which will be denoted $A^{-1/2}$.

3.2 The Filter Formulas

A signal extraction filter F should reduce the noise process to stationarity, and this condition can be expressed as $F = G\Delta(N)$ for some matrix G. In order to avoid nonstationarity in the error process $FY - S$, we likewise require $1 - F = H\Delta(S)$ for some matrix H, where 1 is the $n \times n$ identity matrix. These two criteria together will be called the *signal extraction conditions*. We also say that a signal estimate FY is *dynamic-matching* if the estimate has the same nonstationary differencing operator as S, and $\Delta(S)FY$ has approximately the same covariance structure of $\Delta(S)S$, namely Γ_U. A further desirable quality is that F be centro-symmetric, i.e., $F_{ij} = F_{n-i+1,n-j+1}$ (see Dagum and Luati (2004), McElroy (2008)). In particular, the centro-symmetry property implies that the asymmetric filters corresponding to the first and last rows of F are transposes of one another, and the middle filter given by the central row of F (when n is odd) is a symmetric sequence.

There is no unique filter matrix with these properties. Part of the problem is that there is no matrix satisfying a dynamic-replicating condition (i.e., that the covariance matrix of $\Delta(S)FY$ is exactly equal to Γ_U) unless the noise process is stationary. (Though in this case, the filter matrix below gives a natural solution that *is* dynamic-replicating.) But more generally when the noise is nonstationary (as is the case with seasonal adjustment), the filter matrix will only be dynamic-matching.

Theorem 1 *Define the following matrices*

$$M = \Delta'(S)\Gamma_U^{-1}\Delta(S) + \Delta'(N)\Gamma_V^{-1}\Delta(N)$$
$$J = 1 - \Gamma_W \Gamma_{\partial U}^{-1/2}\Gamma_W^{-1/2}$$
$$F = M^{-1}\left(\Delta'(N)\Gamma_V^{-1}\Delta(N) - \Delta'(Y)\Gamma_{\partial V}^{-1} J \Delta(Y)\right).$$

Then the matrix F satisfies the signal extraction conditions, is dynamic-matching, and is centro-symmetric. If the noise is stationary F is dynamic-replicating. The error covariance matrix is

$$M^{-1} + M^{-1}\left(\Delta'(Y)\Gamma_{\partial V}^{-1} J \Gamma_W J' \Gamma_{\partial V}^{-1}\Delta(Y)\right) M^{-1}. \qquad (4)$$

Remark 1 It is shown in McElroy (2008) that the minimal mean square error signal extraction filter has error covariance matrix M^{-1}; hence the second term of (4) represents the additional error that results from the dynamic-matching approach. We also see from the formula for F that the DM filter equals the MSE optimal filter minus a matrix term that fully differences the data; hence the MSE optimal signal estimates are adjusted by a stationary series to the DM estimate.

Remark 2 Uniqueness of a dynamic-matching filter cannot be achieved even in the case of stationary signal and noise. In this case, we seek F such that $F\Gamma_Y F' = \Gamma_S$; one class of solutions is given by $F = \Gamma_S^{1/2} R \Gamma_Y^{-1/2}$ for any orthogonal matrix R. Also F will be centro-symmetric so long as R is, by the properties of matrix square roots detailed in the Appendix. Likewise, we can find an entire class of DM filters by inserting R between $\Gamma_{\partial U}^{-1/2}$ and $\Gamma_W^{-1/2}$ in the formula for J in Theorem 1, where R is orthogonal and centro-symmetric. Since the identity matrix is a fairly natural orthogonal centro-symmetric matrix, we take this as the "canonical" choice.

For implementation of these results, one first obtains the component models using either a decomposition or structural approach. The requisite Δ and Γ matrices are then easily formed from the differencing, AR and MA polynomials. Then it is a simple matter to form the filter matrix F from Theorem 1, utilizing singular value decompositions to compute the requisite matrix square roots. Repeating this procedure for every desired signal, we obtain dynamic-matching estimates for all the components of interest. For example, if the data has seasonal (S), trend (T), and irregular (I) components, then

$$Y = \tilde{S} + \tilde{T} + \tilde{I} + R,$$

where $\tilde{S}$, etc. denotes the DM estimate of S. Here R is a remainder component – unlike with WK smoothing, the filter matrices do not sum up to the identity matrix. However, R will be stationary (in a broad sense), since each of the errors $\tilde{S} - S$, $\tilde{T} - T$, and $\tilde{I} - I$ will be stationary. For the application of seasonal adjustment, we see three possible ways of defining a seasonally adjusted component:

$$Y - \tilde{S}, \qquad \tilde{T} + \tilde{I}, \qquad \widetilde{T + I}.$$

Only the last estimate will have the desired dynamic-matching properties, in general. In this case, the residual can be lumped in with the seasonal $\tilde{S}$ as an undesirable portion of the series.

