DEUTSCHE
BUNDESBANK
EUROSYSTEM

Discussion Paper

Deutsche Bundesbank
No 07/2026

Redesigning the classical automatic selection of X-11 seasonal filters

Karsten Webel

Discussion Papers represent the authors' personal opinions and do not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

Editorial Board: Daniel Foos
Stephan Jank
Thomas Kick
Martin Kliem
Malte Knüppel
Christoph Memmel
Hannah Paule-Paludkiewicz

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank, Press and Public Relations Division, at the above address or via email: www.bundesbank.de/contact

Internet http://www.bundesbank.de

Reproduction permitted only if source is stated.

DOI https://doi.org/10.71734/DP-2026-7
ISBN 978-3-98848-064-4
ISSN 2941-7503

Bundesbank Discussion Paper No 07/2026

Redesigning the classical automatic selection of X-11 seasonal filters*

Karsten Webel
Deutsche Bundesbank

Abstract

The classical X-11 seasonal adjustment method for monthly and quarterly time series is equipped with routines for data-driven selections of both Henderson trend-cycle filters and $3 \times k$ seasonal moving averages, currently involving up to three candidate filters in either case. Although these routines have a long-standing tradition that can be traced back at least to 1960, they have not been adopted in a recent JDemetra+ implementation of a modified X-11 method tailored to the specifics of infra-monthly time series, such as the coexistence of multiple seasonal patterns with potentially fractional periodicities. Focusing on seasonal moving averages, we seek to fill this gap by suggesting a generic redesign of the legacy selection concept based upon the so-called moving seasonality ratio. This blueprint utilises a broader set of candidate seasonal filters and, unlike the original setting, a set of common approaches for deriving the requisite asymmetric variants. Considering intersections of multiple approach-specific selection rules stabilises the final filter choice and, what is more, naturally provides the warranted thresholds controlling the potential recalculation of the moving seasonality ratio from suitably shortened detrended observations. Our proposed redesign is illustrated using one specific rule based upon threshold quartiles and real-time data for three German macroeconomic time series sampled at quarterly, monthly, and daily intervals. The last example also highlights the need for additional intermediate steps in the calculation of the moving seasonality ratio when the data contain complex seasonal dynamics.

Keywords: asymmetric linear filters, concurrent revision policy, JDemetra+, moving seasonality ratio, nonparametric seasonal adjustment, real-time data.

JEL classification: C01, C02, C14, C22, C40, C50.

*Contact address: Karsten Webel, Deutsche Bundesbank, Central Office, Directorate General Data & Statistics and Research Centre, Postfach 10 06 02, 60006 Frankfurt am Main, Germany. Phone: +49 69 9566 32702. E-mail: karsten.webel@bundesbank.de. The author thanks Dominique Ladiray and Steve Matthews for providing historical STATCAN papers about the X-11 seasonal adjustment method, Lea Hengen for programming support, and Guy Mélard for valuable comments and suggestions. Discussion Papers represent the authors' personal opinions and do not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem.

1 Motivation

In 1967, the U.S. Bureau of the Census introduced the X-11 seasonal adjustment method for monthly and quarterly time series (Shiskin, Young, and Musgrave, 1967), which quickly made a name for itself in official statistics and therefore has become the cornerstone of many subsequent seasonal adjustment programs, such as X-11-ARIMA (Dagum, 1978, 1980), X-11-ARIMA/88 (Dagum, 1988), X-12-ARIMA (Findley, Monsell, Bell, Otto, and Chen, 1998), X-13ARIMA-SEATS (U.S. Census Bureau, 2025) and JDemetra+ (Grudkowska, 2017; Smyk, Barthelemy, Filgueira, Perez, Sande, and Webel, 2025). These programs have progressively implemented numerous ground-breaking enhancements of the original X-11 routines; besides advanced seasonality diagnostics, those enhancements have especially been centred around automatic procedures that should facilitate mass production, ranging from data-driven selection rules for trend-cycle and seasonal filters to log/level tests and routines for outlier detection and ARIMA model identification, which significantly improved regression-based data pretreatment. However, despite some early experiments with a variant for weekly time series (Shiskin, 1969), the X-11 method has remained applicable only to monthly and quarterly data throughout this development process, presumably due to a lack of demand for considering other (less traditional) data types in official statistics.

During the past decade, however, this narrative has changed slowly but surely as a result of the ongoing digital transformation in many statistical agencies around the world, which, for example, have explored new ways to blend traditional survey data with administrative and digital data sources (Jarmin, 2019; Radermacher, 2019, 2020). In many countries, the demand for weekly and especially daily economic data additionally skyrocketed in the wake of the COVID-19 pandemic outbreak in 2020 when the timely tracking of rapid economic changes became crucial virtually overnight. In sum, these circumstances have sparked a general interest in developing tailored modelling and seasonal adjustment methods for unconventional data types that often exhibit multi-faceted seasonal dynamics, including fractional periodicities, see, for instance, the overviews given in Ladiray, Palate, Mazzi, and Proietti (2018) and Webel (2022). One such tailored method is a modification of the classical X-11 method available in a recent release of JDemetra+ (Webel and Smyk, 2024). This extended X-11 variant essentially inherits the linear filtering operations from the classical method, but is applicable to time series containing any number of seasonal patterns with any set of periodicities; however, it currently facilitates no data-driven selection of trend-cycle or seasonal filters.

In general, such automatic procedures could be built upon classical X-11 noise-to-signal ratios, such as the so-called I/C ratio for picking the Henderson trend-cycle filter (Shiskin et al., 1967) or the moving seasonality ratio (MSR) for picking the seasonal filter (Lothian, 1978, 1984). An alternative strategy that simultaneously selects both filters has emerged after the successful introduction of ARIMA model-based (AMB) seasonal adjustment techniques, which began with the trailblazing contributions of Burman (1980) and Hillmer and Tiao (1982), see also Bell and Hillmer (1984) for an early overview. The key idea of those strategies is to calculate first an MSE-optimal AMB seasonal adjustment, often using relatively simple ARIMA models, such as the so-called Airline model. Afterwards, the combination of X-11 trend-cycle and seasonal filters that produces the seasonally adjusted figures closest to this AMB benchmark is identified. Proximity in this context

is typically measured in the spectral domain through the absolute or squared difference between the transfer functions of the symmetric AMB and (linearly approximated) X-11 seasonal adjustment filters, see, for example, Chu, Tiao, and Bell (2012) and Depoutot and Planas (1998).

Trend-cycle extraction in the extended X-11 variant can be carried out utilising kernel-based local polynomial regression filters (Proietti and Luati, 2008) or reproducing kernel Hilbert spaces (Dagum and Bianconcini, 2008, 2013, 2015), see also Quartier-la-Tente (2024) for an illustration. Encompassing the classical Henderson filters, either framework provides multiple options for specifying the constituent kernel, the polynomial trend degree to be preserved, and the way asymmetric filters should be derived from symmetric ones. Given this vast set of options, seeking an automatic joint selection of X-11 trend-cycle and seasonal filters founded on the AMB benchmark strategy seems to be an inconvenient starting point. The present paper thus focuses on seasonal filters alone and pursues an automatic selection based upon a redesign of the classical MSR criterion. As shown in Table 1, MSR-based filter selection has a long-standing tradition that can be traced back at least to 1960, and the two selection rules, suggested originally in Lothian (1978) for the classical X-11 method, have gradually morphed into the single rule that is currently in effect in X-12-ARIMA, X-13ARIMA-SEATS, and JDemetra+. Along this path, Lothian's original rules have undergone several more or less undocumented changes with respect to the involved seasonal filters and the thresholds separating the selection and recalculation intervals; the latter indifference regions were not foreseen in the original rules, but have been added later to shield the automatic filter selection against frequent revisions when new data become available. However, despite these improvements, one may still argue that those rules lack a common concept for deriving the requisite concurrent seasonal filters. In fact, in the words of Section 3, the asymmetric $3 \times k$ seasonal filters have been calculated using extrapolated averages for $k \in \{1, 3, 5\}$, Musgrave's formula (20) for $k = 9$, and the cut-and-redistribute approach (17) with $l = 5$ for $k = 15$, see also the comments in Ladiray and Quenneville (2004) and Shiskin and Eisenpress (1957).

Our aim is to devise redesigned MSR-based selection rules for the seasonal filters available in the extended X-11 variant. The underlying concept should be generic and as general as possible in order to facilitate easy reproduction as well as further extensions. In particular, we utilise the exact same approach to deriving asymmetric filters for each $3 \times k$ seasonal moving average, consider intersections of multiple approach-specific classical MSR-based selection rules to stabilise the final filter choice, and, as a byproduct, naturally obtain the MSR recalculation intervals through the process of intersecting. The conceptual foundation is laid out in Section 2, which reviews the key steps of the classical X-11 method and the MSR-based filter selection originally developed in Lothian (1978, 1984). Our exposition of X-11 is rather brief and non-technical as detailed descriptions are readily available elsewhere (e.g. Ladiray and Quenneville, 2001); our review of Lothian's ideas, however, is more elaborate as his original research papers are unfortunately difficult to retrieve from public sources nowadays, although they are insightful and therefore deserve to be generally accessible. The details of our proposed generic redesign are discussed in Section 3, leading to a new blueprint for MSR-based seasonal filter selection. Section 4 presents some illustrations, using one particular redesigned rule based upon threshold quartiles and real-time data for three German macroeconomic time series recorded at quarterly, monthly, and daily intervals. Finally, Section 5 concludes.

Table 1: Evolution of MSR-based automatic selection rules for $3 \times k$ seasonal filters

	Selection interval for ... seasonal filter						
Program	3×1	3×3	3×5	3×9	3×15	Stable	Source
	$[0, 1.9)$	$[1.9, 2.7)$	$[2.7, 3.6)$	$[3.6, 4.4)$	$[4.4, 5.0)$	$[5.0, \infty)$	Marris (1960)
X-10	$[0, 1.5)$	$[1.5, 2.5)$	$[2.5, 4.5)$	$[4.5, 6.5)$	$[6.5, 8.5)$	$[8.5, \infty)$	Shiskin et al. (1967)
	$[0, 2.3)$ $[0, 2.1)$	$[2.3, 4.1)$ $[2.1, 3.8)$	$[4.1, 5.2)$ $[3.8, 5.0)$	$[5.2, 6.5)$ $[5.0, 6.9)$		$[6.5, \infty)$ $[6.9, \infty)$	Lothian (1978)
X-11-ARIMA	$[0, 1.5)$	$[1.5, 2.5)$	$[2.5, 7.0)$			$[7.0, \infty)$	Lothian (1978)
X-11-ARIMA/88 X-12-ARIMA X-13ARIMA-SEATS JDemetra+		$[0, 2.5]$	$[3.5, 5.5]$	$[6.5, \infty)$			Dagum, Chhab, and Solomon (1991)

Remarks: **1** The unimplemented rule of Marris (1960) is based upon the MSR variant $\mathbb{E}\left|\Delta\left(si\right)_{i,j}\right| / \mathbb{E}\left|\Delta\,\hat{s}_{i,j}\right|$ in place of (7), where in analogy to (6) a simple 5-term moving average is used as the tentative estimator of the seasonal component. However, the potential benefits of the 7-term moving average are also acknowledged. The extrapolations of the missing values in the seasonal-irregular component near the sample boundaries are also slightly different; in particular, they depend on the length of the $3 \times k$ filter under consideration. **2** The unimplemented rules of Lothian (1978) apply to 10- (top) & 20-year series (bottom). **3** If the MSR happens to fall into one of the indifference intervals where no seasonal filter is selected according to the rule of Dagum et al. (1991), then it is recalculated using a shortened seasonal-irregular component from which the last year has been omitted. If necessary, this recalculation is repeated a maximum of five times in X-11/12/13 and as long as the shortened seasonal-irregular component contains at least 5τ observations in JDemetra+. If still no seasonal filter can be selected, then the 3×5 filter is finally chosen. **4** The implemented selection rules are generally applicable to both monthly and quarterly data.

