Seasonal Adjustment Subject to Frequency Aggregation Constraints

Tucker McElroy* Osbert Pang* Brian Monsell†

Abstract
Quarterly seasonal adjustments in official statistics are often not the result of a direct adjustment of the quarterly series, but instead are an indirect adjustment arising from the aggregation of the seasonally adjusted monthly series. However, the temporal aggregation of nonseasonal monthly series to a quarterly frequency can exhibit seasonality; we provide a rigorous framework for understanding how this occurs. To solve the problem, we build on prior work that uses benchmarking to enforce seasonal adjustment adequacy as temporal aggregation is applied, where adequacy is metrized and supplied as a hard constraint to the benchmarking optimization problem. It is vital to use a seasonality diagnostic that examines a time series at high seasonal lags, and can properly capture type I and type II errors, and therefore we propose to utilize new autoregressive seasonality diagnostics in tandem with the proposed benchmarking procedure. We examine the proposed procedure on X-13ARIMA-SEATS seasonal adjustments of several economic time series.

Key Words: Benchmarking, Temporal Aggregation, Adjustment Adequacy

Disclaimer This report is released to inform interested parties of research and to encourage discussion. The views expressed on statistical issues are those of the authors and not those of the U.S. Census Bureau.

1. Introduction

The problem of the presence of residual seasonality in published GDP has recently been publicized through a series of articles, with concerns first being articulated by Furman (2015), Gilbert et al. (2015), Rudebusch et al. (2015), and Groen and Russo (2015). These critiques have prompted renewed interest in seasonality diagnostics and seasonal adjustment at the Bureau of Economic Analysis (BEA) – see the discussion in Lengerman et al. (2017). McCulla and Smith (2015) review some of the changes BEA has implemented in response to the critiques, whereas Phillips and Wang (2016) points out continuing difficulties. Preliminary findings at BEA determined that residual seasonality could arise from the aggregation of monthly source data to a quarterly frequency (Moulton and Cowan, 2016); this phenomenon was demonstrated through simulations and theoretical models (McElroy, 2016).

When monthly source data is available in a raw form, it can be seasonally adjusted and then aggregrated, resulting in the indirect adjustment of the quarterly data. Alternatively, it can be first aggregated and then adjusted, resulting in the direct adjustment of the quarterly data. Whereas operationally there is more control over outcomes in the latter approach (as the analyst has control over program settings to ensure adequacy), the direct adjustment will typically be different from the indirect adjustment, and in particular is not equal to aggregation of the monthly adjustment.

BEA requires this equality be preserved, because for certain components of GDP, monthly seasonally adjusted numbers are published. In particular, the goods and services portion of Personal Consumption of Expenditure (PCE), which accounts for two thirds of GDP, is published at a monthly frequency and needs to be consistent with quarterly values.

*U.S. Census Bureau, 4600 Silver Hill Road, Washington, DC 20233-9100
†Retired (U.S. Census Bureau)

Further complicating matters, for some components of GDP there is no monthly raw data available, i.e., only the monthly seasonal adjustments have been provided to BEA (by other agencies). Hence there is no way to compute a quarterly aggregate, and therefore no way to compute a direct adjustment; apparently, the indirect adjustment is the only available method, and in some cases the adjustment is inadequate.

This paper addresses these problems by extending the reconciliation methodology advocated in McElroy (2018), which focused upon the related problem of cross-aggregation of time series with the same sampling frequency. The novelty of our method lies in using nonlinear optimization to obtain reconciled series (which satisfy the aggregation constraints across frequency) such that seasonal adjustment adequacy holds. While there exists a substantial literature on the enforcing of temporal aggregation constraints (i.e., benchmarking), the key issue of adequacy – namely, that the final quarterly adjustment does not exhibit seasonality – has not been addressed (cf. Denton (1971), Dagum (1979), Cholette (1984), Dagum and Cholette (2006), Di Fonzo and Marini (2011), Quenneville and Fortier (2012), and Chen (2012)). Although it has been recognized in the seasonal adjustment community that this is a pervasive challenge (cf. Hood and Findley (2001) and Astolfi, Ladiray, and Mazzi (2001)), there is no available method that simultaneously addresses frequency aggregation and adequacy.

Our solution involves nonlinear optimization for each series, whereby monthly adjustments are changed as little as possible, such that they are still adequate and also aggregate to the quarterly adjustment, which is also enforced to be adequate. In the case that no raw monthly data is available, we utilize the seasonal adjustment of the indirect quarterly adjustment (assuming it is inadequate – otherwise we are already done) as the program's quarterly benchmark. We motivate our work by a discussion of the quarterly aggregation phenomenon in Section 2, followed by a rigorous treatment of how dynamics are altered by aggregation in Section 3. Section 4 proposes the benchmarking methodology, and Section 5 illustrates with a couple of applications to economic data. Section 6 offers some concluding remarks.

2. The Phenomenon of Seasonality in Frequency Aggregated Series

The phenomenon of "frequency-aggregated seasonality" refers to the situation where a change in sampling frequency via flow aggregation results in the manifestation of seasonality. For example, we may have a monthly time series that exhibits no seasonality, and yet when aggregated to a quarterly frequency displays seasonality. A second case arises when the monthly series is indeed seasonal, is seasonally adjusted, and the quarterly aggregate of the monthly adjustment displays seasonality. In either case, direct adjustment of the quarterly series results in removal of seasonality, but typically this direct adjustment will no longer be an aggregate of the monthly adjustment. If these types of aggregation relationships are desirable, then a modification of the direct adjustment is needed.

Such phenomena have been empirically observed in PCE goods and services data, as noted and discussed in McElroy (2018). A possible explanation from economics is suggested by budgets and quotas: a certain amount of resources is allocated on a quarterly basis, but is expended in the three months of that quarter as needs arise. Hence there is no seasonal pattern on a monthly basis, but the rationing becomes apparent at a quarterly frequency; naturally, this indicates a negative correlation between months that lie within the same quarter. We proceed to describe a process with these properties. Suppose that a monthly series $\{X_t\}$ for $t \in \mathbb{Z}$ can be written in terms of latent processes $\{S_t\}$ and $\{N_t\}$,

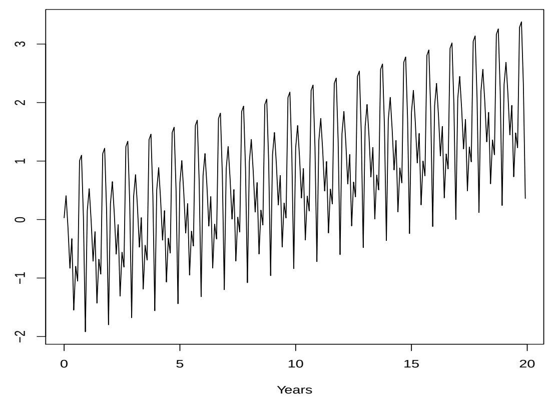

Figure 1: Simulated monthly series with salient seasonality (and trend).

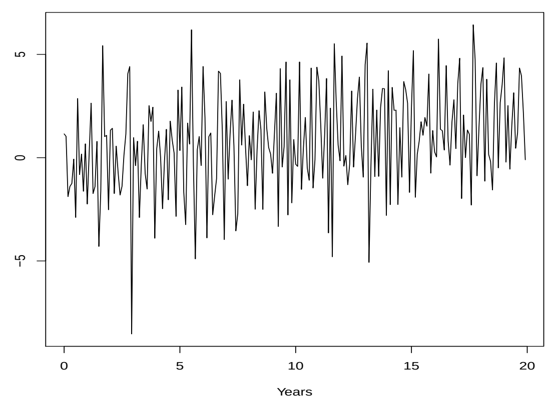

Figure 2: Simulated monthly series, consisting of monthly seasonality plus quota noise.

representing unobserved seasonal and nonseasonal dynamics respectively:

$$X_t = S_t + N_t. \tag{1}$$

We suppose that $\{S_t\}$ has apparent seasonality, but that the variability in the so-called "quota" noise $\{N_t\}$ is so high that no seasonality is apparent in $\{X_t\}$ (see McElroy (2018) for a related discussion). Figure 1 shows a simulated $\{S_t\}$, together with a time trend, and Figure 2 shows the result of adding quota noise. There is no apparent seasonality in this monthly series. In particular, there is no oscillatory pattern present in the sample autocorrelations (Figure 3), and the values at lags 12, 24, and 36 are low. As for the spectral density, there are only minor peaks present, and these are not located exactly at the seasonal frequencies denoted by vertical lines in Figure 4. (See Findley, Lytras, and McElroy (2017) for a discussion of seasonality diagnostics.)