4 Correcting Negative Seasonality

The purpose of this section is re-examine the suite of Foreign Trade series from section 2, this time with SA components obtained from the DM method described in section 4. As before, we fit Box-Jenkins airline models to each regression-adjusted series, but this time we apply the DM filters. The resulting seasonally adjusted series are then examined through ACF plots, with focus on lags

12, 24 and 36 (Tables 1 and 2). For the Import series, 12 were seasonal using WK filters, and for 11 of these the seasonality was removed by using DM filters. In contrast, of the seven nonseasonal WK SAs, none were made more seasonal by the DM procedure. This latter situation is a sort of "Type A" error – namely that a series with no seasonality in its WK SA is made worse by the DM filtering – and is more serious than the "Type B" error, defined as the failure of DM filters to render SAs nonseasonal when they were seasonal under the WK procedure. It is comforting that no Type A errors occur; the Type B error rate is 1/12. Laying statistical significance aside, one finds that for almost every series and lag, the absolute value of the ACF is reduced moving from WK to DM.

For the Export series, 11 WK SAs had seasonality and 5 of these were corrected, so the Type B error rate is 6/11 (and no Type A errors were made). Putting Imports and Exports together with the Shoe Series, we see in Table 3 that the overall Type B error rate is 1/3. Under WK, 60% of the SAs are seasonal, whereas only 20% are using the DM filters. Based on this limited exercise and the theoretical consideration of section 2, it seems likely that improvements to model specifications will reduce the incidence of Type B errors, while we cannot expect any improvements to the WK SAs. This is because negative seasonality is fundamental to WK filtering due to the structure of these filters, whereas the DM filters can produce component estimates that replicate component dynamics so long as one has decent estimates of the covariance structures in the data. In any event, the improvements due to using DM filters are considerable, with no cost in terms of Type A error.

Note that Type A and B errors are not like Type I and II errors in standard statistical inference, because the former can always be ascertained in practice. Hence if application of DM filtering happens to not remove seasonality (which includes the Type A error case that the WK SA was completely nonseasonal), one can fiddle with the model specification in an effort to achieve complete nonseasonality.

The four series with noninvertible airline models all had a unit root for the seasonal moving average parameter, indicating that the ARIMA modeling equation could be simplified through cancelation, compensating through inclusion of seasonal dummies. In these cases, the SA estimates generated by the DM filters had bizarre patterns. This was handled by fixing the seasonal moving average parameter at the value .98, which is suitably close to unity. In practice we recommend changing the model in such cases, but for the sake of uniformity of comparisons we stuck with the airline model for all 40 series.

Given that DM filtering is easy to implement (it requires about as much code as an implementation of matrix-based WK filtering as described in McElroy (2008), and is just as fast to run in practice) and does a superior job of reducing seasonality, what are the drawbacks? From (4) we can expect an increase in mean squared error (after all, the WK filters minimized MSE!); based on Figure 3, the DM SAs seem to be a bit noisier (as are the trends, though this is not displayed here). Using the formula (4), we can plot both MSEs for the Shoe series: see Figure 5. As expected, the

DM method has uniformly higher MSE. The increase in MSE, given by the second term in (4) is close to constant; the relative discrepancy in the middle of the series is 19.7% – so there is roughly a 20% increase in MSE due to use of DM over WK (the same calculation with standard errors reduces to a 9.4% increase). This type of assessment only requires a knowledge of the fitted model. If the increase in MSE is deemed to be too large for a particular application, then the practitioner can stay with the WK approach. We also examined revision error variances for concurrent filters (as described in McElroy and Gagnon (2008)), but in the cases examined there was less than a 3% discrepancy between DM and WK.

Another comparison between WK and DM is afforded by comparing squared gain and phase delay plots for concurrent filters for both series. As discussed in Findley and Martin (2006), gain plots can be used to assess how a filter attenuates the variance of a time series at various frequencies, whereas the phase delay gives information about how much lag a filter induces on a time series at each frequency. These are displayed, for the concurrent filter used on the Shoe series, in Figure 6. The reduced width of the gain function's spectral trough for the DM method corresponds to its dynamic-matching properties, since in an approximate sense the squared gain is designed such that its product with the data spectrum f_y will yield the target spectrum f_{SA} (see the discussion in Section 2). The phase delays are similar, but the DM filter has slightly less delay in the frequency band $(0, .2)$, which correspond to trend-cycle frequencies. Thus the increased error of the DM method is to some extent counter-acted by the reduced lag in the low frequency component, which is of course highly desirable.