2 Theoretical background

Let $\{x_{i,j}\}$ denote a time series with τ observations per year, so that $x_{i,j}$ is the observation for the j-th period of the i-th year. Ignoring calendar variation and outliers for the sake of convenience, we assume that this series can be decomposed either additively or multiplicatively into a trend-cyclical component $\{t_{i,j}\}$, a seasonal component $\{s_{i,j}\}$, and an irregular component $\{\varepsilon_{i,j}\}$, that is

$$x_{i,j} = t_{i,j} \circ s_{i,j} \circ \varepsilon_{i,j}, \quad (i,j) \in \{1,\dots,n\} \times \{1,\dots,\tau\}, \tag{1}$$

where $\circ \in \{+,\times\}$ and n is the number of years.

2.1 Essentials of the X-11 method

Iterative linear filtering philosophy The classical X-11 method developed by Shiskin et al. (1967) is a nonparametric approach[1] to the seasonal adjustment of quarterly and monthly time series (so $\tau \in \{4, 12\}$), which has been modelled after the "ratio-to-moving-average" principle (Macaulay, 1931). Its core idea is to run a predetermined sequence of linear filtering operations on the observed time series in order to produce a sequence of preliminary, refined, and final estimates of the three unobservable components appearing on the right-hand side of model (1). These filtering operations are schematised in Figure 1; in essence, they boil down to an iteration of the following 4-step principle:

1. Obtain an estimate of the trend-cyclical component by smoothing the input series with a proper trend-cycle filter; both the input series and the trend-cycle filter depend on the current position in the entire filtering process (see below).

2. Remove the estimated trend-cycle from the observed time series to obtain an estimate of the fused seasonal-irregular component, henceforth also referred to as the SI ratios.

3. Obtain estimates of the seasonal and irregular components by period-wise smoothing the SI ratios either with a $3 \times k$ seasonal moving average or with a stable seasonal filter giving the average SI ratios for each period.

4. Remove the estimated seasonal component from the observed time series to obtain an estimate of the seasonally adjusted data.

This basic principle is passed through twice in the three X-11 iterations labelled B, C, and D; the first two iterations mainly serve to detect and down-weight extreme SI ratios, whereas the last iteration actually aims to provide the final estimates of the trend-cyclical, seasonal and irregular components. In any of these iterations, the input series to the first pass is the observed time series, possibly corrected for prior adjustment factors and/or extreme values, which is then smoothed with a composite $2 \times \tau$ moving average; in the second pass (and for the final trend-cycle estimation in the D iteration), the input series is the most recent estimate of the seasonally adjusted data which is then smoothed with

[1] This does not mean that the X-11 method is distribution-free. In fact, Gaussianity is made use of at some stages in the algorithm, see, for example, Lothian's fundamental model (4) below.

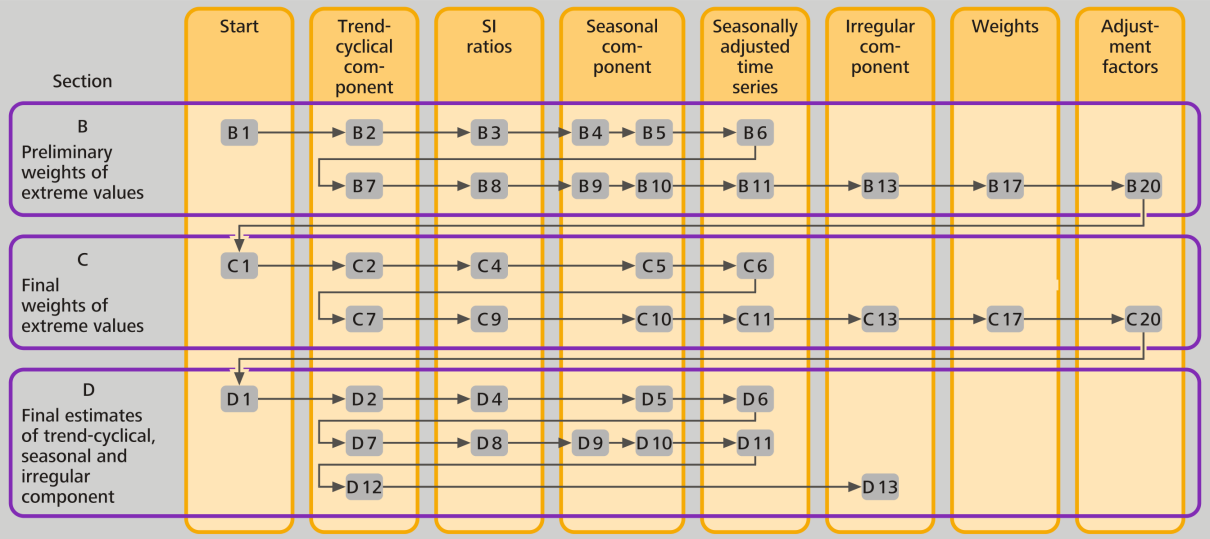

Figure 1: Basic workflow of the X-11 seasonal adjustment method as implemented in JDemetra+.

a Henderson filter. More information about the classical X-11 method and its various output tables can be found in Ladiray and Quenneville (2001).

Seasonal filters and their selection The composite $3 \times k$ seasonal moving averages can be compactly characterised by the z-transform of their coefficients:

$$\mathcal{S}_{3\times k}(z) = \mathcal{S}_3^\star(z)\,\mathcal{S}_k^\star(z) \quad \text{with} \quad \mathcal{S}_m^\star(z) = \frac{1}{m}\,\frac{z^{\tau m}-1}{z^\tau - 1}\,z^{-\tau(m-1)/2}, \tag{2}$$

where z is a complex number and $k \in \{1, 3, 5, 9, 15\}$. When disentangling the SI ratios in Step 3, users can specify one filter for all periods or different filters for different periods, but they may also opt for a data-driven selection between the 3×3, 3×5, and 3×9 filters.[2] The latter approach first requires the calculation of a noise-to-signal (I/S) ratio, also known as the moving seasonality ratio (MSR), and then picks the appropriate filter according to the most recent rule given in Table 1. The global MSR criterion is given by

$$\text{MSR} = \frac{1}{\tau}\sum_{j=1}^{\tau} \text{MSR}_j,$$

where the period-specific local variants are obtained as

$$\text{MSR}_j = \frac{\sum_{i=2}^{n} |\hat{\varepsilon}_{i,j} - \hat{\varepsilon}_{i-1,j}|}{\sum_{i=2}^{n} |\hat{s}_{i,j} - \hat{s}_{i-1,j}|} \tag{3}$$

[2]The automatic selection affects only the seasonal filter finally applied during the second pass in the D iteration; as for the other five passes, the so-called X-11 default filters are applied, that is, the 3×3 filter is applied during the first pass in the B, C, and D iterations, while the 3×5 filter is applied during the second pass in the B and C iterations.

in the additive version of model (1) and as

$$\mathrm{MSR}_j = \frac{\sum_{i=2}^{n} \left| \hat{\varepsilon}_{i,j} \,/ \hat{\varepsilon}_{i-1,j} - 1 \right|}{\sum_{i=2}^{n} \left| \hat{s}_{i,j} \,/ \hat{s}_{i-1,j} - 1 \right|}$$

in the multiplicative version of model (1), where in either case $\{\hat{s}_{i,j}\}$ and $\{\hat{\varepsilon}_{i,j}\}$ are tentative estimators of the seasonal and irregular components detailed in (6) below.

2.2 Lothian's approach to data-driven filter selection

The key idea of the automatic selection of X-11 seasonal filters according to Lothian (1978, 1984) is to pick shorter (longer) $3 \times k$ filters when the degree of moving seasonality is relatively higher (lower) compared with the irregular variation. The derivation of the selection rules comprises three essential steps:

1. Formulate a Gaussian model for the SI ratios and tentative linear estimators for the seasonal and irregular components. Within this framework, both the numerator and denominator of (3) can be written as an unknown expectation of a folded Gaussian random variable.

2. Under the assumption that the moving portion of seasonality evolves linearly near the sample boundaries, either unknown expectation can be rewritten as a nonlinear function of a common unknown nuisance parameter.

3. Under the assumption that two competing concurrent linear estimators for the seasonal component produce the same mean squared error, the unknown nuisance parameter can be expressed as a nonlinear function of the known weights of the two competing filters.

Selected technical details that are important for understanding our proposed redesign are discussed next.

Gaussian model Let $\{(si)_{i,j}\}$ denote the SI ratios[3] and assume that $n > 3$. Sticking to the additive version of model (1) and assuming a Gaussian white noise irregular component, it follows that

$$(si)_{i,j} = s_{i,j} + \varepsilon_{i,j}, \quad \varepsilon_{i,j} \overset{iid}{\sim} \mathcal{N}\left(0, \sigma_\varepsilon^2\right), \tag{4}$$

where $\sigma_\varepsilon^2 < \infty$. The seasonal component is further assumed to evolve as the sum of a stable portion $\{s_j\}$ and a moving portion $\{\tilde{s}_{i,j}\}$ according to

$$s_{i,j} = s_j + \tilde{s}_{i,j}, \tag{5}$$

[3]In practice, X-11 removes the Henderson trend-cycle estimate stored in output table D7 and considers the SI ratios stored in output table D8, possibly corrected with the replacement values for extreme SI ratios stored in output table D9.

where the stable portion satisfies $\sum_{j=1}^{\tau} s_j = 0$ in each year i, while the moving portion fulfills the conditions $\sum_{j=1}^{\tau} \tilde{s}_{i,j} = 0$ in each year i and $\sum_{i=1}^{n} \tilde{s}_{i,j} = 0$ in each period j.