The quota noise $\{N_t\}$ is constructed as a nonstationary process where every third random variable equals the negative of the sum of the previous two:

$$N_t = -N_{t-1} - N_{t-2} \qquad \text{if } t \bmod 3 = 0.$$

(Alternative constructions can be envisioned: the covariance matrices of the three random

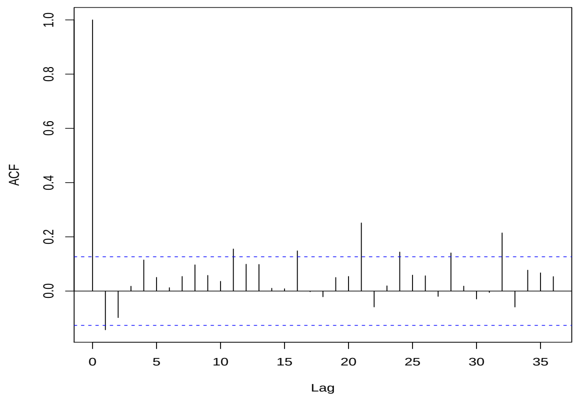

Figure 3: Autocorrelation plot of noisy monthly series. (Big spikes at lags 12, 24, and 36 would indicate seasonality.)

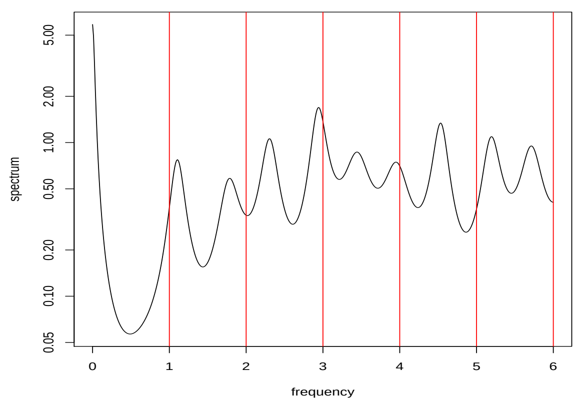

Figure 4: Spectral density plot of noisy monthly series. (Big spikes at seasonal frequencies 1, 2, 3, 4, 5, 6 in red would indicate seasonality.)

variables corresponding to a given quarter can be defined to have $[1, 1, 1]'$ as a null vector, or an approximate null vector.) Flow aggregation means that every three monthly observations are summed to a quarterly value; applying such an operator to $\{N_t\}$, we clearly obtain zero, so that the flow aggregation of $\{X_t\}$ equals the flow aggregation of $\{S_t\}$ – which is likely to still be seasonal, as is discussed in Section 3 below. In our example, the flow aggregation of $\{X_t\}$ (and $\{S_t\}$) is displayed in Figure 5; the quarterly seasonality is highly significant and apparent.

In practice we find that examples tend to be less dramatic: for instance, it may occur that a given monthly series exhibits very weak seasonality, and hence is not adjusted, but its quarterly aggregation displays a much greater degree of seasonality – requiring action. In such a case, indirect adjustment is tantamount to quarterly aggregation (since the original monthly series is not adjusted at all), and hence is inadequate.

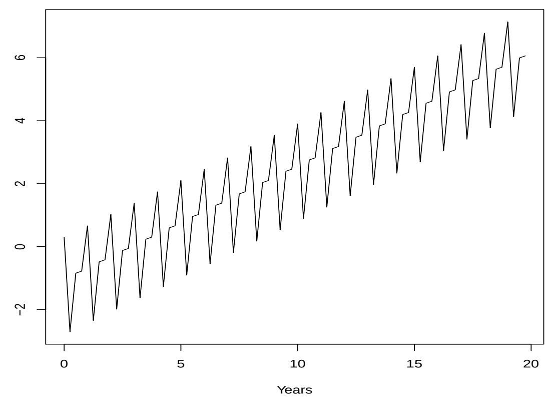

Figure 5: Quarterly aggregation of simulated monthly series with quota noise.

3. How Frequency Aggregation Alters ARIMA Structure

We here discuss how the structure of an ARIMA process is altered by down-sampling and frequency aggregation. Let a given high-frequency process $\{X_t\}$ have an ARIMA structure, such that $\delta(B)X_t = Y_t$ is stationary and $\phi(B)Y_t = \theta(B)Z_t$, where ϕ and θ are degree p and q polynomials with roots outside the unit circle, and δ is a degree d polynomial with roots on the unit circle. (We can relax this slightly, so that θ is permitted to have roots on the unit circle.) Set $\varphi(z) = \delta(z)\,\phi(z)$, the pseudo-autoregressive polynomial. For a positive integer s, we consider embedding a given high-frequency process as a s-variate low-frequency process $\{\mathbf{X}_n\}$, where the jth component of $\mathbf{X}_n$ is X_{ns+j}. This jth component is called the jth *season's series*; considering all such series together in vector format is called the *seasonal vector series*. This type of embedding has been studied by many authors: see Gladyshev (1961), Tiao and Grupe (1980), Osborn (1991), and Franses (1994).

The embedding results of this section expand on those of Tiao and Grupe (1980) by developing the algebra; we first set out some notation. The backshift operator B acts on the high-frequency time index via $B\,X_t = X_{t-1}$; the corresponding low-frequency lag operator is defined to be $L = B^s$, so that $L\,\mathbf{X}_n = \mathbf{X}_{n-1}$. The $s \times s$ identity matrix is denoted 1_s, and we use underlines to denote matrices. Given a scalar Laurent series $\psi(z)$, its embedded Laurent series is defined to be the $s \times s$-dimensional matrix Laurent series $\underline{\psi}(L) = \sum_k \underline{\psi}_k L^k$, where the r, m coefficients of $\underline{\psi}_k$ are given by $\psi_{r-m+s*k}$; cf. equation (2.7) of Tiao and Grupe (1980). Also, if $\psi(B)$ is a polynomial of degree q, then the degree of the matrix polynomial $\underline{\psi}(L)$ is $\lfloor q/s \rfloor + 1$.

Illustration: Consider the polynomials $d(B) = 1 - B$ and $U(B)$ with $s = 4$. The embedding yields

$$\underline{d}(L) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} L$$

$$\underline{U}(L) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix} + \begin{bmatrix} 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} L.$$

Even though the ideal of matrix polynomials is not Abelian, the seasonal embedding mapping is a homomorphism with respect to polynomial products, which is an important property. If $\mathcal{W}$ denotes the embedding, and $a(B)$ and $b(B)$ are two scalar polynomials, then $\mathcal{W}(a \cdot b) = \mathcal{W}(a) \cdot \mathcal{W}(b)$. As a result, if a scalar polynomial $c(B)$ factorizes into a product $c(B) = a(B) \cdot b(B)$, then automatically $\underline{c}(L) = \underline{a}(L) \cdot \underline{b}(L)$. Furthermore, because the scalar product is Abelian, it follows from the homomorphism property that $\underline{c}(L) = \underline{b}(L) \cdot \underline{a}(L)$.

Proposition 1 *Let $\mathcal{W}$ be the mapping of $a(B)$ to $\underline{a}(L)$, with $\underline{a}_k(r,m) = a_{r-m+ks}$. Then $\mathcal{W}$ is a homomorphism.*

Proof of Proposition 1. Let e_r and e_m denote unit vectors. The kth coefficient of $\underline{a}(L) \cdot \underline{b}(L)$ is $\sum_j \underline{a}_j \, \underline{b}_{k-j}$. The (r,m)th entry of such is

$$\begin{aligned} e_r' \sum_j \underline{a}_j \, \underline{b}_{k-j} e_m &= \sum_j \sum_h \underline{a}_j(r,h) \, \underline{b}_{k-j}(h,m) \\ &= \sum_{j,h} a_{r-h+sj} \, b_{h-m+s(k-j)} \\ &= \sum_\ell a_\ell \, b_{r-m-\ell+sk} \end{aligned}$$

by change of variable. This is recognized as c_{r-m+sk}, or $\underline{c}_k(r,m)$. $\square$

The seasonal embedding can be applied to a scalar process as follows. Suppose $\delta(B) X_t = \psi(B) Z_t$, where $\{Z_t\}$ is white noise of variance σ^2, and $\psi(z) = \theta(z)/\phi(z)$. (Below we consider an extension to non-ARIMA processes.) Then the seasonal embedding is

$$\underline{\delta}(L) \mathbf{X}_n = \underline{\psi}(L) \mathbf{Z}_n,$$

where $\{\mathbf{Z}_n\}$ is vector white noise with variance matrix equal to $\sigma^2 1_s$. We can use these results to describe the ARIMA model for the down-sampled and frequency-aggregated processes. First, the matrix polynomial $\underline{\delta}(L)$ can be written in terms of its adjoint and determinant as follows:

$$\underline{\delta}^\sharp(L) \, \underline{\delta}(L) = \det \underline{\delta}(L) \, 1_s. \tag{2}$$

This equation implicitly defines the adjoint $\underline{\delta}^\sharp(L)$, which has dimension equal to $D(s-1)$, where D is the degree of $\underline{\delta}(L)$. We note in passing that the two matrix polynomials commute with one another. Also, the degree of $\det \underline{\delta}(L)$ is Ds. As discussed above, $D = \lfloor d/s \rfloor + 1$, so the degree of the determinantal polynomial is $d + d_s$ where $0 \leq d_s < d$ is the remainder occurring when d is divided by s.