5 Conclusion

In this paper we consider the problem of residual seasonality in seasonal adjustment estimates generated from model-based WK methods. The observation that MSE optimal filters tend to induce seasonal spectral troughs in SA estimates is not new – Nerlove (1964) seems to be the first to observe the phenomenon, whereas Bell and Hillmer (1984) gave an explanation. Until their paper, there was concern in the seasonal adjustment community that WK filtering was somehow defective; Bell and Hillmer (1984) showed that the dips were to be expected. The comment on their paper by Ansley and Wecker (1984) further shows how this issue can be corrected by taking the square root of the WK filter's frequency response function. (Also see Wecker (1979) for earlier work on this concept.) However, note that their results are derived under the assumptions of stationary signal and noise, and they focus on bi-infinite symmetric filters.

Given this background, we have sought to: (1) provide additional documentation of residual negative seasonality in "typical" economic time series[2]; (2) make a case for dynamic-matching ap-

[2]Nerlove (1964) also documents this via plotting spectral estimates obtained via the "recoloring" approach of using the pseudo-periodogram; however, the statistical significance of the troughs – and hence the phenomenon – is

proaches to seasonal adjustment; (3) present a simple matrix approach to obtain DM SA estimates; (4) document the pros and cons of the DM method on the same economic series. Theorem 1 shows that the DM method has reasonable properties, and takes little extra time over the direct matrix approach for WK filtering (although the method will be slower than using a state space smoother). For the sake of making comparisons, the airline model was used for all 40 time series; but even in this case, which is somewhat unfavorable to the DM approach, there are appreciable improvements to be seen in the reduction of negative seasonality. The cost in terms of increased MSE and benefit in terms of improved phase delay in the low frequency band can be assessed directly, as discussed in Section 4.

We also mention some directions for future research. The DM method of Theorem 1 need not be limited to seasonal adjustment, but can also be used for trend and cycle estimation, for example. It would also be interesting to extend the dynamic-matching philosophy to the forecasting of nonstationary time series.

Appendix

Proof of Theorem 1. We must show that the stated F satisfies: $F = G\Delta(N)$, $1 - F = H\Delta(S)$, and is centro-symmetric. Utilizing (2), we see that

$$G = M^{-1}\left(\Delta'(N)\Gamma_V^{-1} - \Delta'(Y)\Gamma_{\partial V}^{-1} J\Delta(S)\right)$$
$$H = M^{-1}\left(\Delta'(S)\Gamma_U^{-1} + \Delta'(Y)\Gamma_{\partial V}^{-1} J\Delta(N)\right).$$

To prove centro-symmetry, define the transverse-transpose of a square matrix A to be A^* with jkth entry $A_{n-k+1,n-j+1}$. Then A is centro-symmetric iff $A = A^*$. Based on the discussion in section 4.1 of McElroy (2008), it suffices to demonstrate centro-symmetry of $\Delta'(Y)\Gamma_{\partial V}^{-1}J\Delta(Y)$. Now from the definition of the matrix square root and the elementary properties $(AB)^* = B^*A^*$ and $A^{*\prime} = A'^*$, we find that $(A^{1/2})^* = (A^*)^{1/2}$. It follows that if A is centro-symmetric, so is $A^{1/2}$. So let $\tilde{\Delta}$ be a $n \times n$ matrix with jkth entry δ_{j-k}, so that $[0\,1]\tilde{\Delta} = \Delta(Y)$. Then

$$\begin{aligned}
&\Delta'(Y)\Gamma_{\partial V}^{-1}J\Delta(Y) \\
&= \tilde{\Delta}' \begin{bmatrix} 0 \\ 1 \end{bmatrix} \Gamma_{\partial V}^{-1}\left(1 - \Gamma_W \Gamma_{\partial U}^{-1/2}\Gamma_W^{-1/2}\right)[0\,1]\,\tilde{\Delta} \\
&= \tilde{\Delta}' \begin{bmatrix} 0 & 0 \\ 0 & \Gamma_{\partial V}^{-1}\left(1 - \Gamma_W \Gamma_{\partial U}^{-1/2}\Gamma_W^{-1/2}\right) \end{bmatrix} \tilde{\Delta}.
\end{aligned}$$

Applying the operator $*$ yields

$$\tilde{\Delta}^* \begin{bmatrix} \left(1 - \Gamma_W^{-1/2}\Gamma_{\partial U}^{-1/2}\Gamma_W\right)\Gamma_{\partial V}^{-1} & 0 \\ 0 & 0 \end{bmatrix} \tilde{\Delta}'^* = \Delta'(Y)\left(1 - \Gamma_W^{-1/2}\Gamma_{\partial U}^{-1/2}\Gamma_W\right)\Gamma_{\partial V}^{-1}\Delta(Y).$$

---

not ascertained.