Tentative estimators and their expectations Tentative estimators of the seasonal and irregular components in model (4) are given by

$$\hat{s}_{i,j} = \frac{1}{7} \sum_{k=-3}^{3} (si)_{i+k,j} \quad \text{and} \quad \hat{\varepsilon}_{i,j} = (si)_{i,j} - \hat{s}_{i,j}. \tag{6}$$

The population parameter that the sample-based local MSR (3) attempts to estimate is then given by

$$\text{MSR}_j = \frac{\mathbb{E}\,|\Delta\,\hat{\varepsilon}_{i,j}|}{\mathbb{E}\,|\Delta\,\hat{s}_{i,j}|}, \tag{7}$$

where $\Delta\, z_{i,j} = z_{i,j} - z_{i-1,j}$. The two expectations in (7) can generally be derived without any further assumption (see Lothian, 1978); for the sake of brevity, however, we now already impose an additional linearity assumption on the moving portion of seasonality near the sample boundaries, which would otherwise be required later on anyway. This assumptions reads

$$\tilde{s}_{i,j} = \alpha_j + \beta_j \times i, \quad (\alpha_j, \beta_j) \in \mathbb{R}^2. \tag{8}$$

Plugging (5) into (4), it immediately follows under assumption (8) that

$$\Delta\,\hat{s}_{i,j} \sim \mathcal{N}\left(\beta_j, \frac{2}{49}\,\sigma_\varepsilon^2\right),$$

and hence

$$\mathbb{E}\,|\Delta\,\hat{s}_{i,j}| = \frac{2\sigma_\varepsilon}{7\sqrt{\pi}} \exp\left(-\frac{49\beta_j^2}{4\sigma_\varepsilon^2}\right) + \beta_j\,\Phi_\star\left(\frac{7\beta_j}{\sqrt{2}\sigma_\varepsilon}\right) \tag{9}$$

according to the definition of the folded Gaussian distribution (Leone, Nelson, and Nottingham, 1961), where $\Phi_\star(x) = \Phi(x) - \Phi(-x)$ and $\Phi(\cdot)$ is the cumulative distribution function of the standard Gaussian distribution. Similarly, it can be shown under assumption (8) that

$$\Delta\,\hat{\varepsilon}_{i,j} \sim \mathcal{N}\left(0, \frac{100}{49}\,\sigma_\varepsilon^2\right),$$

and hence

$$\mathbb{E}\,|\Delta\,\hat{\varepsilon}_{i,j}| = \frac{10\sqrt{2}\sigma_\varepsilon}{7\sqrt{\pi}}. \tag{10}$$

Alternative MSR formula Plugging (9) and (10) into (7), it is easily seen that each local MSR population parameter can be rewritten as a nonlinear function in the single nuisance parameter $\delta_j = |\beta_j|/\sigma_\varepsilon$. More precisely, each local MSR criterion is then given by

$$\text{MSR}_j = \frac{10\sqrt{2}}{7\sqrt{\pi}} \times \left[\frac{2}{7\sqrt{\pi}} \exp\left(-\frac{49}{4}\,\delta_j^2\right) + \delta_j\,\Phi_\star\left(\frac{7}{\sqrt{2}}\,\delta_j\right)\right]^{-1}, \tag{11}$$

which is a monotonically decreasing function in $|\beta_j|$ when the irregular variance σ_ε^2 is kept fixed. Therefore, the global maximum of (11) is reached when no moving seasonality is present ($\beta_j = 0$), so that the range of (7) becomes $0 < \text{MSR}_j \leq 5\sqrt{2} \approx 7.071$ within this modelling framework (which may not always hold in practice).

Nuisance parameter Generalising the 7-term moving average estimator in (6), consider the class of linear estimators for the seasonal component given by

$$\hat{\underline{s}}_{i,j} = \sum_{k=1}^{m} w_{i,k}(si)_{k,j}, \tag{12}$$

where the time-dependent weights $\mathbf{w}_i = (w_{i,1}, \ldots, w_{i,m})^\top$ have length $m \leq n$ and satisfy $\sum_{k=1}^{m} w_{i,k} = 1$ for each year $i \in \{1, \ldots, n\}$. The mean squared error (MSE) of such an estimator is given by

$$\text{MSE}\left(\hat{\underline{s}}_{i,j}\right) = \sigma_\varepsilon^2 \left[\delta_j^2 \left(i - \sum_{k=1}^{m} w_{i,k} \times k\right)^2 + \sum_{k=1}^{m} w_{i,k}^2\right], \tag{13}$$

where the two summands in brackets resemble the MSE portions attributable to squared bias and dispersion, respectively. Now, take two concurrent linear estimators of the form (12), say $\{\hat{\underline{s}}_{n,j}^{(\text{A})}\}$ and $\{\hat{\underline{s}}_{n,j}^{(\text{B})}\}$ with known filter weights $\mathbf{w}_n^{(\text{A})} = \left(w_{n,1}^{(\text{A})}, \ldots, w_{n,m_\text{A}}^{(\text{A})}\right)^\top$ and $\mathbf{w}_n^{(\text{B})} = \left(w_{n,1}^{(\text{B})}, \ldots, w_{n,m_\text{B}}^{(\text{B})}\right)^\top$, respectively, and assume that they produce the exact same MSE (13). Under this assumption, the nuisance parameter δ_j in (11) can be rewritten as a nonlinear function of $\mathbf{w}_n^{(\text{A})}$ and $\mathbf{w}_n^{(\text{B})}$ according to

$$\delta_j^2 = \frac{\sum_{k=1}^{m_\text{B}} \left[w_{n,k}^{(\text{B})}\right]^2 - \sum_{k=1}^{m_\text{A}} \left[w_{n,k}^{(\text{A})}\right]^2}{\left[m_\text{A} - \sum_{k=1}^{m_\text{A}} w_{n,k}^{(\text{A})} \times k\right]^2 - \left[m_\text{B} - \sum_{k=1}^{m_\text{B}} w_{n,k}^{(\text{B})} \times k\right]^2}. \tag{14}$$

Finite sample corrections The tentative estimators in (6) cannot be calculated near the sample boundaries in model (4) since the requisite SI ratios for the out-of-sample years $i \in \{-2, -1, 0\} \cup \{n+1, n+2, n+3\}$ are not available. For each period, the missing values at either end are therefore imputed with the average of the three closest SI ratios actually available, that is we set

$$(si)_{k,j} = \frac{1}{3}\sum_{i=1}^{3}(si)_{i,j}, \quad k \in \{-2, -1, 0\},$$

and

$$(si)_{k,j} = \frac{1}{3}\sum_{i=n-2}^{n}(si)_{i,j}, \quad k \in \{n+1, n+2, n+3\},$$

in (6), essentially creating asymmetric variants of the symmetric 7-term moving average. Those artificial extrapolations of course affect the distributions of the transformed tentative estimators near the sample end-points. For example, the boundary distributions of the transformed tentative estimator of the seasonal component are given by

$$\Delta\,\hat{s}_{i,j} \sim \mathcal{N}\left(\tilde{\beta}_j(i), \frac{4}{147}\,\sigma_\varepsilon^2\right),$$

where

$$\tilde{\beta}_j(i) = \begin{cases} \beta_j(i+1)/7, & i \in \{2,3,4\} \\ \beta_j(n-i+3)/7, & i \in \{n-2, n-1, n\} \end{cases},$$

which results in the following boundary corrections to (9):

$$\mathbb{E}\,|\Delta\,\hat{s}_{i,j}| = \frac{2\sqrt{2}\sigma_\varepsilon}{7\sqrt{3\pi}}\,\exp\left(-\frac{147\,\tilde{\beta}_j^2(i)}{8\sigma_\varepsilon^2}\right) + \tilde{\beta}_j(i)\,\Phi_\star\left(\frac{7\sqrt{3}\,\tilde{\beta}_j(i)}{2\sigma_\varepsilon}\right).$$

Similar boundary corrections can be made to (10), which, together with the aforementioned corrections to (9), translate into separate rescaling factors for the expectations of the transformed tentative seasonal and irregular estimators, assuming $\beta_j = 0$ yet again. These rescaling factors are n-dependent constants for $n \in \{4,5,6\}$ and the following functions of n for $n \geq 7$:

$$r_s(n) = \frac{n-1}{n-\left(7-\sqrt{24}\right)} \quad \text{and} \quad r_i(n) = \frac{n-1}{n-\left(7-0.2\,\sqrt{894}\right)}.$$

The three ratios of constants for $n \in \{4,5,6\}$ and the function ratio

$$r(n) = \frac{r_i(n)}{r_s(n)} = \frac{n-\left(7-\sqrt{24}\right)}{n-\left(7-0.2\,\sqrt{894}\right)}, \quad n \geq 7, \tag{15}$$

then constitute a sequence of n-dependent overall rescaling factors for the local MSR (11). It is easily seen that these factors have an asymptotic upper bound since $r(n) < (n-2)/(n-1) < 1$; Figure 2 shows both sequences for $n \in \{4,\ldots,50\}$.

3 Redesigned elements

We seek to redesign the classical MSR-based automatic selection of X-11 seasonal filters with regard to three key elements: the involved $3 \times k$ seasonal filters, the calculation of asymmetric variants, and the derivation of the MSR recalculation intervals. More specifically, we consider multiple approaches of obtaining asymmetric seasonal filters, each of which will produce a Lothian-type selection rule utilising the exact same derivation concept for all $3\times k$ filters. This does not only pay tribute to the advanced options available for the trend-cycle filters in the extended X-11 variant, it—more importantly—also enables a comprehensible computation of the MSR recalculation intervals via intersecting the various Lothian-type selection intervals, thus further stabilising the seasonal filter finally selected by the redesigned MSR-based rules.

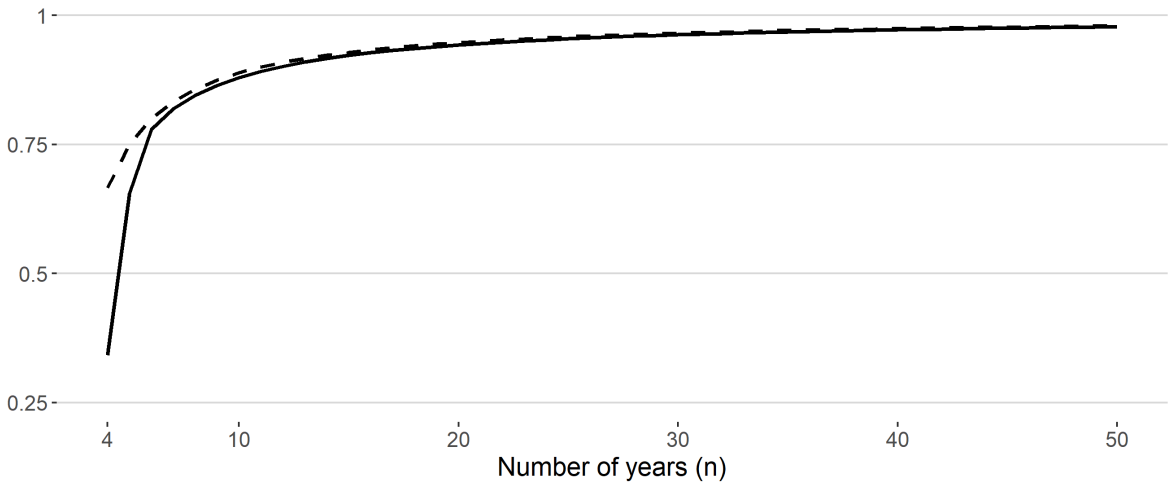

Figure 2: Rescaling factors (15) (*solid line*) and asymptotic upper bound $(n-2)/(n-1)$ (*dashed line*).

3.1 Symmetric $3 \times k$ seasonal filters

Going back to the roots documented in Table 1, we reconsider both the 3×1 and 3×15 moving averages and introduce the 3×23 moving average as a substitute for the stable seasonal filter. The latter is picked since it is the natural extension of the $3 \times k$ filters lineup available in the classical X-11 method.