Consider Lemma 1 of Tiao and Grupe (1980), and apply this to the embedding of $\delta(B)$. It follows that if ζ is a root of $\delta(z)$, then ζ^s has the property that $\underline{\delta}(\zeta^s)$ is a singular matrix with row eigenvector $[1, \zeta, \ldots, \zeta^{s-1}]$. Hence we can conclude that ζ^s is a root of $\det \underline{\delta}(z)$, and this accounts for exactly d of the $Ds = d + d_s$ roots. (Note that taking the sth power of the various roots ζ of $\delta(z)$ can result in repeated roots.) The converse is also true: if some ξ is a root of $\det \underline{\delta}(z)$, then necessarily $\underline{\delta}(\xi)$ is singular, indicating it has a left null-vector, i.e., a left eigenvector with associated eigenvalue of zero – the only such eigenvalues are given by the roots of $\delta(z)$. Hence, the actual degree of $\det \underline{\delta}(z)$ is d, because the final d_s coefficients are all zero. As a result, we obtain

$$\det \underline{\delta}(z) = \prod_{j=1}^d (1 - z/\zeta_j^s).$$

Next, using (2) we see that each component of $\mathbf{X}_n$ can be differenced to stationarity by application of $\det \underline{\delta}(L)$, because

$$\det \underline{\delta}(L)\, \mathbf{X}_n = \underline{\delta}^\sharp(L)\, \underline{\psi}(L) \mathbf{Z}_n$$

and the right hand side is a stationary process. Denoting this stationary process by $\mathbf{Y}_n$, we can obtain an adjoint polynomial for $\underline{\phi}(L)$ in the same manner, only now the roots are outside the unit circle. Let P be the degree of $\underline{\phi}(L)$, which satisfies $P = \lfloor p/s \rfloor + 1$. Using the homomorphism property of Proposition 1, we have

$$\mathbf{Y}_n = \underline{\delta}^\sharp(L)\, \underline{\phi}(L)^{-1}\, \underline{\theta}(L)\, \mathbf{Z}_n$$
$$\det \underline{\phi}(L)\, \mathbf{Y}_n = \underline{\delta}^\sharp(L)\, \underline{\phi}^\sharp(L)\, \underline{\theta}(L)\, \mathbf{Z}_n.$$

Now we can deduce the ARIMA model structure of down-sampled or aggregated components by left-multiplying by an appropriate row vector: for the jth sub-series (corresponding to season $1 \leq j \leq s$), we have $X_{ns+j} = e_j' \mathbf{X}_n$, which is an ARIMA process written as

$$\det \underline{\delta}(L)\; \det \underline{\phi}(L)\, X_{ns+j} = \vartheta^{(j)}(L)\, \epsilon_n,$$

where $\vartheta^{(j)}(z)$ is a polynomial of degree $D(s-1)+P(s-1)+Q$ of unit leading coefficient, where $Q = \lfloor q/s \rfloor + 1$ and ϵ_n is a white noise of variance σ_j^2. The polynomial is obtained by spectral factorization, and satisfies

$$\vartheta^{(j)}(L)\, \vartheta^{(j)}(L^{-1})\, \sigma_j^2 = e_j'\, \underline{\delta}^\sharp(L)\, \underline{\phi}^\sharp(L)\, \underline{\theta}(L)\, \underline{\theta}(L^{-1})'\, \underline{\phi}^\sharp(L^{-1})'\, \underline{\delta}^\sharp(L^{-1})'\, e_j\, \sigma^2.$$

If instead we consider frequency-aggregation, then we must apply the row vector ι, which is the sum of the e_j vectors. The process $\overline{X}_{ns} = \iota' \mathbf{X}_n$ is an ARIMA process written as

$$\det \underline{\delta}(L)\; \det \underline{\phi}(L)\, \overline{X}_{ns} = \vartheta^{(0)}(L)\, \epsilon_n,$$

where $\vartheta^{(0)}(z)$ is a polynomial of degree $D(s-1)+P(s-1)+Q$ of unit leading coefficient, and ϵ_n is a white noise of variance σ_0^2. The polynomial is obtained by spectral factorization, and satisfies

$$\vartheta^{(0)}(L)\, \vartheta^{(0)}(L^{-1})\, \sigma_0^2 = \iota'\, \underline{\delta}^\sharp(L)\, \underline{\phi}^\sharp(L)\, \underline{\theta}(L)\, \underline{\theta}(L^{-1})'\, \underline{\phi}^\sharp(L^{-1})'\, \underline{\delta}^\sharp(L^{-1})'\, \iota\, \sigma^2.$$

Clearly, other linear combinations of $\mathbf{X}_n$ have dynamics determined in a similar manner. Next, we may ask how the dynamics of the down-sampled or frequency-aggregated processes may be inferred from these new ARIMA structures. It is clear that the new pseudo-autoregressive operator given by $\det \underline{\delta}(L)\; \det \underline{\phi}(L)$ is obtained from $\varphi(z)$ by simply raising each root to the power s. In McElroy (2019) it is argued that the root structure of the pseudo-autoregressive polynomial governs the dynamics of the process, although the moving average structure also plays a role (it can retard or even annihilate persistent oscillatory effects).

A high-frequency process exhibiting seasonality will typically have an autoregressive root of the form $\rho^{-1} e^{i\omega}$ for $\rho \in (0,1)$ (but close to one) and $\omega = 2\pi j/r$ for some $1 \leq j \leq r$, where r is the number of seasons. In passing to a low-frequency process, we down-sample or aggregate by first embedding in a process with s seasons. Hence, these autoregressive roots get mapped to $\rho^{-s} e^{2\pi i j s/r}$, so depending on the relationship of s to r the seasonality could be completely removed, or perhaps just shifted. For example, if $r = 12$ and $s = 3$, as occurs when passing from monthly to quarterly data, we have $2\pi j s/r = \pi j/2$, which takes the values of $\pi/2$ or π depending on j. Hence, the six seasonal frequencies are automatically mapped to the two quarterly frequencies through the autoregressive root structure.

In terms of the ideas discussed in the previous section, our analysis can now be extended to the case that $\{X_t\}$ satisfies (1), and the latent seasonal process $\{S_t\}$ follows some ARIMA model – but the quota noise $\{N_t\}$ is heteroscedastic. Note that by embedding the quota noise, we obtain a stationary process $\mathbf{N}_n$: the s-dimensional process $\{\mathbf{N}_n\}$ is by assumption a white noise sequence with covariance matrix Σ_N satisfying $\iota' \, \Sigma_N \, \iota = 0$ (or more generally, equals a small number). Therefore, the frequency-aggregated process $\overline{X}_{ns} = \iota' \mathbf{X}_n$ still has the ARIMA structure described above, even when quota noise has been added to obfuscate $\{S_t\}$.

4. Benchmarking Methodology

The benchmarking problem involves a time series sampled at a high and low frequency, which for simplicity we suppose to be the monthly and quarterly frequency, respectively. We wish to adjust both the monthly and quarterly data such that the resulting seasonal adjustments are *adequate*, which means that according to some seasonality diagnostic there is no seasonality. The monthly series is denoted $\{X_{t,m}\}$ and the quarterly is denoted $\{X_{i,q}\}$, where $t, i \in \mathbb{Z}$ and are related via $t = 3\,i + j$ for $j = 1, 2, 3$. The data satisfy the following frequency aggregation property:

$$X_{i,q} = X_{3i+1,m} + X_{3i+2,m} + X_{3i+3,m} \tag{3}$$

for the ith quarter. Direct adjustments of the monthly and quarterly series will utilize an N symbol, for nonseasonal, i.e., $\{N_{t,m}\}$ and $\{N_{i,q}\}$ respectively. These direct adjustments need not satisfy (3). If they do not, we seek modifications $\{Y_{t,m}\}$ and $\{Y_{i,q}\}$ that satisfy (3), are close to the direct adjustments, and are adequate. If we have available $\{X_{t,m}\}$, we can compute all other quantities, viz. $\{X_{i,q}\}$, $\{N_{t,m}\}$, and $\{N_{i,q}\}$. However, in some cases only $\{N_{t,m}\}$ is available, in which case we define $\{N_{i,q}\}$ as follows: (i) aggregate $\{N_{t,m}\}$ and test for seasonality; (ii) if adequate we are done, but otherwise seasonally adjust and declare the result to be $\{N_{i,q}\}$.