But this is the transpose of the original expression, which establishes the centro-symmetry of $\Delta'(Y)\Gamma_{\partial V}^{-1} J \Delta(Y)$. To establish dynamic-matching, observe that

$$\begin{aligned}
\Delta(S)FY &= \Delta(S)G\Delta(N)Y \\
&= \Gamma_U \Delta'(N)\Gamma_W^{-1}\Delta(S)\Gamma_V \left(\Gamma_V^{-1} - \Delta'(S)\Gamma_{\partial V}^{-1} J \Delta(S)\right) \Delta(N)Y \\
&= \Gamma_U \Delta'(N)\Gamma_W^{-1} \left(1 - \Delta(S)\Gamma_V \Delta'(S)\Gamma_{\partial V}^{-1} J\right) \Delta(Y)Y \\
&= \Gamma_U \Delta'(N)\Gamma_W^{-1}(1 - J)W \\
&= \Gamma_U \Delta'(N)\Gamma_{\partial U}^{-1/2}\Gamma_W^{-1/2} W.
\end{aligned}$$

In the second equality we use the fact (see McElroy (2008)) that $\Delta(S)M^{-1}\Delta'(N) = \Gamma_U \Delta'(N)\Gamma_W^{-1}\Delta(S)\Gamma_V$. The above random vector has covariance matrix

$$\Gamma_U \Delta'(N)\Gamma_{\partial U}^{-1}\Delta(N)\Gamma_U = \Gamma_U^{1/2}\left(\Gamma_U^{1/2}\Delta'(N)\Gamma_{\partial U}^{-1}\Delta(N)\Gamma_U^{1/2}\right)\Gamma_U^{1/2}.$$

The matrix in the center of the right-hand expression is idempotent, with rank equal to $n - d_N$. Therefore for large n it is an approximate identity matrix, and the covariance is approximately Γ_U. When the noise is stationary, the covariance reduces to exactly Γ_U. Finally, we compute the signal extraction error covariance matrix. The error process is

$$\epsilon = FY - S = (F - 1)S + FN = GV - HU,$$

which has covariance matrix

$$\Gamma_\epsilon = G\Gamma_V G' + H\Gamma_U H'.$$

We can use the above formulas for G and H to compute this, and we easily obtain (4). □

References

[1] Ansley, C. and Wecker, W. (1984) On Dips in the Spectrum of a Seasonally Adjusted Time Series. Comment to "Issues involved with the seasonal adjustment of economic time series" by Bell and Hillmer. *Journal of Business and Economics Statistics* **2**, 323–324.

[2] Bell, W. and Hillmer, S. (1984) Issues involved with the Seasonal Adjustment of Economic Time Series. *Journal of Business and Economic Statistics* **2**, 291–320.

[3] Box, G. and Jenkins, G. (1976) *Time Series Analysis,* San Francisco, California: Holden-Day.

[4] Brockwell, P. and Davis, R. (1991) *Time Series: Theory and Methods.* New York: Springer.

[5] Dagum, E. and Luati, A. (2004) A linear transformation and its properties with special applications in time series filtering. *Linear Algebra and its Applications* **338**, 107–117.

[6] Findley, D. and Martin, D. (2006) Frequency Domain Analyses of SEATS and X-12-ARIMA Seasonal Adjustment Filters for Short and Moderate-Length Time Series. *Journal of Official Statistics* **22**, 1–34.

[7] Golub, G. and Van Loan, C. (1996) *Matrix Computations.* Baltimore and London: The Johns Hopkins University Press.

[8] Grether, D. and Nerlove, M. (1970) Some Properties of Optimal Seasonal Adjustment. *Econometrica* **38**, 682–703.

[9] Hillmer, S. and Tiao, G. (1982) An ARIMA-model-based approach to seasonal adjustment. *Journal of the American Statistical Association* **77**, 63–70.

[10] McElroy, T. (2008) Matrix Formulas for Nonstationary ARIMA Signal Extraction. *Econometric Theory* **24**, 1-22.

[11] McElroy, T. and Gagnon, R. (2008) Finite Sample Revision Variances for ARIMA Model-Based Signal Extraction. *Journal of Official Statistics* **24**, 451–467.

[12] McElroy, T. and Sutcliffe, A. (2006) An Iterated Parametric Approach to Nonstationary Signal Extraction. *Computational Statistics and Data Analysis, Special Issue on Signal Extraction* **50**, 2206–2231.

[13] Nerlove, M. (1964) Spectral Analysis of Seasonal Adjustment Procedures. *Econometrica* **32**, 241–286.