To simplify notations from Section 2, the weights of each symmetric $3 \times k$ moving average will be denoted by $\mathbf{w}^{(k)}$. It is also noted that this filter has length $m_k = k+2$ and is therefore applicable at the middle positions $i \in \{h_k, \ldots, n - h_k\}$ in model (4), where $h_k = (k+1)/2$ is the filter's half-length, or bandwidth.

3.2 Asymmetric variants

Following these simplified notations, the weights of any asymmetric $3 \times k$ moving average will be denoted by $\mathbf{w}_f^{(k)}$, where $f \in \{0, \ldots, h_k - 1\}$ is the number of future observations available at the current sample boundary after the current time point. Thus, an asymmetric $3 \times k$ moving average has length $m_{k,f} = h_k + 1 + f < m_k$ and is therefore applicable at the boundary position $i = n - f$ in model (4).

We now seek to find principles for determining the weights $\mathbf{w}_f^{(k)}$ directly from $\mathbf{w}^{(k)}$ in a unified manner for all k. Both simple and sophisticated approaches are readily available, but we will focus here on the simpler ones for reasons detailed below. Table 2 provides some examples for $f = 0$ and $k \in \{3, 5, 9\}$, demonstrating that even these principles can generate a broad variety of asymmetric filters covering both balanced and localised variants (see, for example, the concurrent 3×5 and 3×9 filters obtained with the cut-and-redistribute approach (17) for $l = 0$ versus $l = 1$).

Table 2: Weights $\mathbf{w}_0^{(k)}$ for asymmetric variants of selected concurrent $3 \times k$ seasonal filters

	Weight applicable in filter ... to position ...												
	3×3			3×5				3×9					
Variant	$n-2$	$n-1$	n	$n-3$	$n-2$	$n-1$	n	$n-5$	$n-4$	$n-3$	$n-2$	$n-1$	n
CAN	0.1667	0.3333	0.5000	0.1111	0.2222	0.3333	0.3333	0.0667	0.1333	0.2000	0.2000	0.2000	0.2000
CAR-0-CON	0.2222	0.3333	0.4444	0.1667	0.2333	0.3000	0.3000	0.1111	0.1481	0.1852	0.1852	0.1852	0.1852
CAR-0-LIN	0.1667	0.3333	0.5000	0.1067	0.2133	0.3200	0.3600	0.0582	0.1164	0.1746	0.1958	0.2169	0.2381
CAR-0-QDE	0.1190	0.3492	0.5317	0.0533	0.2133	0.3467	0.3867	0.0094	0.1001	0.1811	0.2153	0.2397	0.2544
CAR-0-QIN	0.1349	0.3175	0.5476	0.0800	0.1857	0.3200	0.4133	0.0419	0.0936	0.1551	0.1893	0.2332	0.2869
CAR-1-CON	0.1111	0.2222	0.6667	0.0667	0.1333	0.2000	0.6000	0.0370	0.0741	0.1111	0.1111	0.1111	0.5556
CAR-2-CON	0.1111	0.3889	0.5000	0.0667	0.1333	0.4000	0.4000	0.0370	0.0741	0.1111	0.1111	0.3333	0.3333
CAR-2-LIN	0.1111	0.3333	0.5556	0.0667	0.1333	0.3333	0.4667	0.0370	0.0741	0.1111	0.1111	0.2593	0.4074
CAR-2-QIN	0.1111	0.2889	0.6000	0.0667	0.1333	0.2800	0.5200	0.0370	0.0741	0.1111	0.1111	0.2000	0.4667
LPRF	0.1788	0.3333	0.4879	0.0758	0.2030	0.3303	0.3909	−0.0580	0.0467	0.1514	0.2190	0.2866	0.3543
MUS	0.2123	0.3333	0.4544	0.1433	0.2255	0.3078	0.3234	0.0508	0.1120	0.1731	0.1972	0.2214	0.2455
X-11	0.1852	0.4074	0.4074	0.1500	0.2833	0.2833	0.2833	0.0507	0.1121	0.1735	0.1969	0.2212	0.2456

Remark: The X-11 variant replicates the legacy weights (in decimal form) of the concurrent $3 \times k$ seasonal filters implemented in X-11-ARIMA according to Ladiray and Quenneville (2001).

Cut-and-normalise (CAN) approach Following Gasser and Müller (1979), the obsolete weights $w_j^{(k)}$ with $j \in \{m_{k,f}+1, \ldots, m_k\}$ are dropped from the symmetric filter $\mathbf{w}^{(k)}$, and the remaining weights are simply divided by their sum. More formally, the elements of any asymmetric filter $\mathbf{w}_f^{(k)}$ are thus obtained as

$$w_{f,j}^{(k)} = w_j^{(k)} \left(\sum_{i=1}^{m_{k,f}} w_i^{(k)} \right)^{-1}, \tag{16}$$

where $j \in \{1, \ldots, m_{k,f}\}$.

Cut-and-redistribute (CAR) approach According to Kendall and Ord (1990), the summed obsolete weights dropped from $\mathbf{w}^{(k)}$ are uniformly reallocated to the last l entries of $\mathbf{w}_f^{(k)}$, enforcing $l \in \{1, \ldots, m_{k,f}\}$. The modified weights of the asymmetric filter are thus given by

$$w_{f,j}^{(k)} = w_j^{(k)} + \frac{1}{l} \sum_{i=m_{k,f}+1}^{m_k} w_i^{(k)}, \tag{17}$$

where $j \in \{m_{k,f} - l + 1, \ldots, m_{k,f}\}$, while the other entries of $\mathbf{w}_f^{(k)}$ equal those of $\mathbf{w}^{(k)}$.

The rationale of (17) is to impute in classical X-11 fashion each missing future observation that prevents the application of the companion symmetric filter with a simple l-term average of the l most recent observations actually available, which is why the CAR version (17) will subsequently be referred to as the CAR-l-CON approach (with CON indicating constant imputation weights). When $l > 1$, this rationale motivates a generalised CAR approach in which the constant imputation weights l^{-1} in (17) are replaced with a sequence $\{c_j\}$ of position-dependent imputation weights satisfying $\sum_j c_j = 1$. Those weights should typically increase with j to put more weight on more recent observations in the imputation. Changing the index set from $\{m_{k,f} - l + 1, \ldots, m_{k,f}\}$ to $\{1, \ldots, l\}$ for convenience, a sequence of linearly increasing weights $\{c_j\}$ is obtained from setting

$$c_j = \frac{2j}{l(l+1)}, \quad j \in \{1, \ldots, l\}. \tag{18}$$

The case of using the imputation weights (18) in (17) will be called the CAR-l-LIN approach. A sequence of increasing quadratic weights, meaning that the rate of change between adjacent quadratic weights is increasing, is obtained from setting

$$c_j = \frac{6j^2}{l(l+1)(2l+1)}, \quad j \in \{1, \ldots, l\}. \tag{19}$$

The case of using the imputation weights (19) in (17) will be named the CAR-l-QIN approach. Those weights can also be used to construct recursively a sequence of corresponding decreasing quadratic weights, meaning that the rate of change between adjacent quadratic weights is decreasing in the reversed order of the QIN rates; this case will be referred to as the CAR-l-QDE approach, noting, however, that the QIN and QDE weights are the same for $l = 2$.

In what follows, we will be interested in working with only those CAR approaches

that are equally applicable to all $3 \times k$ seasonal moving averages. The short length of the concurrent 3×1 filter ($m_{1,0} = 2$), however, restricts the reallocation parameter to $l \in \{1, 2\}$, which in turn yields strongly localised concurrent $3 \times k$ filters for higher values of k, regardless of which imputation weights are used in (17) (see, for instance, the CAR-1-CON weights for the concurrent 3×9 filter reported in Table 2). For that reason, we will additionally consider the case $l = m_{k,f}$, or $l = 0$ for short, in which the summed dropped weights from $\mathbf{w}^{(k)}$ are uniformly reallocated across all entries of the asymmetric filter.

Musgrave (MUS) surrogates The concept derived in Musgrave (1964a,b) has been predicated on the idea of minimising the mean squared revision error between preliminary and final seasonal estimates obtained from applying asymmetric and symmetric linear filters, respectively. Although this minimisation problem has originally been solved numerically, the optimal weights can be written explicitly as a function of the entries in $\mathbf{w}^{(k)}$. For example, Findley et al. (1998) give the elements of $\mathbf{w}_f^{(k)}$ as

$$w_{f,j}^{(k)} = w_j^{(k)} + \frac{1}{m_{k,f}} \sum_{i=m_{k,f}+1}^{m_k} w_i^{(k)} + \frac{\left(j - \frac{m_{k,f}+1}{2}\right) \delta^2}{1 + \frac{m_{k,f}\left(m_{k,f}^2-1\right)}{12} \delta^2} \sum_{i=m_{k,f}+1}^{m_k} \left(i - \frac{m_{k,f}+1}{2}\right) w_i^{(k)}, \quad (20)$$

where $j \in \{1, \ldots, m_{k,f}\}$, noting that Musgrave's ideas can also be applied to find asymmetric Henderson trend-cycle filters (see the related work of Doherty, 2001; Kenny and Durbin, 1982; Laniel, 1986; Quenneville, Ladiray, and Lefrançois, 2003). Reinterpreting model (4)–(5) as a trend-cycle extraction problem under linearity assumption (8) for the j-th period, the nuisance parameter in (20) can be written as $\delta = |\beta|/\sigma_\varepsilon$; unsurprisingly, it can also be linked to the I/C noise-to-signal ratio (employed for the automatic selection of Henderson trend-cycle filters in the classical X-11 method) via $\delta^2 = 4/(\pi \times I/C^2)$. To identify the nuisance parameter providing the Musgrave weights (20) closest in the mean squared error sense to the legacy X-11 weights of the concurrent 3×9 filter (see Table 2), we perform a grid search over $I/C \in [9, 10]$ in incremental steps of 0.01 and find the minimiser $I/C = 9.85$ and, hence, $\delta^2 \approx 0.0131$, which will be used in our subsequent calculations.

Selected other approaches We also experimented with some other and more sophisticated derivation principles, but ran into several problems with those and therefore decided to refrain from considering them any further. For example, the symmetric $3 \times k$ seasonal moving averages can be embedded in the kernel-based local polynomial regression filters framework (Proietti and Luati, 2008) via specifying a uniform ($k = 1$) or trapezoidal ($k > 1$) kernel under the requirement of preserving linear movements in the middle of the sample span and constants near the sample boundaries. However, the resulting asymmetric variants tend to be strongly localised and may even contain negative entries, especially for the longer concurrent filters (see the LPRF weights reported in Table 2). These features occasionally translate into negative and hence erroneous transition values for the squared nuisance parameter δ_j^2 in (14), presumably because Lothian's framework is implicitly built on non-negative weights for either competing concurrent filter.

Reproducing kernel Hilbert spaces (Dagum and Bianconcini, 2008, 2013, 2015) are also

capable of replicating the $3 \times k$ seasonal moving averages, using uniform and trapezoidal kernels once again; as for the latter, however, the associated probability density defining the inner product of the Hilbert space is not unique, leaving some element of randomness.