We ensure the above criteria hold by minimizing the discrepancy between $\{Y_{t,m}\}$ and $\{N_{t,m}\}$, and between $\{Y_{i,q}\}$ and $\{N_{i,q}\}$, while imposing (3) and adequacy of both $\{Y_{t,m}\}$ and $\{Y_{i,q}\}$. Actually, we can just plug (3) into the optimization criterion, yielding for each quarter i

$$\begin{aligned} \mathcal{L}(Y_{3i+1,m}, Y_{3i+2,m}, Y_{3i+3,m}) = & \left(N_{i,q} - \sum_{j=1}^{3} Y_{3i+j,m} \right)^2 / N_{i,q} \\ & + \sum_{j=1}^{3} \left(N_{3i+j,m} - Y_{3i+j,m} \right)^2 / N_{3i+j,m}. \end{aligned} \tag{4}$$

Adequacy is checked by applying a diagnostic δ to both putative solutions for the monthly and quarterly series, which is compared to a threshold α. In McElroy (2018) the QS diagnostic (Maravall, 2012) was utilized, but some problems with spurious detections of seasonality have raised concerns about this method. Here we instead use the root diagnostic of McElroy (2019), which can be adapted to different sampling frequencies, and offers a p-value for rejection of the null hypothesis that seasonality is present to a given degree. Specifically, we examine the p-values as a function of seasonal persistence ρ at frequency $2\pi/4$ (the quarterly seasonal frequency) and $2\pi j/12$ for $j = 1, 2, \ldots, 5$ (the monthly seasonal frequencies), and demand that all ρ, such that the p-value is less than a given α, satisfy $\rho < .98$, this value corresponding to a substantial degree of oscillation in the autocorrelation function. In other words, we wish to enforce that

$$\max_{\rho \in (.98,1)} p(\rho) \leq \alpha, \tag{5}$$

where $p(\rho)$ denotes the p-value as a function of ρ determined by the null hypothesis. Such a condition says that the null hypothesis of seasonality of degree ρ can be rejected at level α for all $\rho \in (.98, 1)$. This threshold .98 can be altered; lowering it demands that even weaker forms of seasonality must also be screened out, so that we are more exacting of our requirements on the adjustment, whereas raising the threshold means we are more relaxed in our standards. Our notation for the corresponding constraints is

$$\delta\{Y_{1,m}, \ldots, Y_{3i+3,m}\} \leq \alpha, \quad \delta\{Y_{1,q}, \ldots, Y_{i,q}\} \leq \alpha, \tag{6}$$

where δ indicates the maximum of p-values (5) computed on either the monthly or quarterly data. Hence, we seek to minimize (4) subject to (6). One approach is to use Lagrangian techniques with inequality constraints (cf. Kuhn and Tucker (1951)), or utilize the introduction of slack variables. Instead, following an approach reviewed in Smith and Coit (1995), we convert the constrained minimization problem into a penalized minimization, and iteratively increase the penalty. This is accomplished by introducing tuning parameters $\omega_m, \omega_q > 0$ and minimizing

$$\begin{aligned}\sum_i \mathcal{H}(Y_{3i+1,m}, Y_{3i+2,m}, Y_{3i+3,m}) = & \sum_i \mathcal{L}(Y_{3i+1,m}, Y_{3i+2,m}, Y_{3i+3,m}) \\ & + \omega_m \left(\min\left[\alpha - \delta\{Y_{1,m}, \ldots, Y_{3i+3,m}, \ldots\}, 0\right]\right)^2 \\ & + \omega_q \left(\min\left[\alpha - \delta\{Y_{1,q}, \ldots, Y_{i,q}, \ldots\}, 0\right]\right)^2.\end{aligned}$$

The rationale is that each of the two penalty terms is zero if and only if $\delta \leq \alpha$; so if $\delta > \alpha$, a large positive value is added to the objective function, and such solutions will tend to be rejected. In practice, if inadequate solutions are obtained then one must adjust ω_m or ω_q upwards. Our own implementation uses an iterative scheme, whereby the optimization problem is initialized with prior solutions and we gradually increment ω_m and ω_q (starting at an initial value of zero, which enforces no adequacy) until (6) is satisfied.

Now it is possible that an adequate solution is obtained immediately, where $\omega_m = \omega_q = 0$. This can happen because the mere attempt to impose aggregation constraints can yield adequate adjustments, since the new monthly values $Y_{3i+j,m}$ are forced to resemble the given direct monthly adjustments $N_{3i+j,m}$, which are adequate themselves. More generally, if the initial solution with $\omega_m = \omega_q = 0$ yields inadequate monthly or quarterly reconciled series, then we must gradually increment these penalties, repeating until the diagnostics' conditions are satisfied. For this procedure to be numerically feasible, it is essential for the diagnostics to be quickly computed. The root diagnostics (McElroy, 2019) require no nonlinear optimization, and the critical values can be speedily obtained using Monte Carlo.

Because the optimization portion of the reconciliation procedure is the major bottleneck, it is preferable to use an algorithm that will find a solution quickly. For the examples in this paper, we use the Bound Optimization by Quadratic Approximation (BOBYQA) algorithm of Powell (2009), as implemented in the minqa (v1.2.4; Bates et al., 2014) package in R. As the name of the algorithm suggests, the algorithm seeks to minimize an objective function $F(x)$, where x is an n-dimensional input, subject to some bounds on the elements of x; the bounds, however, can be arbitrarily large. The algorithm iteratively constructs a quadratic approximation Q to the objective function F such that, over sets of interpolation points y that are chosen and adjusted automatically, $Q(y) = F(y)$. The model updates by minimizing the Frobenius norm of the the change to the second derivative matrix of Q; it requires no calculation of the first derivative of the objective function itself.

5. Applications

We have examined data for 49 economic series taken from various surveys conducted by the U.S. Census Bureau. These series measure quantities such as inventories and shipments, construction spending, or imports and exports. Using the same range for ρ as described above and an α level of 0.1 for the root diagnostic, the majority of these series are such that the null hypothesis of seasonality of degree ρ is rejected at both a monthly level and an aggregated quarterly level. The applications used for illustration are drawn from the remaining subset of series for which the raw monthly series appears to have some weak seasonality present, but is deemed nonseasonal (and thus left as is), while the resulting quarterly aggregate is more noticeably seasonal. For these examples, we have used a discrete set of values for ρ, incrementing from 0.98 to 1 by 0.001, and an α level of 0.1 for the root diagnostic. Also, instead of using a fixed order p for the autoregressive polynomial in the root diagnostic, we have opted to allow the order p to be chosen using a selection criterion (e.g., AIC); McElroy (2019) discusses the advantages and disadvantages of these two approaches. Our optimization starts with an initial value of zero for both ω_m and ω_q, incrementing each by 1000 should the solution fail to satisfy the conditions required for adequacy.

5.1 Example 1: Imports of Steelmaking and Ferroalloying Materials

We first examine International Trade data for imports of steelmaking and ferroalloying materials.[1] The span we will consider is from January 1991 through December 2005. Figure 6 shows the monthly series and its quarterly aggregate. Figure 7 provides the corresponding spectral density and autocorrelation functions for the differenced, log transformed versions of both the monthly series and its quarterly aggregate. Visually, the spectral density for the monthly series does not seem to suggest that there is any seasonality present, nor are the values of the autocorrelation function particularly pronounced at the seasonal lags. For the aggregated series, however, there is a peak that is close to the quarterly frequency in the spectral density, and there is a large value for the autocorrelation function at the second seasonal lag (lag 8). Additionally, computing the QS statistic for the differenced, log transformed versions of both series suggests that while the monthly series is not a strong candidate for seasonal adjustment (p-value of 0.137), the quarterly aggregated series is seasonal (p-value of 0.008).

We apply the root diagnostic to the monthly series, the quarterly aggregate, the monthly seasonal adjustment, and indirect and direct quarterly seasonal adjustments. Note that since the monthly series is not adjusted, the results for monthly series and monthly seasonal adjustment should be more or less the same, as should the results for the quarterly aggregate and the indirect quarterly seasonal adjustment – since the critical values are determined by Monte Carlo, some minor fluctuations should be expected.

[1] Available online at `https://www.bea.gov/international/detailed-trade-data` as historical data under the IDS-0182 heading. Numbers are subject to error arising from a variety of sources, including nonsampling error. For more information, refer to `https://www.census.gov/foreign-trade/index.html` or `https://www.bea.gov/resources/methodologies/us-international-economic-accounts-concepts-methods`.

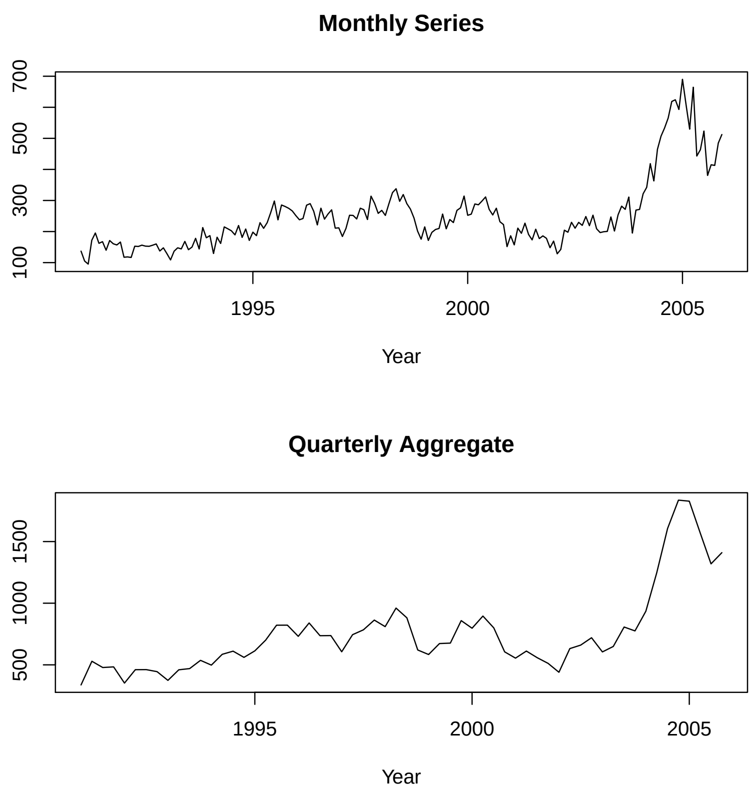

Figure 6: Monthly series for imports of steelmaking and ferroalloying materials and its quarterly aggregate. Data sourced from Bureau of Economic Analysis, U.S. Trade in Goods.