[14] Wecker, W. (1979) A New Approach to Seasonal Adjustment. *Proceedings of the American Statistical Association, Business and Economic Statistics Section.* 322–323.

Table 1. Negative Seasonality for Imports

	ACF for WK SA			ACF for DM SA		
Series	12	24	36	12	24	36
m00120	-.155	-.236*	.053	-.083	-.194	.054
m00190	-.300*	-.312*	.101	-.139	-.283*	.065
m12060	-.258*	-.157	-.073	-.034	-.130	-.128
m12135	-.131	-.153	.127	-.059	-.100	.132
m12150	-.225*	-.097	-.083	-.163	-.075	-.070
m12540	-.022	-.077	.044	-.020	-.074	.046
m21110	-.097	-.110	-.116	-.057	-.097	-.111
m21160	-.045	-.160	-.193	.016	-.110	-.177
m21180	-.181	-.097	.013	-.135	-.075	.017
m21610	-.237*	-.168	-.003	-.173	-.137	-.004
m22020	-.269*	-.170	.100	-.211	-.138	.099
m3000	-.252*	-.146	-.038	-.084	-.091	-.024
m3010	-.158	.097	-.216*	-.004	.117	-.210
m40020	-.273*	-.250*	.050	-.060	-.156	.043
m40040	-.218*	-.219*	-.021	-.142	-.207	-.041
m40110	-.191	-.114	-.032	-.191	-.113	-.032
m41140	-.052	-.070	-.053	.023	-.036	-.058
m41310	-.339*	-.038	-.124	-.176	-.011	-.102
m42110	-.362*	-.014	-.069	-.148	.042	-.032

Table 1: Sample ACF values at indicated lags 12, 24, and 36 for 19 Import series. WK SA refers to seasonal adjustments estimated through WK filters, while DM SA refers to the use of the Dynamic-Matching method. Asterisks flag values that are significant.

Table 2. Negative Seasonality for Exports

	ACF for WK SA			ACF for DM SA		
Series	12	24	36	12	24	36
x00300	.079	.004	-.149	.079	.005	-.150
x10140	-.308*	-.113	-.068	-.167	-.069	-.041
x11020	-.123	-.161	-.093	-.043	-.124	-.081
x12550	-.097	-.218*	-.017	.052	-.153	-.047
x12600	-.297*	-.131	.092	-.223*	-.097	.097
x12770	-.123	-.093	-.266*	.002	.041	-.217
x13200	-.077	-.142	-.028	-.040	-.127	-.032
x21000	-.337*	-.081	-.000	-.288*	-.068	-.000
x21030	-.257*	.031	-.101	-.257*	.033	-.100
x21150	-.188	-.125	-.050	-.151	-.110	-.047
x21500	-.126	-.237*	.035	-.064	-.197	.039
x3020	-.149	-.158	-.167	-.049	-.142	-.185
x3022	-.103	-.208	-.020	.005	-.143	-.026
x3	.008	-.234*	-.170	.056	-.194	-.155
x40000	-.218	-.253*	-.117	-.122	-.233*	-.122
x40030	-.358*	-.031	.008	-.265*	-.006	.035
x41020	-.160	-.188	-.100	-.069	-.157	-.099
x41120	-.154	-.140	-.142	-.071	-.121	-.150
x41140	-.226*	.034	-.055	-.226*	.034	-.055
x42100	-.121	-.126	.076	.019	-.040	.066

Table 2: Sample ACF values at indicated lags 12, 24, and 36 for 20 Export series. WK SA refers to seasonal adjustments estimated through WK filters, while DM SA refers to the use of the Dynamic-Matching method. Asterisks flag values that are significant.

Table 3. Error Chart

	DM Seasonal	DM Nonseasonal	Total
WK Seasonal	8	16	24
WK Nonseasonal	0	16	16
Total	8	32	40

Table 3: Tallies of SAs deemed to be significantly seasonal/nonseasonal before (WK) and after (DM) correction.