Generalising Musgrave's mean squared revision error criterion for trend-cycle filters, recent advances that seek to minimise the asymmetric filter's phase shift by taking into account its fidelity, smoothness, and timeliness properties (Grun-Rehomme, Guggemos, and Ladiray, 2018) suffer from the same issue; even an operational procedure based upon prefixed weights for those three properties seems infeasible in practice as the optimal weight allocation remains unclear. We are also unaware of an (operational) extension to the case of $3 \times k$ seasonal moving averages.

3.3 MSR recalculation intervals

A minor blemish of the classical MSR-based selection rule currently implemented in many contemporary seasonal adjustment programs (see Table 1) is the somewhat mysterious emergence of its MSR recalculation intervals; while the purpose of introducing these intervals has clearly been the avoidance of too frequent revisions to the automatic filter choice (see e.g. Dagum et al., 1991), it nowadays seems rather difficult to retrieve any information about the origin of the implemented thresholds controlling the potential MSR recalculations from publicly accessible documents. To lighten this opaqueness, our proposed redesign is founded on multiple Lothian-type selection rules where each of them utilises one specific approach to deriving the requisite concurrent variants of all involved $3 \times k$ seasonal moving averages; the final selection interval for each seasonal filter is then obtained through averaging in an appropriate manner the corresponding approach-specific selection intervals, which by construction will also create the final MSR recalculation intervals in a natural way. This intuitively requires the approach-specific selection intervals to display a moderate amount of variation but not too much dissimilarity.

3.4 Final selection rules

Overall, the blueprint for our proposed redesign of the classical MSR-based selection of X-11 seasonal filters consists of the following steps:

1. Retrieve the weights $\mathbf{w}^{(k)}$ of the symmetric $3 \times k$ seasonal moving averages from the z-transform (2), where $k \in \{1, 3, 5, 9, 15, 23\}$.

2. For each symmetric $3 \times k$ seasonal moving average, derive the weights $\mathbf{w}_0^{(k)}$ of the concurrent filter according to each of the following principles:

   - the CAN approach (16),
   - the CAR-l-CON approach (17) with $l \in \{0, 1, 2\}$,
   - the CAR-0-LIN approach (18),
   - the CAR-0-QIN approach (19) and the associated CAR-0-QDE approach,
   - the MUS approach (20) with $\delta^2 = 0.0131$.

Table 3: Examples of deriving redesigned MSR-based automatic selection rules for $3 \times k$ seasonal moving averages

	Lothian-type selection intervals					
Variant	3×1	3×3	3×5	3×9	3×15	3×23
CAN	$[0, 1.4)$	$[1.4, 2.8)$	$[2.8, 4.6)$	$[4.6, 5.9)$	$[5.9, 6.3)$	$[6.3, \infty)$
CAR-0-CON	$[0, 1.6)$	$[1.6, 3.2)$	$[3.2, 4.9)$	$[4.9, 6.0)$	$[6.0, 6.3)$	$[6.3, \infty)$
CAR-0-LIN	$[0, 1.5)$	$[1.5, 2.7)$	$[2.7, 4.4)$	$[4.4, 5.7)$	$[5.7, 6.2)$	$[6.2, \infty)$
CAR-0-QDE	$[0, 1.4)$	$[1.4, 2.4)$	$[2.4, 4.0)$	$[4.0, 5.5)$	$[5.5, 6.1)$	$[6.1, \infty)$
CAR-0-QIN	$[0, 1.4)$	$[1.4, 2.3)$	$[2.3, 4.0)$	$[4.0, 5.5)$	$[5.5, 6.1)$	$[6.1, \infty)$
CAR-1-CON	$[0, 1.4)$	$[1.4, 1.8)$	$[1.8, 3.5)$	$[3.5, 5.3)$	$[5.3, 6.1)$	$[6.1, \infty)$
CAR-2-CON	$[0, 1.2)$	$[1.2, 2.4)$	$[2.4, 3.5)$	$[3.5, 5.3)$	$[5.3, 6.1)$	$[6.1, \infty)$
MUS	$[0, 1.6)$	$[1.6, 3.1)$	$[3.1, 4.4)$	$[4.4, 5.2)$	$[5.2, 6.1)$	$[6.1, \infty)$
	Final selection intervals					
Intersection	3×1	3×3	3×5	3×9	3×15	3×23
[MAX, MIN]	$[0, 1.2]$	$[1.6, 1.8]$	$[3.2, 3.5]$	$[4.9, 5.2]$	$[6.0, 6.1]$	$[6.3, \infty)$
[QR3, QR1]	$[0, 1.4]$	$[1.5, 2.4]$	$[2.9, 3.9]$	$[4.4, 5.3]$	$[5.8, 6.1]$	$[6.2, \infty)$

Remarks: **1** Variant-specific selection intervals are based upon the finite sample corrections (15) with $n = 13$. **2** QR1 and QR3 denote the first and third median-unbiased sample quartiles as given in Definition 8 of Hyndman and Fan (1996), that is, they equal $\hat{Q}_8(0.25)$ and $\hat{Q}_8(0.75)$, respectively.

3. For each pair of neighbouring concurrent $3 \times k$ seasonal moving averages and each derivation principle, calculate the transition values of the squared nuisance parameter δ_j^2 under the equal-MSE condition via (14).

4. Obtain the corresponding MSR thresholds for selecting the $3 \times k$ seasonal moving averages under the given derivation principle from substituting these transition values into (11), using Lothian's finite sample correction (15) with $n = m_{23,0} = 13$.

5. For each $3 \times k$ seasonal moving average, calculate an appropriate intersection of its Lothian-type selection intervals, taking all derivation principles into account.

Table 3 provides two examples of such redesigned selection rules. The top panel shows the Lothian-type selection intervals obtained for the six candidate $3 \times k$ seasonal moving averages under the aforementioned eight principles for deriving their asymmetric variants. It turns out that the selection intervals for the 3×1 and 3×23 filters are quite similar across the derivation principles, but those for the other four filters show indeed a decent amount of variation.

The bottom panel shows how the final selection intervals can be obtained through intersecting the eight Lothian-type selection rules. According to the first intersection approach, the left and right endpoints of any final selection interval are given by the maximum lower and minimum upper thresholds defining the corresponding Lothian-type selection interval. For example, the (ordered) eight thresholds for transitioning from the 3×5 to the 3×9 filter are 3.5, 3.5, 4.0, 4.0, 4.4, 4.4, 4.6, and 4.9, so their minimum (3.5)

becomes the right endpoint of the final selection interval for the 3×5 filter, whereas their maximum (4.9) becomes the left endpoint of the final selection interval for the 3×9 filter. This [MAX, MIN] intersection, however, produces in general very narrow final selection intervals and hence very wide MSR recalculation intervals, which is clearly undesired.

The second intersection approach follows the same logic, but the maximum and minimum thresholds are now replaced with the third quartile of the lower thresholds and the first quartile of the upper thresholds, respectively. In the previous transition scenario, the first quartile of the eight thresholds is 3.9, which becomes the right endpoint of the final selection interval for the 3×5 filter, whereas the third quartile (4.4) becomes the left endpoint of the final selection interval for the 3×9 filter. This [QR3, QR1] intersection much better balances the widths of the final selection and MSR recalculation intervals, although the final selection interval for the 3×15 filter is still relatively narrow. The symmetric widths of the MSR recalculation intervals are another opportune feature. They amount to 0.1 for the two outer transitions (3×1 to 3×3 and 3×15 to 3×23) and to 0.5 for the three inner transitions.

4 Illustration

We now consider three German macroeconomic time series to illustrate the [QR3, QR1] rule, including a comparison with the classical MSR-based selection rule whenever applicable (recall Table 1). The three series are quarterly gross domestic product (GDP), monthly industrial production, and daily realised electricity consumption; each of them is available as a real-time data set storing the different versions of the unadjusted time series as they have been released by the respective data provider at various consecutive time points in the past. Accordingly, the different release versions of a time series share the same starting date, but have an increasing sample span over time, resulting in a triangular shape of the entire real-time data set. In what follows, release versions will be called vintages and be indexed by release period. For example, the Q2 2005 vintage of GDP refers to quarterly GDP as released in the second quarter of 2005, covering the Q1 1991 through Q1 2005 sample span. Table 4 summarises basic information about the data vintages and providers of our three examples.

The main purpose of this empirical exercise is to evaluate the filter selection made by the [QR3, QR1] rule. Does this redesigned rule tend to pick the same seasonal filters as the classical MSR-based rule, or are there significant discrepancies? And how do both filter selection strategies compare in terms of seasonal adjustment adequacy? To answer these questions, we are now going to seasonally adjust each vintage in each real-time data set under a concurrent revision policy. The examples on quarterly GDP and monthly industrial production will then enable us to compare the seasonal filters selected by both rules across the vintages and to check the corresponding real-time seasonally adjusted figures for systematic over- and under-adjustment issues, measured through sample first-order seasonal autocorrelations. The example on daily electricity consumption will showcase that additional intermediate steps in the MSR calculation may be required when the raw data display more complex forms of seasonality.

The general setting for the seasonal adjustment specification in each example is as follows:

Table 4: Basic information about real-time unadjusted data

	GDP	Industrial production	Electricity consumption
Periodicity	Quarterly	Monthly	Daily
Unit	Chain-linked index	Laspeyres index	Gigawatt hours (GWh)
First vintage	Q2 2005	Nov 2005	W01 2022
Span	Q1 1991–Q1 2005	Jan 1991–Sep 2005	1 Jan 2015–30 Dec 2021
Sample size	57	177	2,556
Last vintage	Q1 2025	Feb 2025	W29 2025
Span	Q1 1991–Q4 2024	Jan 1991–Dec 2024	1 Jan 2015–13 Jul 2025
Sample size	136	408	3,847
# vintages	80	232	185
Data provider	Federal Statistical Office	Federal Statistical Office	Federal Network Agency

Remarks: **1** The chain-linked GDP data are calculated using previous-year prices and the annual overlap method. The GDP vintages do not contain flash estimates, which are available as of the Q3 2020 release. **2** The base year of the industrial production index is typically updated every five years; the real-time vintages cover the base years 2000, 2005, 2010, 2015, and 2021. **3** Realised electricity consumption covers electricity supplied to the network for the general supply, excluding electricity supplied to the railroad network, electricity supplied to internal industrial and closed distribution networks, and electricity consumed by the producers. The vintages for this daily series have been updated once a week, see Webel (2025) for details.

- The multiplicative form of the unobserved components (UC) decomposition (1) is used.
- Prior to UC estimation, each series is corrected for deterministic working-day and/or holiday-related variations using a linear time series regression with ARIMA errors. The resulting calendar factor estimates are removed permanently, that is, they will not be contained in the final seasonally adjusted estimates.
- The ARIMA model for the regression disturbances of the aforementioned pretreatment model is identified automatically for the monthly and quarterly series, using the JDemetra+ implementation of the procedure described in Gómez and Maravall (2001). Forecasts of the unadjusted data are generated from the entire model over a one-year horizon; no backcasts are generated. In contrast, the linearisation of the daily series utilises the so-called extended Airline model without forecast extensions (see e.g. Webel and Smyk, 2024, for details).
- Either data pretreatment framework also involves the automatic detection of additive outliers (AO) and level shifts (LS) based upon t-tests with critical values that have been adapted to sample size (see Table 7.22 in U.S. Census Bureau, 2025, for examples). In contrast to calendar factor estimates, the estimated outlier effects are removed temporarily, that is, after UC estimation, the estimated AO and LS effects will be assigned to the final irregular and trend-cycle estimates, respectively.