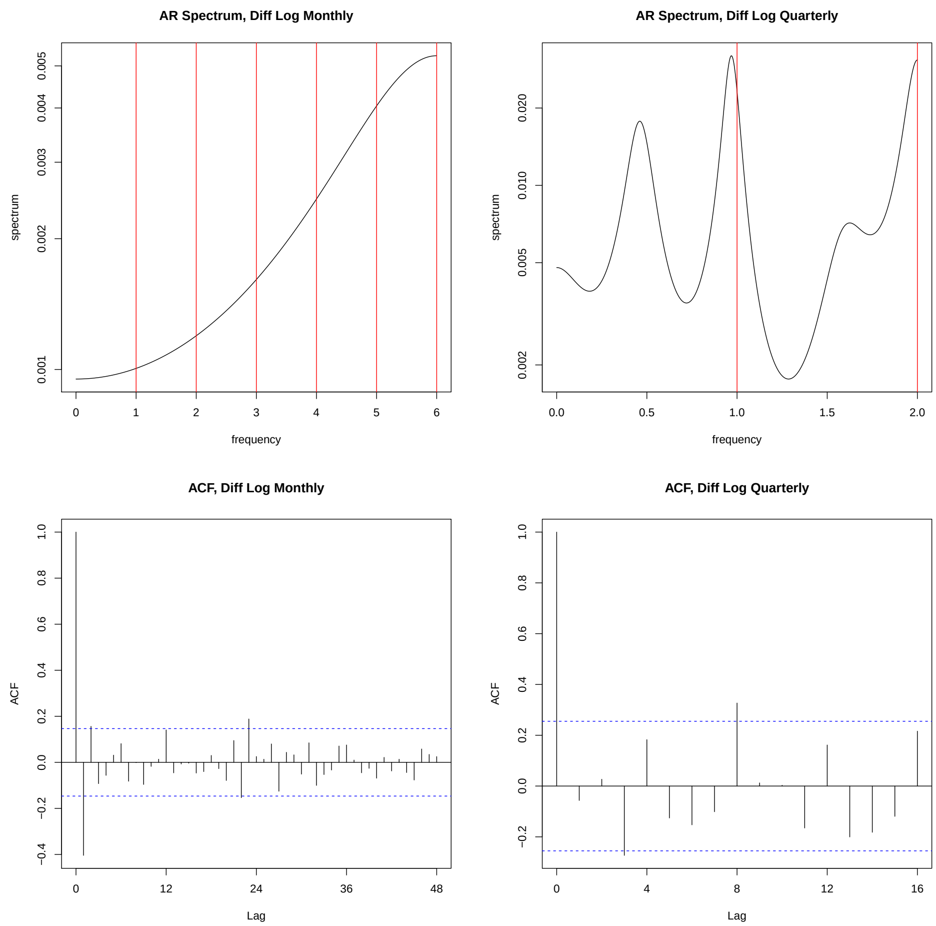

Figure 7: Spectral densities and autocorrelation functions for monthly import series and quarterly aggregate. Data sourced from Bureau of Economic Analysis, U.S. Trade in Goods.

Table 1 shows the values of ρ for which the specified series is deemed seasonal by the root diagnostic. The empty set for the monthly series indicates that the series is not considered seasonal for any of the tested values of ρ, while the quarterly aggregate appears to be seasonal for most of the tested values. Hence, we apply the optimization step, finding an adequate solution with $\omega_m = \omega_q = 0$. The last two rows of Table 1 show that the resulting reconciled monthly series and its quarterly aggregate are not viewed as seasonal by the root diagnostic. Figure 8 visually supports this, as neither the spectrum nor the autocorrelation function suggests that there is seasonality still present in either the reconciled monthly or quarterly series. Figures 9 and 10 provide comparisons of the reconciled monthly series to the original monthly series and the reconciled quarterly series to the original quarterly aggregate. For the most part, we see that the reconciliation (the red lines) does not make drastic alterations to the original monthly and quarterly series (the black lines); rather, it seems to attenuate some of the more extreme fluctuations observed in the original series.

Series	ρ
Monthly	$\varnothing$
Qtrly Agg	$[0.980, 0.994]$
Monthly SA	$\varnothing$
Indirect Qtrly SA	$[0.980, 0.994]$
Direct Qtrly SA	$\varnothing$
Reconciled Mthly	$\varnothing$
Reconciled Qtrly	$\varnothing$

Table 1: Values of ρ for which the root diagnostic applied to the given series of import data has a p-value exceeding $\alpha = 0.1$. Data sourced from Bureau of Economic Analysis, U.S. Trade in Goods.

5.2 Example 2: State and Local Construction Spending for Hospitals

Next, we look at state and local construction spending for hospitals, where the quantity being measured is the value of construction put in place.[2] The span for consideration here is January 1993 through March 2010. Figure 11 displays the monthly series and its quarterly aggregate. In addition, we display the spectral densities and autocorrelation functions for the corresponding differenced, log transformed series in Figure 12. The monthly series seems to be weakly seasonal, as the peaks in the spectral density are slightly offset from the seasonal frequencies associated with monthly series, and the values of the autocorrelation function for the first two seasonal lags are large. Similarly, the peak in the spectral density for the quarterly aggregate is also slightly offset from the relevant frequency for quarterly series, although there is a fairly pronounced spike in the autocorrelation function at the second seasonal lag. The QS diagnostic for both of these series yields a p-value of 0, so the diagnostic indicates that both the monthly series and its quarterly aggregate are actually seasonal.

Table 2 shows the values of ρ for which the specified series is deemed seasonal by the root diagnostic. The empty set for the monthly series indicates that the series is not considered seasonal for any of the tested values of ρ, while the quarterly aggregate appears

[2]Numbers for the Value of Construction Put in Place (VIP) Survey were found on `https://www.census.gov/construction/c30/historical_data.html`. These numbers are subject to sampling and nonsampling error; methodology for this survey can be found at `https://www.census.gov/construction/c30/methodology.html`.

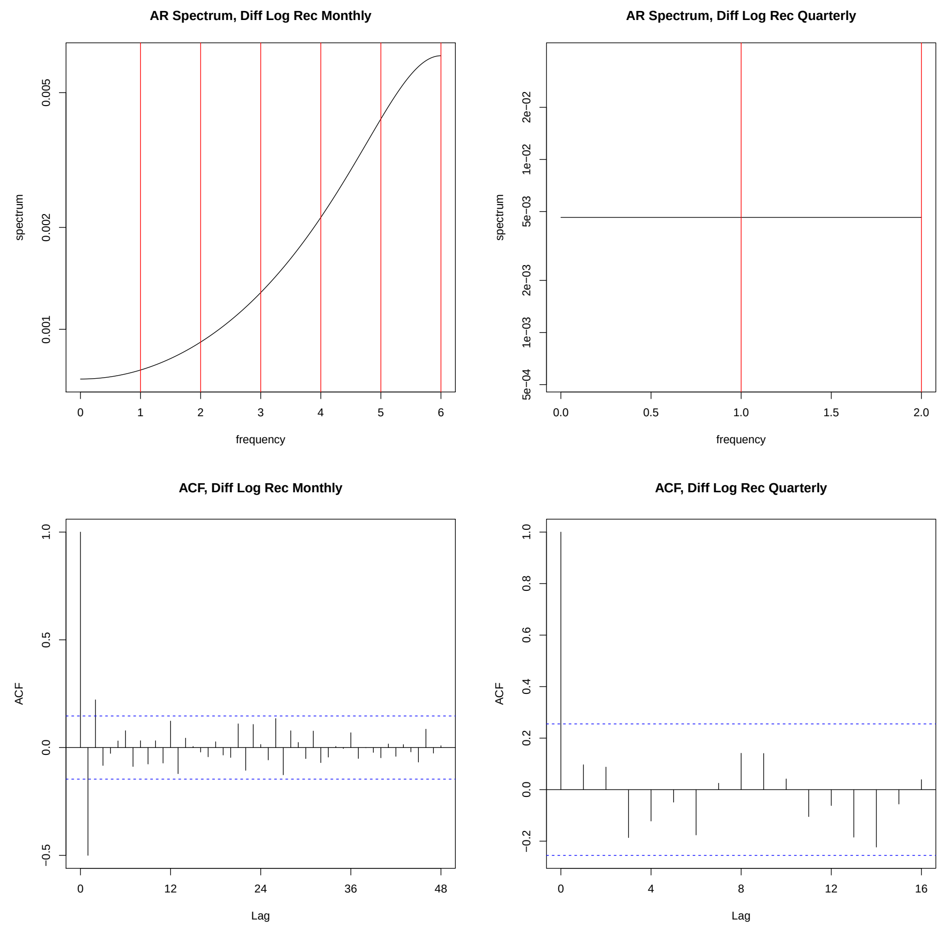

Figure 8: Spectral densities and autocorrelation functions for reconciled monthly and quarterly import series. Reconciliation applied to data sourced from Bureau of Economic Analysis, U.S. Trade in Goods.