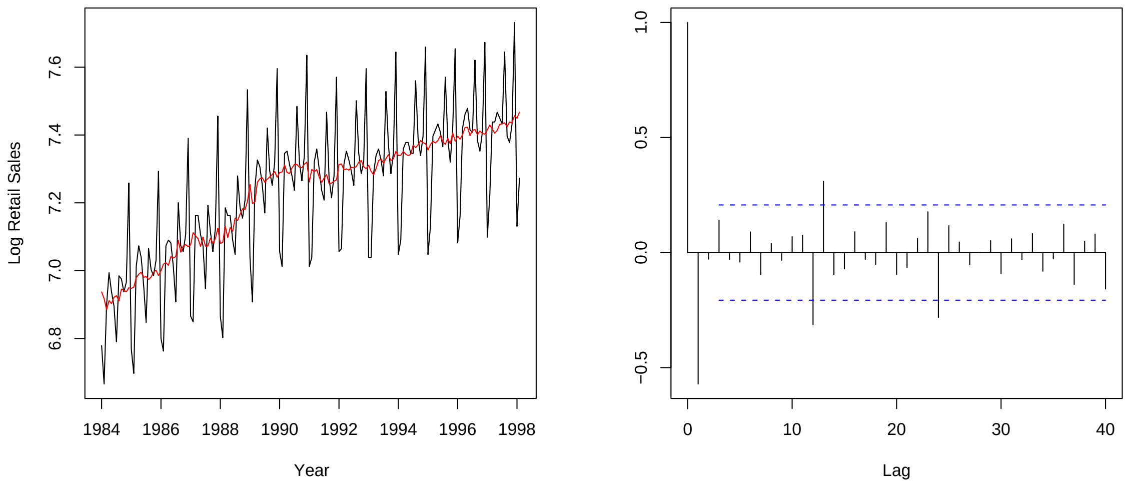

Figure 1: The left panel displays the trading day and outlier adjusted Shoe series in logarithms in black, with the WK SA component in red. The right panel displays the sample ACF plot for the twice-differenced WK SA component. Notice the significant negative correlations at lags 12 and 24.

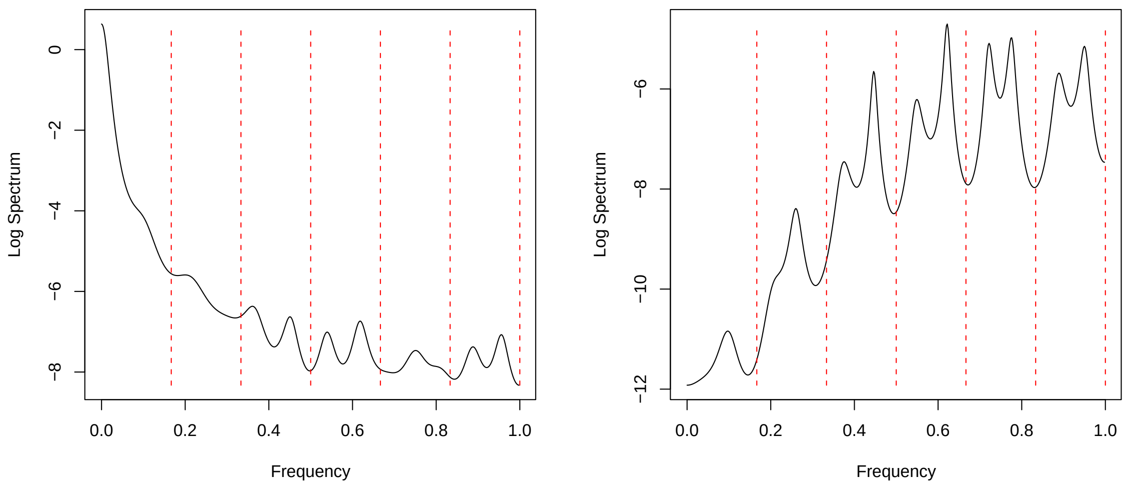

Figure 2: The left panel displays the AR(30) spectrum estimate (in logs) of the WK SA component. The right panel displays the AR(30) spectrum estimate (in logs) of the twice-differenced WK SA component. Vertical red lines in both panels indicate the seasonal frequencies.