- Concurrent seasonal adjustments are run across the data vintages; this means the entire pretreatment model is re-identified and re-estimated from each vintage and, using the freshly precleaned data, updated estimates for each UC in model (1) are obtained as well.

The specific settings regarding the X-11 filtering operations in each example will be detailed below in the respective subsections.

4.1 Quarterly GDP

The results for the quarterly GDP real-time data are summarised in Figure 3. Panel (a) shows the unadjusted series released in the Q2 2005 vintage; Panel (b) displays the corresponding autoregressive (AR) model-based estimate of order $p = 30$ for the spectral density of the differenced logged series. Both the time series plot and the spectral plot clearly indicate the presence of seasonal dynamics in the quarterly GDP data. In particular, the estimated spectral density exhibits visible peaks at the two seasonal frequencies, which are marked as grey verticals in Panel (b).

The X-11 seasonal adjustments utilise the 5-term Henderson trend-cycle filter and an automatic detection and replacement of extreme SI ratios based upon quarter-specific standard deviations and default σ-limits. Panel (c) shows the automatically selected seasonal moving averages when following the classical versus redesigned MSR-based rules. According to the classical rule, the 3×3 filter is picked for the GDP series released between the Q1 2012 through Q2 2019 vintages, whereas the 3×5 filter is picked for the other data releases. The redesigned rule tends to select the same seasonal filters in general, but occasionally picks slightly longer ones for GDP series released in earlier vintages, leading to a total of eleven filter revisions over the 80 data vintages (compared to just two filter revisions suggested by the classical rule). To check for systematic seasonal adjustment inadequacy, Panel (d) compares the kernel density estimates for the sample first-order seasonal autocorrelation, $\hat{\rho}_4$, in the differenced logged seasonally adjusted real-time data. In general, an estimated density having most of its mass centred around zero and within the grey range that marks the approximate large-lag standard error of $\hat{\rho}_4$ signals seasonal adjustment adequacy; in contrast, an estimated density having most of its mass distributed over large positive (negative) values of $\hat{\rho}_4$ indicates a systematic under-adjustment (over-adjustment) issue. It is seen that the seasonally adjusted GDP series obtained from applying the classical and redesigned seasonal filter selections equally tend to suffer from a minor under-adjustment issue as the corresponding unimodal density estimates peak near 0.05, which is well within the range indicating adequacy (the approximate large-lag standard error of $\hat{\rho}_4$ is 0.125), but also display non-ignorable mass over the range $\hat{\rho}_4 \in [0.1, 0.3]$. Hard-coding the seasonal filter picked by either automatic selection rule somewhat mitigates this issue, but generally does not result in a significant improvement.

Overall, the classical and redesigned rules often select the same seasonal filter and hence lead to very similar seasonally adjusted GDP series. An example is shown in Panel (a), using the unadjusted data from the Q2 2005 vintage and the 3×9 seasonal filter picked by the redesigned rule; the corresponding AR(30) estimate for the spectral density of the differenced logged seasonally adjusted GDP series is plotted in Panel (b), indicating no severe symptoms of residual seasonality in this particular case. In general, however, such

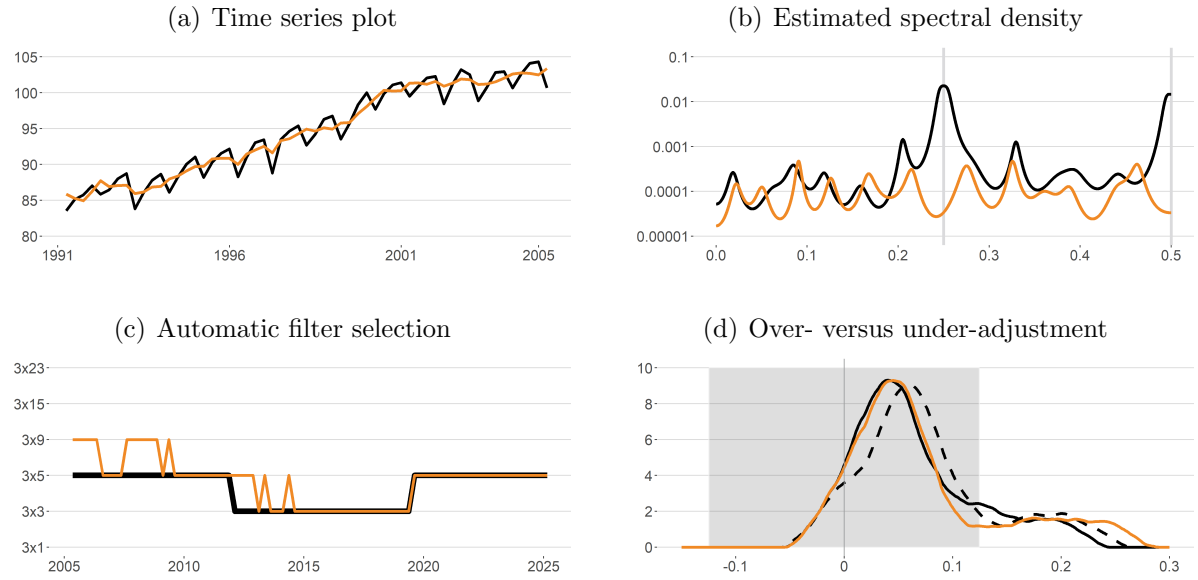

Figure 3: Results for quarterly GDP. Panel (a) shows the unadjusted data (*black line*, Q2 2005 vintage) and the seasonally adjusted data obtained via the [QR3, QR1] rule (*orange line*). Panel (b) shows the AR(30) estimate for the spectral density of the two differenced logged series (in decibels, abscissa in units of cycles per time). Grey verticals mark the two seasonal frequencies. Panel (c) shows the $3 \times k$ seasonal filters picked by the classical (*black line*) and [QR3, QR1] (*orange line*) rules over the Q2 2005 through Q1 2025 vintages. Panel (d) shows Epanechnikov kernel density estimates for $\hat{\rho}_4$ obtained from the differenced logged seasonally adjusted data with (*solid black line*) and without (*dashed black line*) hard-coding the seasonal filter picked by the classical rule. The shaded area marks ± 1 Bartlett approximation to the large-lag standard error of $\hat{\rho}_4$.

an under-adjustment issue could be tackled by shortening the seasonal filter currently in use.

4.2 Monthly industrial production

Figure 4 summarises the results for the monthly German industrial production index. Panel (a) shows the unadjusted data released in the November 2005 vintage, and Panel (b) provides the corresponding AR(30) estimate of the spectral density calculated from the differenced logged series. Seasonality can be easily spotted in the time series plot and is also signalled by the spectral peaks located at the six seasonal frequencies. The estimated spectral density also has an additional local peak at the (radian) frequency 0.348, which is one of the two frequencies associated with trading-day variation in monthly time series (Cleveland and Devlin, 1980).

The 17-term Henderson trend-cycle filter and an automatic detection and down-weighting of extreme SI ratios, using month-specific standard deviations and default σ-limits, are specified for the X-11 seasonal adjustments. Panel (c) shows that both the classical and redesigned rule automatically pick the 3×5 seasonal filter over all data vintages except the final twelve ones, over which the slightly longer 3×9 filter is selected by the redesigned rule. Unsurprisingly, the real-time seasonally adjusted estimates obtained

Figure 4: Results for monthly industrial production. Panel (a) shows the unadjusted data (*black line*, 2000 = 100, November 2005 vintage) and the seasonally adjusted data obtained via the [QR3, QR1] rule (*orange line*). Panel (b) shows the AR(30) estimate for the spectral density of the two differenced logged series (in decibels, abscissa in units of cycles per time). Grey verticals mark the six seasonal frequencies. Panel (c) shows the $3 \times k$ seasonal filters picked by the classical (*black line*) and [QR3, QR1] (*orange line*) rules over the November 2005 through February 2025 vintages. Panel (d) shows Epanechnikov kernel density estimates for $\hat{\rho}_{12}$ obtained from the differenced logged seasonally adjusted data with (*solid black line*) and without (*dashed black line*) hard-coding the seasonal filter picked by the classical rule. The shaded area marks ± 1 Bartlett approximation to the large-lag standard error of $\hat{\rho}_{12}$.

with the two selection rules show little to almost no differences; in particular, the two corresponding kernel density estimates for the sample first-order seasonal autocorrelation, $\hat{\rho}_{12}$, in the differenced logged seasonally adjusted industrial production indices, which are shown in Panel (d), display a bimodal shape and are heavily concentrated in the range $\hat{\rho}_{12} \in [-0.1, 0]$ with 0.065 being the approximate large-lag standard error of $\hat{\rho}_{12}$, signalling mild symptoms of systematic over-adjustment, which is often commonly accepted as the more forgivable type of seasonal adjustment inadequacy and could be tackled by specifying a longer seasonal filter. An example of the seasonally adjusted industrial production index in plotted in Panel (a), using the unadjusted data published in the November 2005 vintage; the corresponding AR(30) estimate for the spectral density of the differenced logged seasonally adjusted series is shown in Panel (b).