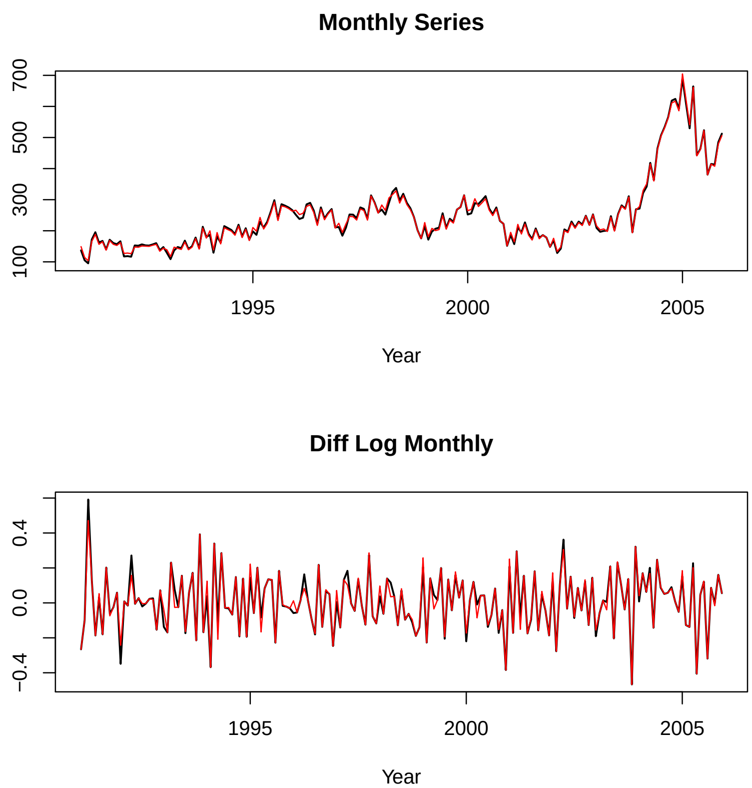

Figure 9: Comparisons of monthly series (black) to reconciled monthly series (red) for import series. Data sourced from Bureau of Economic Analysis, U.S. Trade in Goods.

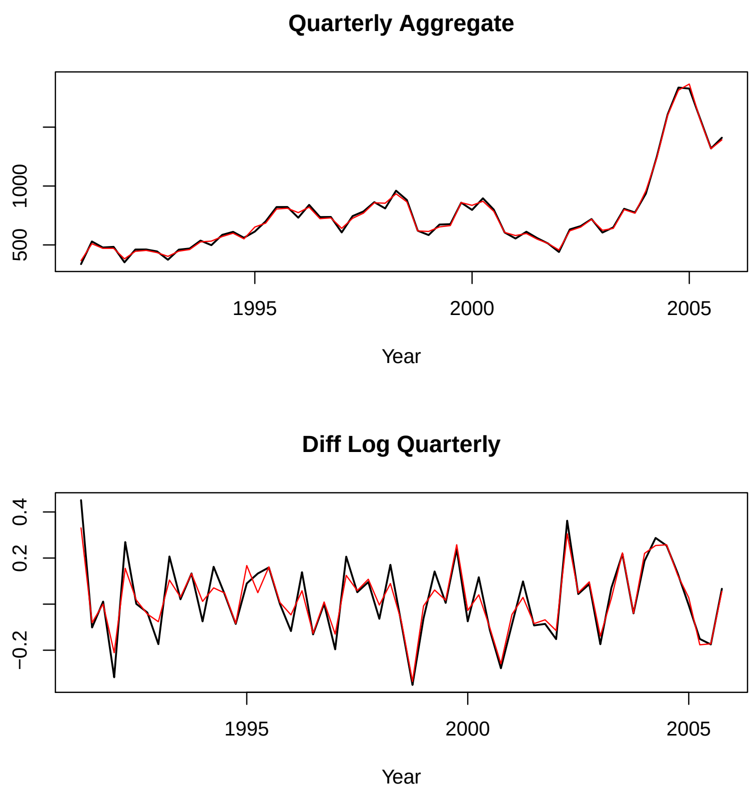

Figure 10: Comparisons of quarterly aggregate (black) to reconciled quarterly aggregate (red) for import series. Data sourced from Bureau of Economic Analysis, U.S. Trade in Goods.

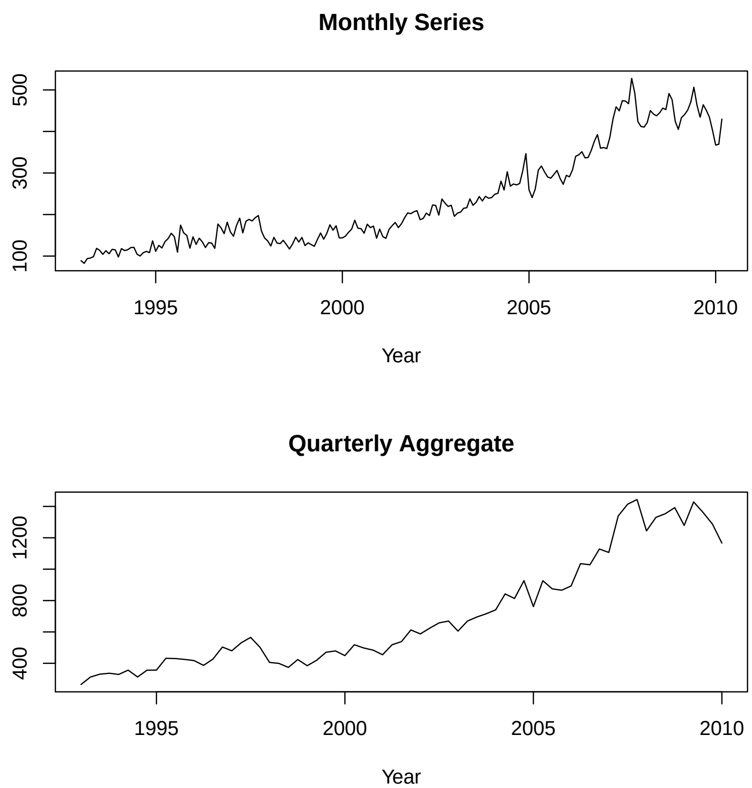

Figure 11: Monthly construction spending series and its quarterly aggregate. Data sourced from U.S. Census Bureau, VIP Survey.

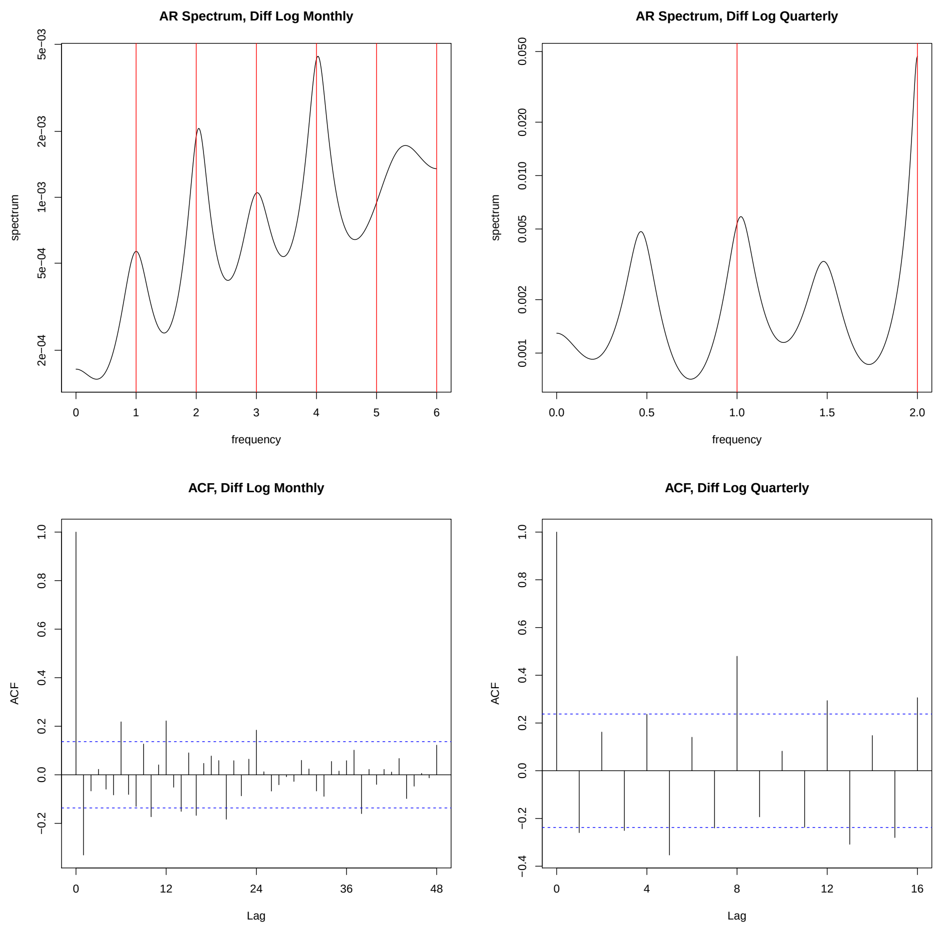

Figure 12: Spectral densities and autocorrelation functions for monthly construction spending series and quarterly aggregate. Data sourced from U.S. Census Bureau, VIP Survey.

to be seasonal for the tested values on the lower end of the interval. Again, we apply the optimization step, finding an adequate solution with $\omega_m = \omega_q = 0$. The last two rows of Table 2 show that the resulting reconciled monthly series and its quarterly aggregate are not classified as being seasonal by the root diagnostic. Looking at Figure 13, we see that in the reconciled series the peaks that were previously seen in the spectral densities have been eliminated. Also, the values of the autocorrelation functions at the seasonal lags have decreased relative to what they were for the original series. Figures 14 and 15 provide comparisons of the reconciled monthly series to the original monthly series, the reconciled quarterly series to the original quarterly aggregate, and the reconciled quarterly series to the direct seasonal adjustment of the quarterly aggregate. We see that the changes to the original monthly series induced by the reconciliation process are minor. The quarterly aggregate is more affected, though; it can be seen that some of the sharper changes in the original aggregated series have been damped during the reconciliation process.