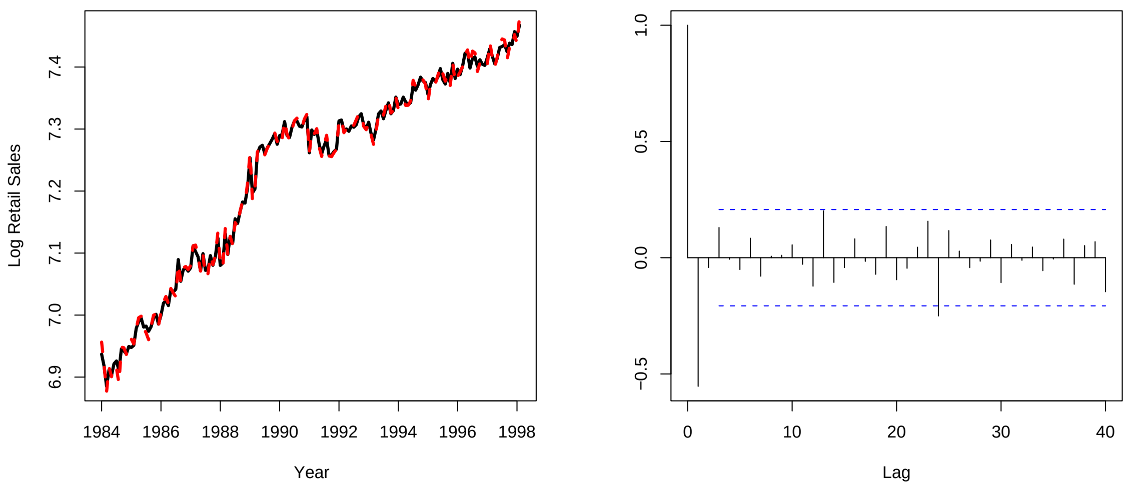

Figure 3: The left panel displays the two SAs together, with the WK SA in black (solid) and the dynamic-matching SA in red (dotted). The right panel displays the sample ACF plot for the twice-differenced dynamic-matching SA component. Notice there is no longer a significant negative correlation at lag 12, though some negative correlation remains at lag 24.

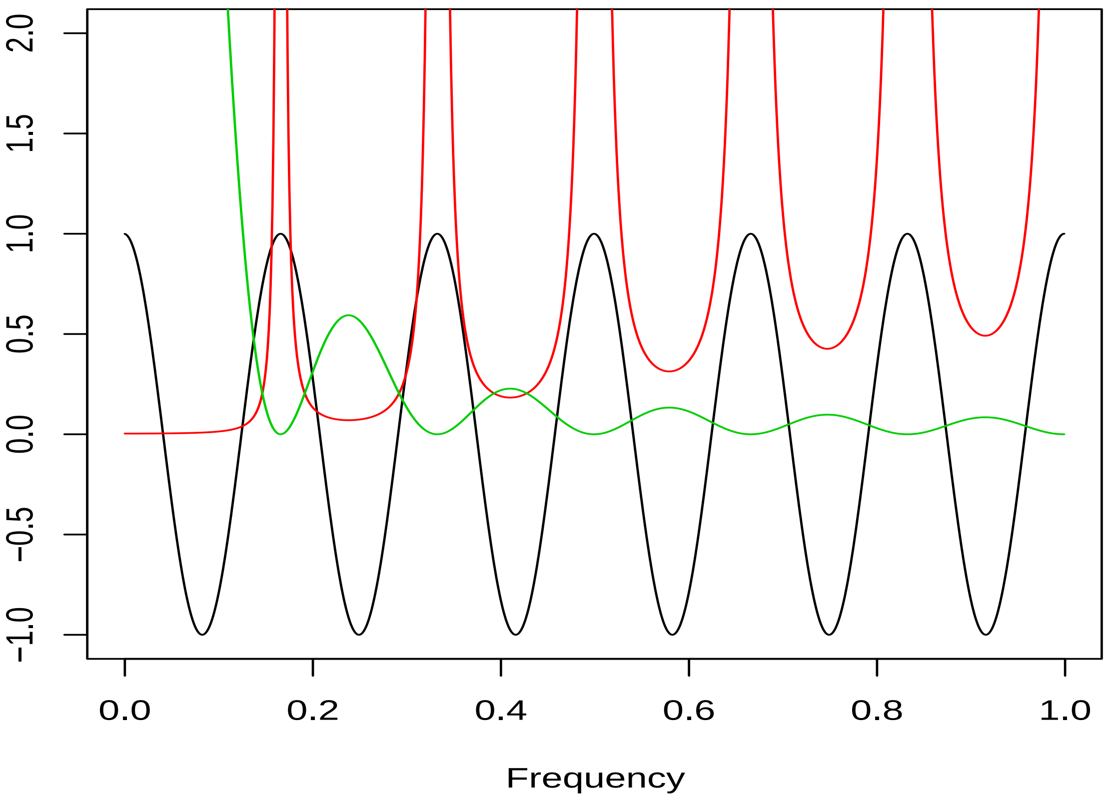

Figure 4: This figure plots the weight function $\cos 12\lambda$ in black, for frequencies $\lambda \in [0, 1]$. In red is the pseudo-spectral density f_1, which has peaks of varying thickness and height at the seasonal frequencies. In green is the spectral density f_2, which has troughs of varying depths at the seasonal frequencies.