4.3 Daily electricity consumption

The results for the daily realised electricity consumption series are summarised in Figure 5. Using the W01 2022 release vintage, the time series plot in Panel (a) highlights the presence of a characteristic U-shaped day-of-the-year (DOY) pattern in the unadjusted data. The COVID-19 pandemic outbreak in March 2020 and its recovery phase are also clearly visible in both unadjusted and seasonally adjusted electricity consumption. The

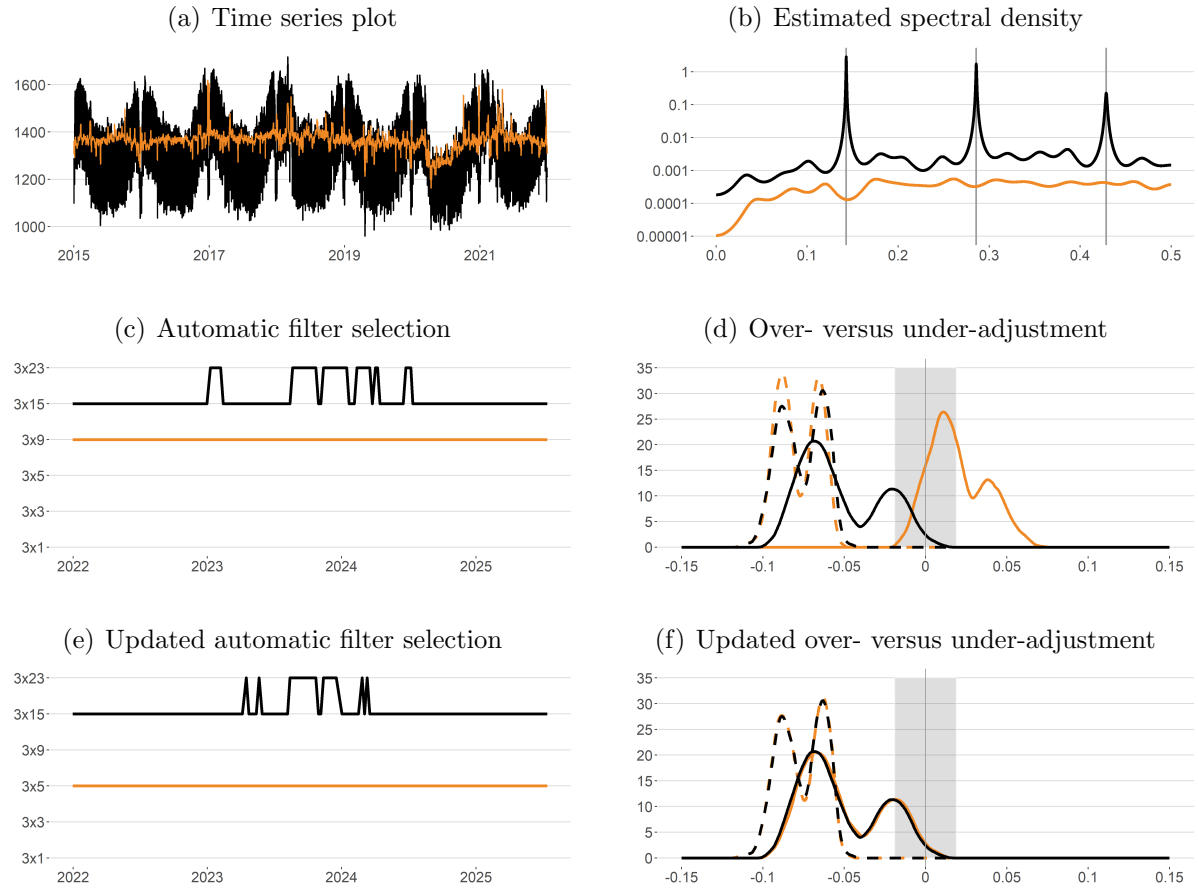

Figure 5: Results for daily electricity consumption. Panel (a) shows the unadjusted data in GWh units (*black line*, W01 2022 vintage) and the seasonally adjusted data obtained via the [QR3, QR1] rule when extracting the day-of-the-week (DOW) pattern (*orange line*). Panel (b) shows the AR(30) estimate for the spectral density of the two differenced logged series (in decibels, abscissa in units of cycles per time). Grey verticals mark the three day-of-the-week frequencies. Panels (c) and (e) show the $3 \times k$ seasonal filters picked by the [QR3, QR1] rule when extracting the DOW (*orange line*) and day-of-the-year (*black line*, DOY) patterns over the W01 2022 through W29 2025 vintages; the difference between the two scenarios is whether *(e)* or not *(c)* the DOY dynamics in (21) are recognised when selecting the seasonal filter for DOW extraction. Panels (d) and (f) show corresponding Epanechnikov kernel density estimates for $\hat{\rho}_7$ (*solid lines*) and $\hat{\rho}_{365}$ (*dashed lines*) obtained from the differenced logged seasonally adjusted data using the $3 \times k$ seasonal filters picked by the [QR3, QR1] rule (*orange lines*) and the X-11 default filters (*black lines*) for DOW extraction. The shaded area marks ± 1 Bartlett approximation to the large-lag standard error of $\hat{\rho}_7$.

estimated spectral density shown in Panel (b) reveals the presence of further and even stronger seasonal movements that will be referred to as the day-of-the-week (DOW) pattern, owing to the associated spectral peaks being located at the three frequencies 1/7, 2/7, and 3/7 (expressed in units of cycles per time).

The sequential extraction of the DOW and DOY patterns with the extended X-11 method utilises trend-cycle estimation with kernel-based local polynomial regression filters, see Proietti and Luati (2008) for details and Webel and Smyk (2024) for an overview. More specifically, for either extraction, a kernel-based Henderson filter is configured to preserve cubic polynomials in the middle of the sample span and "linear-constant" polynomials near the sample boundaries, which produces the requisite asymmetric trend-cycle filters within the given modelling framework; the length of the symmetric filter is set to the smallest odd integer larger than the periodicity of the seasonal pattern to be extracted. In addition, the preliminary and final weights of extreme SI ratios are calculated with default σ-limits in the B and C iterations (recall Figure 1). Note, however, that currently no replacement values for extreme SI ratios are produced in output table D9; for that reason, the deviations between the filtered linearised electricity consumption series stored in output table D1 and the estimated trend-cycle stored in output table D7 are taken as a proxy for the extreme-value corrected SI ratios that enter the MSR calculation.[4]

When selecting the seasonal moving averages with the [QR3, QR1] rule as before, Panel (c) shows that the 3×9 filter is picked across all vintages for extracting the DOW pattern, whereas with regard to DOY extraction the longer 3×15 filter is chosen for most vintages with the exception of some 2023 and 2024 vintages where the even longer 3×23 filter is picked.[5] In practice, however, the rather short sample spans of the considered real-time vintages prevent the application of those long moving averages for DOY estimation, so that we resort to the X-11 default filters for this purpose. According to Panel (d), the resulting kernel density estimates for the sample lag-7 and lag-365 autocorrelations in the differenced logged seasonally adjusted data are bimodal and mainly concentrated in the ranges $\hat{\rho}_7 \in [-0.025, 0.075]$ and $\hat{\rho}_{365} \in [-0.1, -0.05]$, respectively, indicating in particular a systematic and non-ignorable amount of residual DOW seasonality. Given that the approximate large-lag standard error of $\hat{\rho}_7$ is 0.019, the 3×9 filter is apparently too long for adequately capturing the DOW swings in unadjusted electricity consumption, confirming the conclusion reached in a recent large-scale revision analysis of this as well as other daily and weekly time series (Webel, 2025).

This undesirable outcome can, in fact, be attributed to the coexistence of two seasonal patterns in the unadjusted data, which has been neglected so far despite the visual evidence shown in Panels (a) and (b) of Figure 5. In essence, for daily electricity consumption, the seasonal component in UC model (1) factorises according to

$$s_{i,j} = s_{i,j}^{(7)} \circ s_{i,j}^{(365)}, \tag{21}$$

[4]Since the SI ratios stored in output table D8 are in general calculated as the difference/ratio between the linearised data stored in output table B1 and the trend-cycle estimate stored in output table D7, our proxy is essentially obtained from replacing in this calculation the linearised data with a filtered version thereof that acknowledges the final weights for extreme values stored in output table C17.

[5]For the sake of convenience, DOW-adjusted detrended observations recorded on 29 February are disregarded temporarily to enforce an integer seasonal periodicity, namely $\tau = 365$, when handling the DOY pattern.

where $\{s_{i,j}^{(7)}\}$ and $\{s_{i,j}^{(365)}\}$ represent the DOW and DOY patterns, respectively, invalidating Lothian's (implicit) key assumption of a single seasonal pattern being present in the raw data. Ignoring the DOY dynamics when automatically selecting the seasonal moving average designated for DOW estimation will thus inevitably set free the following painful chain reaction:

- When ignored, the DOY pattern in (21) erroneously intermingles with the "true" irregular component from model (1) with $\tau = 7$, constituting a fused DOY-irregular DOW remainder component in Lothian's fundamental Gaussian model (4).

- Under the assumption of orthogonal components, which Lothian's framework is based upon, the variance of the fused DOY-irregular remainder component must be higher than that of the "true" irregular component. Hence, the volatility of Lothian's tentative estimate for the irregular component in (6) is spuriously inflated.

- Since this inflated variance estimate appears in the numerator of the MSR criterion, such as the additive version (3), an upward bias is introduced to the calculated MSR value.

- Since larger values of the MSR criterion tend to go along with the automatic selection of longer seasonal filters, the $3 \times k$ seasonal moving average automatically selected for DOW estimation without proper recognition of the DOY movements in (21) is most probably too long.

To remedy this deficiency, we pragmatically add an intermediate step after obtaining the tentative estimate of the DOW pattern according to (6); during this step, tentative estimates of the DOY pattern and the irregular component are calculated from the fused DOW remainder component, using a reshaped version of model (4) with $\tau = 365$ and the default estimators in (6). The tentative estimates of the DOW pattern and the disentangled irregular component are then used to calculate the MSR criterion upon which the seasonal filter designated for DOW extraction is finally picked.

Panel (e) shows that under this pragmatic recognition of the DOY dynamics in (21) the shorter 3×5 filter is automatically selected for DOW extraction across all data vintages. The subsequent selection of the DOY extraction filter is also affected as the 3×15 filter is now favoured even more frequently over the relatively longer 3×23 filter. Panel (f) finally reveals that the under-adjustment issue with respect to the DOW dynamics has now disappeared; the main support of the bimodal kernel density estimate for $\hat{\rho}_7$ has shifted significantly from $[-0.025, 0.075]$ to $[-0.1, 0]$, whereas the corresponding estimate for $\hat{\rho}_{365}$ is still concentrated over the range $[-0.1, -0.05]$. Overall, the real-time seasonally adjusted electricity consumption series, obtained from applying our proposed redesigned seasonal filter selection for DOW extraction under proper recognition of the DOY pattern, display mild symptoms of systematic over-adjustment with respect to either seasonal pattern, replicating essentially the seasonally adjusted estimates obtained from using the X-11 default filters.

5 Summary

A new version of the JDemetra+ time series software for official statistics has been released recently, which, amongst many other features, implements a modification of the classical X-11 seasonal adjustment method that is applicable to time series with any seasonal periodicity. At the moment, however, this extension merely consists of core routines and hence does not inherit all the established details from the classical X-11 method; in particular, it has not analogue to the well-known data-driven selection procedure for X-11's legacy $3 \times k$ seasonal moving averages based upon the moving seasonality ratio. Going back to the roots, we therefore suggest a generic redesign of this classical rule. Allowing for a broader set of candidate seasonal filters, the key idea is to consider multiple principles for deriving the requisite asymmetric variants and, applying the exact same derivation principle to all $3 \times k$ filters, to calculate multiple approach-specific classical selection rules. The thresholds defining those rules are eventually averaged, which does not only increase the robustness of the final filter choice but also yields in a natural way the indifference regions where the moving seasonality ratio needs to be recalculated from suitably shortened detrended observations.

Quarterly, monthly, and daily real-time macroeconomic data for Germany are considered under a concurrent seasonal adjustment revision policy for illustrating one specific redesigned selection rule based upon threshold quartiles. The examples on quarterly GDP and monthly industrial production demonstrate that the classical and redesigned rules often select the same seasonal moving averages and that, in the occasional case of divergent choices, the redesigned rule picks slightly longer seasonal filters and hence tends to generate more filter revisions over time. The example on daily electricity consumption stresses the fact that the moving seasonality ratio is likely to be biased upwards when the presence of multiple seasonal patterns in the raw data is not properly taken into account; over time, this lapse may lead to a repeated automatic selection of seasonal filters that are too long, which ultimately increases the probability of running into a systematic under-adjustment inadequacy in the sense of residual seasonality being consistently present in real-time seasonally adjusted estimates. To mitigate such issues, we propose incorporating the required number of additional tentative seasonal estimators in the calculation of the moving seasonality ratio; following this pragmatic strategy for daily electricity consumption results in the consistent data-driven selection of shorter seasonal filters for extracting both the day-of-the-week and day-of-the-year patterns.