Series	ρ
Monthly	$\varnothing$
Qtrly Agg	$[0.980, 0.985]$
Monthly SA	$\varnothing$
Indirect Qtrly SA	$[0.980, 0.985]$
Direct Qtrly SA	$\varnothing$
Reconciled Mthly	$\varnothing$
Reconciled Qtrly	$\varnothing$

Table 2: Values of ρ for which the root diagnostic applied to the given series of construction spending data has a p-value exceeding $\alpha = 0.1$. Data sourced from U.S. Census Bureau, VIP Survey.

6. Conclusion

Residual seasonality in a quarterly series can occur when a monthly series is aggregated to a quarterly frequency. This is true even if the monthly series is nonseasonal. Given a seasonally adjusted monthly series, the detection of seasonality in the corresponding quarterly aggregate is undesirable. But the simplest method of ensuring a quarterly series is nonseasonal, viz. direct seasonal adjustment, has the drawback of not preserving accounting relationships. That is, the quarterly aggregate of a monthly seasonally adjusted series will usually not be equal to the seasonal adjustment of the quarterly aggregated raw series. For economists, this situation is equally undesirable.

What we propose is a procedure that is similar to benchmarking in that it makes small modifications to series, such that the accounting relationship between the monthly series and its quarterly aggregate is preserved. Whereas seasonality has typically not been a consideration when benchmarking has been applied (McElroy (2018) seems to be the first case of incorporating benchmarking with seasonality diagnostics), our idea attempts to ensure that the modifications are done such that the reconciled monthly series and its quarterly aggregate are both nonseasonal (i.e., adequately seasonally adjusted). We do so by viewing the benchmarking problem as a penalized minimization problem – our modified objective function includes penalty terms for solutions that are still deemed seasonal.

We illustrate the procedure with two economic series that are readily available; note that the raw components of GDP are protected against publication, and therefore could not be used as illustrations in this paper. The two monthly series considered are either

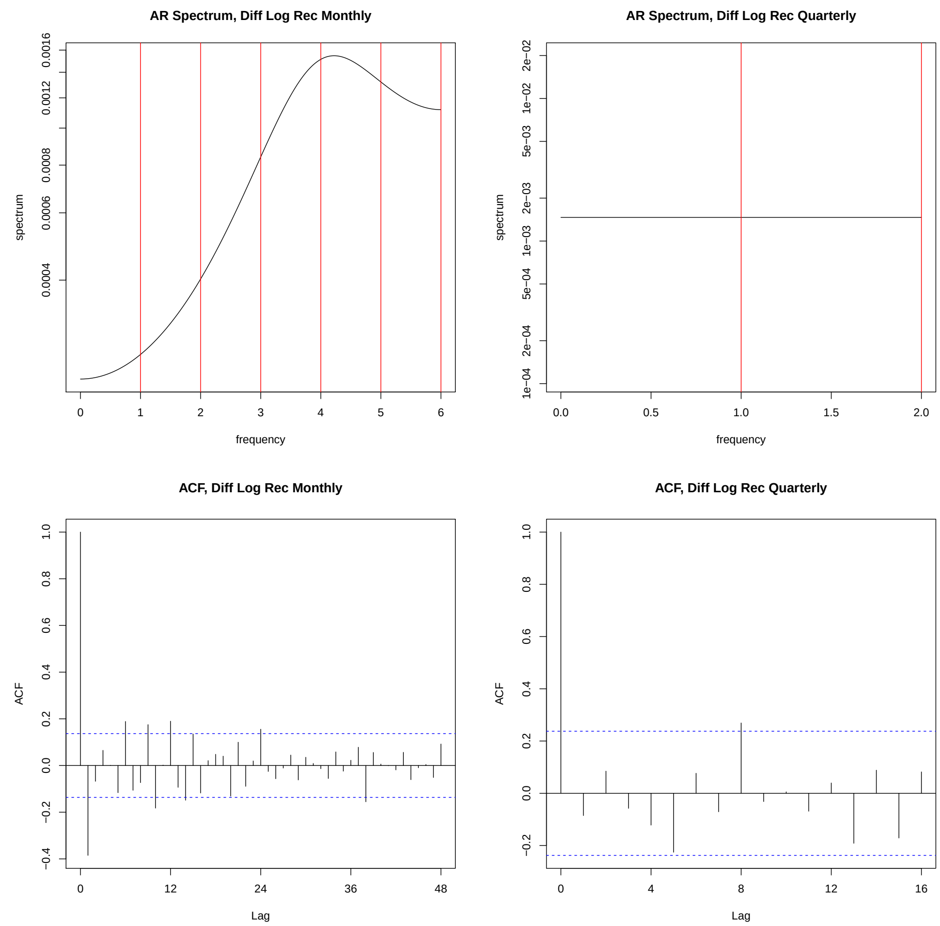

Figure 13: Spectral densities and autocorrelation functions for reconciled monthly and quarterly construction spending series. Reconciliation applied to data sourced from U.S. Census Bureau, VIP Survey.

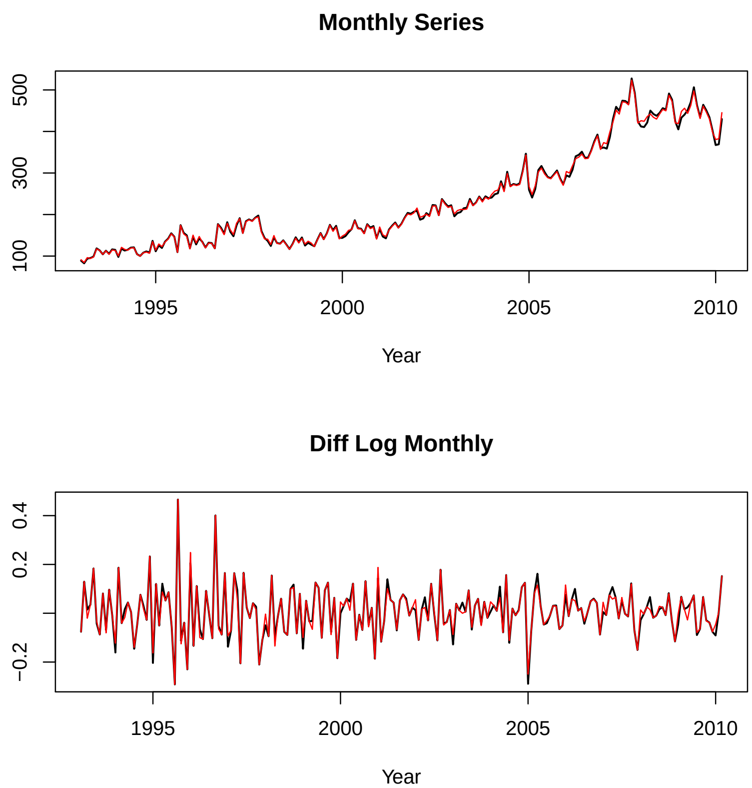

Figure 14: Comparisons of monthly series (black) to reconciled monthly series (red) for construction series. Data sourced from U.S. Census Bureau, VIP Survey.

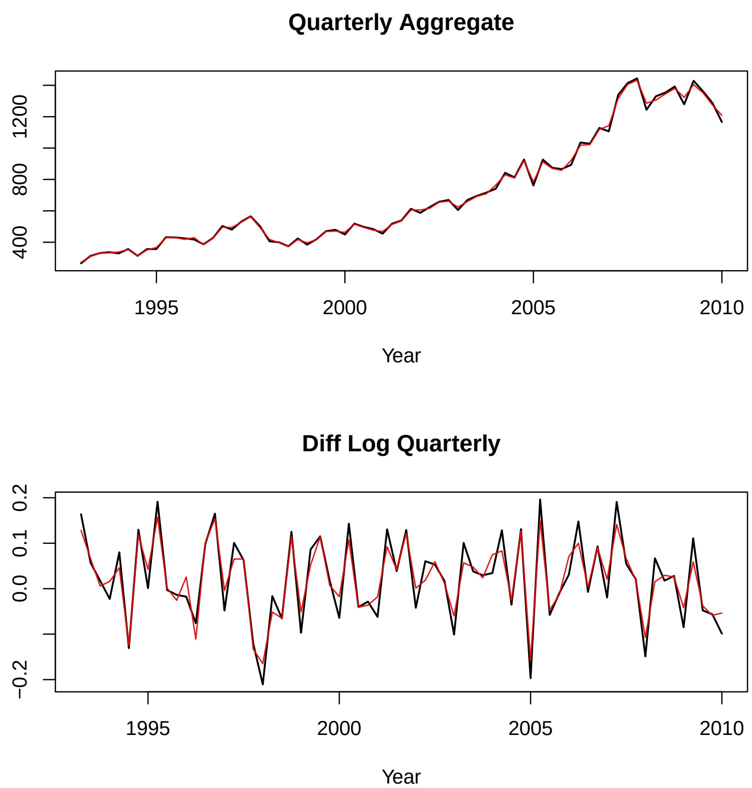

Figure 15: Comparisons of quarterly aggregate (black) to reconciled quarterly aggregate (red) for construction series. Data sourced from U.S. Census Bureau, VIP Survey.

nonseasonal or weakly seasonal, and neither is seasonally adjusted. However, both of their quarterly aggregates are more noticeably seasonal. When we use our procedure on these series, the final outcomes are not considered seasonal for monthly or quarterly frequencies, and the accounting relationship is preserved.