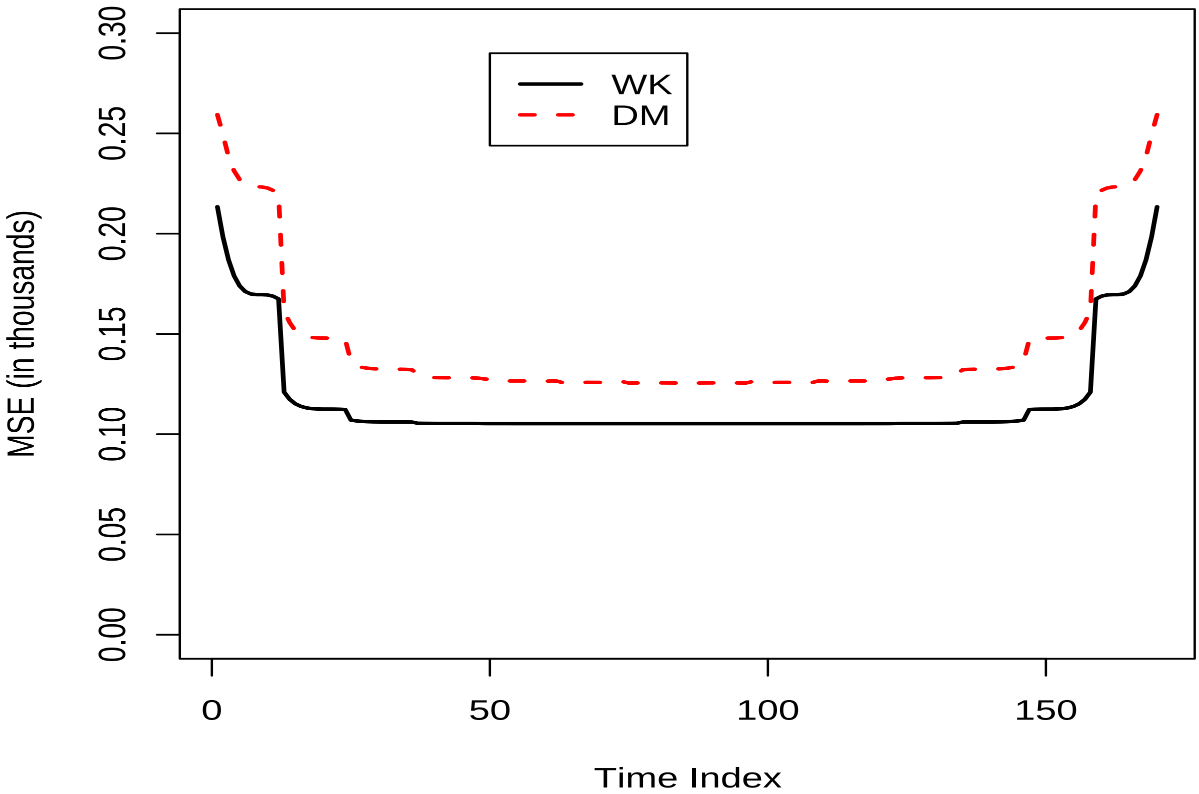

Figure 5: This figure plots the MSEs (as a function of time) for both the WK and DM methods, for the model fitted to the Shoe series.

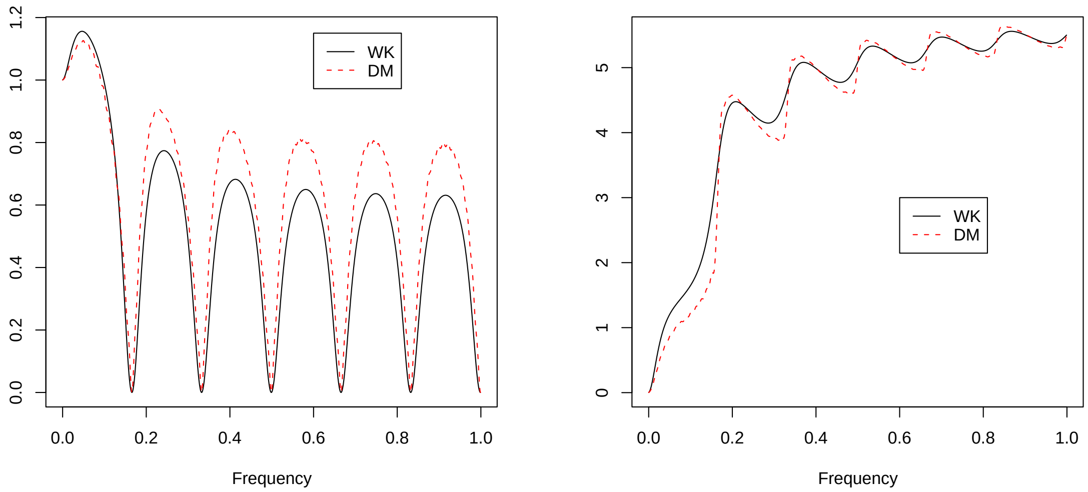

Figure 6: The left panel displays the two squared gains for the concurrent filter used on the Shoe series, with the WK in black (solid) and the DM in red (dotted). The right panel displays the two phase delays for the concurrent filter used on the Shoe series, with the WK in black (solid) and the DM in red (dotted).