Future research could experiment with additional candidate seasonal filters and/or other principles for deriving asymmetric variants to further fine-tune selected elements of our proposed redesign. Given the multitude of options for kernel-based trend-cycle filtering available in the extended X-11 method, another interesting question concerns the potential influence that the actual user specification of these options may have on the subsequent redesigned automatic selection of the seasonal moving averages (recall that the moving seasonality ratio is always calculated from detrended data). Data-driven selection rules that integrate both trend-cycle and seasonal filters could also be developed in this context, utilising, for instance, model-based benchmark seasonal adjustments yet again (see Section 1 for references). Alternatively, one could also think of a more flexible modelling framework that allows for replacing the slope parameter in the linearity assumption (8), which is imposed on the moving portion of seasonality near the sample boundaries,

with a more general measure for the degree of moving seasonality; for example, Stock (2013) has already alluded to the possibility of linking such a measure to an appropriate dispersion parameter within the framework underlying the inverse Canova-Hansen test.

References

Bell, W. R. and S. C. Hillmer (1984, October). Issues Involved With the Seasonal Adjustment of Economic Time Series. *Journal of Business & Economic Statistics 2*(4), 291–320. DOI.

Burman, J. P. (1980, May). Seasonal Adjustment by Signal Extraction. *Journal of the Royal Statistical Society: Series A (General) 143*(3), 321–337. DOI.

Chu, Y.-J., G. C. Tiao, and W. R. Bell (2012, October). A Mean Squared Error Criterion for Comparing X-12-ARIMA and Model-Based Seasonal Adjustment Filters. *Taiwan Economic Forecast and Policy 43*(1), 1–32. PDF.

Cleveland, W. S. and S. J. Devlin (1980, September). Calendar Effects in Monthly Time Series: Detection by Spectrum Analysis and Graphical Methods. *Journal of the American Statistical Association 75*(371), 487–496. DOI.

Dagum, E. B. (1978, September). Modelling, Forecasting and Seasonally Adjusting Economic Time Series with the X-11-ARIMA Method. *Journal of the Royal Statistical Society: Series D (The Statistician) 27*(3–4), 203–216. DOI.

Dagum, E. B. (1980). The X-11-ARIMA Seasonal Adjustment Method. Catalogue 12-564E, Statistics Canada, Ottawa.

Dagum, E. B. (1988). The X-11-ARIMA/88 Seasonal Adjustment Method – Foundations and User's Manual. Time Series Research and Analysis Division, Statistics Canada, Ottawa.

Dagum, E. B. and S. Bianconcini (2008, October). The Henderson Smoother in Reproducing Kernel Hilbert Space. *Journal of Business & Economic Statistics 26*(4), 536–545. DOI.

Dagum, E. B. and S. Bianconcini (2013). A Unified View of Nonparametric Trend-Cycle Predictors Via Reproducing Kernel Hilbert Spaces. *Econometric Reviews 32*(7), 848–867. DOI.

Dagum, E. B. and S. Bianconcini (2015, September). A New Set of Asymmetric Filters for Tracking Short-Term Trend in Real-Time. *Annals of Applied Statistics 9*(3), 1433–1458. DOI.

Dagum, E. B., N. Chhab, and B. Solomon (1991, June). The Autocorrelation of Residuals from the X11ARIMA Method. *Journal of Official Statistics 7*(2), 181–194. PDF.

Depoutot, R. and C. Planas (1998). Comparing Seasonal Adjustment and Trend Extraction Filters with Application to a Model-Based Selection of X11 Linear Filters. Working Paper No 9/1998/A/9, Eurostat, Luxembourg.

Doherty, M. (2001, December). The Surrogate Henderson Filters in X-11. *Australian & New Zealand Journal of Statistics 43*(4), 385–392. DOI.

Findley, D. F., B. C. Monsell, W. R. Bell, M. C. Otto, and B.-C. Chen (1998, April). New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program. *Journal of Business & Economic Statistics 16*(2), 127–152. DOI.

Gasser, T. and H.-G. Müller (1979). Kernel Estimation of Regression Functions. In T. Gasser and M. Rosenblatt (Eds.), *Smoothing Techniques for Curve Estimation*, Volume 757 of *Lecture Notes in Mathematics*, pp. 23–68. Heidelberg: Springer. DOI.

Gómez, V. and A. Maravall (2001). Automatic Modeling Methods for Univariate Series. In D. Peña, G. C. Tiao, and R. S. Tsay (Eds.), *A Course in Time Series Analysis*, pp. 171–201. New York: Wiley. DOI.

Grudkowska, S. (2017). JDemetra+ Reference Manual, Version 2.2. Narodowy Bank Polski, Warsaw. URL.

Grun-Rehomme, M., F. Guggemos, and D. Ladiray (2018). Asymmetric Moving Averages Minimizing Phase Shift. In G. L. Mazzi, D. Ladiray, and D. A. Riester (Eds.), *Handbook on Seasonal Adjustment*, Chapter 15, pp. 391–413. Luxembourg: Publications Office of the European Union.

Hillmer, S. C. and G. C. Tiao (1982, March). An ARIMA Model-Based Approach to Seasonal Adjustment. *Journal of the American Statistical Association 77*(377), 63–70. DOI.

Hyndman, R. J. and Y. Fan (1996, November). Sample Quantiles in Statistical Packages. *The American Statistician 50*(4), 361–365. DOI.

Jarmin, R. S. (2019, Winter). Evolving Measurement for an Evolving Economy: Thoughts on the 21st Century US Economic Statistics. *Journal of Economic Perspectives 33*(1), 165–184. DOI.

Kendall, M. and J. K. Ord (1990). *Time Series* (Third ed.). London: Edward Arnold.

Kenny, P. B. and J. Durbin (1982, January). Local Trend Estimation and Seasonal Adjustment of Economic and Social Time Series. *Journal of the Royal Statistical Society: Series A (General) 145*(1), 1–28. DOI.

Ladiray, D., J. Palate, G. L. Mazzi, and T. Proietti (2018). Seasonal Adjustment of Daily and Weekly Data. In G. L. Mazzi, D. Ladiray, and D. A. Riester (Eds.), *Handbook on Seasonal Adjustment*, Chapter 29, pp. 757–783. Luxembourg: Publications Office of the European Union.

Ladiray, D. and B. Quenneville (2001). *Seasonal Adjustment with the X-11 Method*, Volume 158 of *Lecture Notes in Statistics*. New York: Springer. DOI.

Ladiray, D. and B. Quenneville (2004, October). Implementation Issues on Shrinkage Estimators for Seasonal Factors Within the X-11 Seasonal Adjustment Method. *International Journal of Forecasting 20*(4), 557–560. DOI.

Laniel, N. (1986, July). Design Criteria for 13-Terms Henderson End-Weights. Working Paper TSRA-86-011, Time Series Research and Analysis Division, Methodology Branch, Statistics Canada, Ottawa. PDF.

Leone, F. C., L. S. Nelson, and R. B. Nottingham (1961, November). The Folded Normal Distribution. *Technometrics 3*(4), 543–550. DOI.

Lothian, J. (1978, October). The Identification and Treatment of Moving Seasonality in the X-11 Seasonal Adjustment Method. Research Paper 78-10-004, Seasonal Adjustment and Time Series Analysis Staff, Statistics Canada.

Lothian, J. (1984). The Identification and Treatment of Moving Seasonality in X-11-ARIMA. In *Proceedings of the Business and Economic Statistics Section*, Alexandria, VA: American Statistical Association, pp. 166–171.

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

Marris, S. N. (1960). The Treatment of Moving Seasonality in Census Method II. In *Seasonal Adjustment on Electronic Computers*, pp. 257–309. Paris: Organisation for Economic Cooperation and Development.

Musgrave, J. C. (1964a). A Set of End Weights to End all End Weights. Working Paper, U.S. Department of Commerce, Bureau of the Census, Washington, D.C. PDF.

Musgrave, J. C. (1964b). Alternative Sets of Weights for Proposed X-11 Seasonal Factor Curve Moving Averages. Working Paper, U.S. Department of Commerce, Bureau of the Census, Washington, D.C. PDF.

Proietti, T. and A. Luati (2008, December). Real Time Estimation in Local Polynomial Regression, with Application to Trend-Cycle Analysis. *Annals of Applied Statistics 2*(4), 1523–1553. DOI.

Quartier-la-Tente, A. (2024, December). Improving Real-Time Trend Estimates Using Local Parametrization of Polynomial Regression Filters. *Journal of Official Statistics 40*(4), 685–715. DOI.

Quenneville, B., D. Ladiray, and B. Lefrançois (2003, October–December). A Note on Musgrave Asymmetrical Trend-Cycle Filters. *International Journal of Forecasting 19*(4), 727–734. DOI.

Radermacher, W. J. (2019, November). Governing-by-the-Numbers/Statistical Governance: Reflections on the Future of Official Statistics in a Digital and Globalised Society. *Statistical Journal of the IAOS 35*(4), 519–537. DOI.

Radermacher, W. J. (2020). *Official Statistics 4.0 – Verified Facts for People in the 21st Century*. Heidelberg: Springer. DOI.

Shiskin, J. (1969, September). The Census Bureau Seasonal Adjustment Program. *Business Economics 4*(4), 71–73. URL.

Shiskin, J. and H. Eisenpress (1957, December). Seasonal Adjustments by Electronic Computer Methods. *Journal of the American Statistical Association 52*(280), 415–449. DOI.

Shiskin, J., A. H. Young, and J. C. Musgrave (1967). The X-11 Variant of the Census Method II Seasonal Adjustment Program. Technical Paper No 15, U.S. Department of Commerce, Bureau of the Census, Washington, D.C. PDF.

Smyk, A., T. Barthelemy, M. N. Filgueira, F. A. Perez, L. S. Sande, and K. Webel (2025, March). JDemetra+ Documentation. URL.

Stock, J. H. (2013). Comment on "Unseasonal Seasonals?". *Brookings Papers on Economic Activity 2013*(Fall), 111–119. PDF.

U.S. Census Bureau (2025, April). X-13ARIMA-SEATS Reference Manual, Version 1.1. Time Series Research Staff, Center for Statistical Research and Methodology, U.S. Census Bureau, Washington, D.C. PDF.

Webel, K. (2022). A Review of Some Recent Developments in the Modelling and Seasonal Adjustment of Infra-Monthly Time Series. Discussion Paper 31/2022, Deutsche Bundesbank, Frankfurt am Main. PDF.

Webel, K. (2025, June). Revisions in Concurrent Seasonal Adjustments of Daily and Weekly Economic Time Series. *Statistical Journal of the IAOS 41*(2), 265–293. DOI.

Webel, K. and A. Smyk (2024, December). Seasonal Adjustment of Infra-Monthly Time Series with JDemetra+. *Journal of Official Statistics 40*(4), 783–828. DOI.