One important facet of the methodology is the determination of autoregressive order p in the root diagnostic. Although results work fairly well when there are no unit roots present in the data – and either using a large fixed p or determining p via AIC can yield sensible diagnostics – the asymptotic framework deteriorates when roots are close to unit, say of magnitude .998 or greater. In fact, the root diagnostics of McElroy (2019) have low power against unit root alternatives, because the asymptotic theory is predicated upon stationarity; either method of choosing p is likely to give fallacious results in such a scenario. However, these are cases where the seasonality is more obvious anyways, such that seasonal adjustment would more obviously be routinely applied.

Given this reality, and the absence of a diagnostic that simultaneously addresses the detection of nonstationary (unit root) seasonality as well as stationary (dynamic) seasonality, we recommend that the root diagnostic be used in tandem with practitioner judgment as to the seasonality apparent in a given monthly series, with other exploratory tools – such as autocorrelation plots and spectral density plots – being utilized to screen out the important unit root case. We note that the discrimination of nonstationary versus stationary seasonality is acknowledged to be a challenging problem, and therefore our examples here have focused on the cases where the seasonality present in the monthly series is clearly of a weaker, dynamic variety.

When the root diagnostic determines that the monthly series (or monthly seasonally adjusted series) and the quarterly aggregated counterpart are nonseasonal, the optimization algorithm is not invoked. However, when the optimization step is required, then the practicality of this proposed procedure depends on the speed with which a solution can be found. Our examples used the BOBYQA algorithm from the outset, but as noted earlier, it is possible to obtain an adequate solution using initial values of 0 for both ω_m and ω_q. When both are 0, the penalty terms in the objective function disappear, turning this into a fairly straightforward problem of optimizing the $\mathcal{L}$ expression seen in (4), which can be solved directly. Hence, we can design it so that the optimization algorithm is only required if ω_m or ω_q is positive-valued, with the exact solution used in the case where both are 0. This will yield a fast first step with an adequate outcome in many cases.

REFERENCES

Astolfi, T., Ladiray, D., and Mazzi, G.L. (2001), "Seasonal Adjustment of European Aggregates: Direct versus Indirect Approach," *working paper, Eurostat*, Luxembourg.

Bates, D., Mullen, K.M., Nash, J.C., and Varadhan, R. (2014), "minqa: Derivative-free Optimization Algorithms by Quadratic Approximation," *R package version 1.2.4*, `https://CRAN.R-project.org/package=minqa`.

Chen, B. (2012), "A Balanced System of U.S. Industry Accounts and Distribution of the Aggregate Statistical Discrepancy by Industry," *Journal of Business and Economic Statistics*, 30(2), 202–221.

Cholette, P.A. (1984), "Adjusting Sub-annual Series to Yearly Benchmarks," *Survey Methodology*, 10(1), 35–49.

Dagum, E.B. (1979), "On the Seasonal Adjustment of Economic Time Series Aggregates: A Case Study of the Unemployment Rate, Counting the Labor Force," *National Commission on Employment and Unemployment Statistics*, Appendix, 2: 317–344.

Dagum, E.B., and Cholette, P.A. (2006), *Benchmarking, Temporal Distribution, and Reconciliation Methods for Time Series Data*, Springer-Verlag, New York. Lecture Notes in Statistics 186.

Denton, F. (1971), "Adjustment of Monthly or Quarterly Series to Annual Totals: An Approach Based on Quadratic Minimization," *Journal of the American Statistical Association*, 82, 99–102.

Di Fonzo, T., and Marini, M. (2011), "Simultaneous and Two-step Reconciliation of Systems of Time Series: Methodological and Practical Issues," *Journal of the Royal Statistical Society C*, 60(2), 143–164.

Findley, D.F., Lytras, D., and McElroy, T.S. (2017), "Detecting Seasonality in Seasonally Adjusted Monthly Time Series," *Census Bureau Research report* 2017-03.

Franses, P.H. (1994) "A Multivariate Approach to Modeling Univariate Seasonal Time Series," *Journal of Econometrics*, 63, 133–151.

Furman, J. (2015), "Second Estimate of GDP for the First Quarter of 2015," *Council of Economic Advisers Blog*, May 29, 2015.

Gilbert, C.E., Morin, N.J., Paciorek, A.D., and Sahm, C.R. (2015), "Residual Seasonality in GDP," *FEDS Notes*, May 14, 2015.

Gladyshev, E.G. (1961), "Periodically Correlated Random Sequences," *Soviet Mathematics*, 2, 385–388.

Groen, J., and Russo, P. (2015), "The Myth of First-Quarter Residual Seasonality," *Liberty Street Economics*, June 8, 2015.

Hood, C.H., and Findley, D.F. (2001), "Comparing Direct and Indirect Seasonal Adjustments of Aggregate Series," *American Statistical Association Proceedings of the Business and Economics Statistics Section*.

Kuhn, H.W., and Tucker, A.W. (1951) "Nonlinear Programming," *Proceedings of 2nd Berkeley Symposium*. Berkeley: University of California Press, 481–492.

Lengermann, P., Morin, N., Paciorek, A., Pinto, E., and Sahm, C. (2017) "Another Look at Residual Seasonality in GDP," *FEDS Notes*. Washington: Board of Governors of the Federal Reserve System. `https://doi.org/10.17016/2380-7172.2031`

Maravall, A. (2012), "Update of Seasonality Tests and Automatic Model Identication in TRAMO-SEATS," *Bank of Spain (November 2012 draft)*.

McCulla, S. H. and Smith, S. (2015) "Preview of the 2015 Annual Revision of the National Income and Product Accounts," *Survey of Current Business* 95 (June), `https://apps.bea.gov/scb/pdf/2015/06%20June/0615_preview_of_2015_annual_revision_of_national_income_and_product_accounts.pdf`.

McElroy, T. (2016), "Challenges with Seasonal Adjustment," *FESAC meeting slides*, `https://www.census.gov/about/adrm/fesac/meetings/2016-12-09-meeting.html`.

McElroy, T. (2018), "Seasonal Adjustment Subject to Accounting Constraints," *Stat. Neerl.*, 72, 574–589.

McElroy, T. (2019), "A Diagnostic for Seasonality Based upon Autoregressive Roots," *Preprint*.

Moulton, B.R., and Cowan, B.D. (2016), "Residual Seasonality in GDP and GDI: Findings and Next Steps," *Survey of Current Business* 96 (July), `https://apps.bea.gov/scb/pdf/2016/07%20July/0716_residual_seasonality_in_gdp_and_gdi.pdf`.

Osborn, D.R. (1991), "The Implications of Periodically Varying Coefficients for Seasonal Time-series Processes," *Journal of Econometrics*, 28, 373–384.

Phillips, K. and Wang, J. (2016), "Residual Seasonality in U.S. GDP Data," FRB of Dallas Working Paper No. 1608. `https://ssrn.com/abstract=2866843` or `http://dx.doi.org/10.24149/wp1608`.

Powell, M.J. (2009), "The BOBYQA Algorithm for Bound Constrained Optimization without Derivatives," *Report No. DAMTP 2009/NA06*, Centre for Mathematical Sciences, University of Cambridge, UK, `http://www.damtp.cam.ac.uk/user/na/NA_papers/NA2009_06.pdf`.

Quenneville, B., and Fortier, S. (2012), "Restoring Accounting Constraints in Time Series Methods and Software for a Statistical Agency," in *Economic Time Series: Modelling and Seasonality*, eds. W.R. Bell, S.H. Holan and T.S. McElroy, Boca Raton: CRC Press, Taylor and Francis Group, pp. 231–253.

Rudebusch, G.D., Wilson, D., and Mahedy, T. (2015), "The Puzzle of Weak First-Quarter GDP Growth," *FRBSF Economic Letter*, May 18, 2015.

Smith, A.E., and Coit, D.W. (1995), "Penalty Functions," in *Handbook of Evolutionary Computation*, 97(1), p.C5.

Tiao, G.C., and Grupe, M.R. (1980), "Hidden Periodic Autoregressive-Moving Average Models in Time Series Data," *Biometrika*, 67, 365–373.