Data-driven selection criteria for X-13ARIMA-SEATS seasonal adjustment algorithms

Karsten Webel*

Abstract

When choosing a software package for conducting seasonal adjustment as part of their daily routines, decisions of many statistical agencies are based on pragmatic reasons. Then, usually all time series, or at least broad subsets thereof, are seasonally adjusted according to the approach implemented in the chosen software package. Recent releases of X-13ARIMA-SEATS and (J)Demetra+ may change habits as these programs include both a combination of the nonparametric X-11 approach with ARIMA forecasts and the parametric SEATS approach. Hence, users may choose between both methods for each particular time series under review. Accordingly, the question naturally arises which criteria this selection should be based on. We develop a decision tree that relies on both theoretical and empirical issues. In particular, the latter include visual inspection of squared gains of final X-11 and SEATS seasonal adjustment filters as well as calculation of diverse revision measures, which is illustrated using German turnover data. Thereby, we also show that running the SEATS algorithm in X-13ARIMA-SEATS with default options may lead to results that are somewhat misleading.

Key Words: ARIMA model-based approach, signal extraction, unobserved components, X-11 approach

1. Motivation

Many macroeconomic time series sampled usually at a monthly or quarterly frequency share the property that they exhibit a considerable amount of seasonal variation, i.e. fluctuations that recur each year in the same period with similar intensity. Since it is generally agreed that economic developments are best judged from indicators revealing new information, seasonal adjustment has become a standard tool in both official statistics and academic research. Unsurprisingly, different strategies for estimating seasonal movements have been set up over the last decades, see Bell et al. (2012), Ghysels and Osborn (2001) and Hylleberg (1992), among others, for an overview of both theoretical developments and applications to real-world data. In practice, however, the X-11 and SEATS core developed originally by Shiskin et al. (1967) and Gómez and Maravall (1996), respectively, have probably been the two methods applied most frequently. Since both cores have been implemented in individual software packages, such as X-12-ARIMA and TRAMO/SEATS, decisions of many practitioners at statistical agencies on the method employed have been directly depended on the choice of their preferred seasonal adjustment program. Often, the latter has been based on pragmatic reasons, such as employees' individual backgrounds, data users' demands (e.g. for low revisions) and the program's suitability for statistical mass production. Then, basically all time series have been seasonally adjusted according to the approach implemented in the favourite software package.

Recent releases of X-13ARIMA-SEATS and (J)Demetra+ may change habits as they jointly incorporate the X-11 and SEATS core. Users of these programs may thus choose between a nonparametric and the parametric ARIMA model-based (AMB) method for each particular time series under review without the need to switch software packages. This

*Deutsche Bundesbank, Central Office, General Economic Statistics Division, Statistics Department, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt, Germany

immediately raises the question of which criteria one should focus on when making this choice. For that purpose, we suggest a decision tree that combines theoretical considerations regarding conceptual differences between the philosophies underlying the X-11 and AMB approach with empirical findings.

The remainder of this paper is organised as follows. Section 2 provides the notational framework as well as basic ideas of both seasonal adjustment methods considered. The decision tree for choosing between them is presented in Section 3, followed by an empirical illustration using selected turnover data for Germany in Section 4. Section 5 enlarges on this illustration as it demonstrates that employing the SEATS core with default options as implemented in X-13ARIMA-SEATS may lead to results that are somewhat misleading. Finally, Section 6 concludes.

2. Preliminaries

2.1 Notations

Let $\{x_t\}$ denote the time series under review. We assume that it can be decomposed into two orthogonal unobserved components according to

$$x_t = s_t \circ n_t, \tag{1}$$

where signal $\{s_t\}$ and noise $\{n_t\}$ are supposed to capture all nonseasonal and seasonal variations, respectively, of $\{x_t\}$. In most cases, the decomposition is additive or multiplicative, i.e. $\circ \in \{+, \cdot\}$. Either way, seasonal adjustment is interpreted as a signal extraction problem which can be solved by linear filtering.[1] To see this, let $\mathbf{x} = (x_1, \ldots, x_T)^\top$ denote a finite realisation of $\{x_t\}$. We may then express any estimator of the seasonally adjusted time series as

$$\hat{\mathbf{s}} = \mathcal{W}\,\mathbf{x}, \tag{2}$$

where $\hat{\mathbf{s}} = (\hat{s}_1, \ldots, \hat{s}_T)^\top$ and $\mathcal{W} \in \mathbb{R}^{T\times T}$ is a matrix whose t-th row, $\mathcal{W}^{(t)}$, contains the filter weights employed to estimate s_t. We may thus rewrite (2) for each $t \in \{1, \ldots, T\}$ as

$$\hat{s}_t = \mathcal{W}^{(t)}\,\mathbf{x} = \sum_{j=1}^{T} \mathcal{W}_j^{(t)}\, x_j. \tag{3}$$

Hence, weights stored in $\mathcal{W}$ do not only depend on t but also on the number of observations available, T. Assuming for simplicity that the latter is odd, i.e. $T = 2k+1$ for some $k \in \mathbb{N}$, the (symmetric) central seasonal adjustment filter $\mathcal{W}^{(k+1)}$ and the (asymmetric) concurrent seasonal adjustment filter $\mathcal{W}^{(T)}$ are usually of greatest interest. The whole matrix $\mathcal{W}$ can be easily derived from seasonally adjusting the set of T unit vectors. This impulse response method is comprehensively described in Section 3.4 of Ladiray and Quenneville (2001).

Often, the performance of any seasonal adjustment filter occurring in (3) is better judged in the spectral rather than the time domain. To see this, let ψ denote any linear filter transforming data $\{x_t\}$ into output $\{y_t\}$ via $y_t = \sum_j \psi_j\, x_{t-j}$. For monthly data, its gain is defined by

$$g_\psi(\lambda) = \left| \sum_j \psi_j\, e^{-ij\lambda\pi/6} \right|, \tag{4}$$

[1] In practice, seasonal adjustment is a highly nonlinear data transformation as it comprises, for example, temporary removal of outliers which may disturb estimation of any unobserved component. We shall thus assume for convenience that $\{x_t\}$ has already been linearised appropriately.

where $\lambda \in [0, 6]$ is in units of cycles per year. If $\{x_t\}$ is a stationary time series with spectral density $f_x(\lambda)$, then the spectral density of $\{y_t\}$ is given by $f_y(\lambda) = g^2_\psi(\lambda)\, f_x(\lambda)$. The squared gain of ψ thus signals suppression or amplification of its output's variance component of frequency λ when $g^2_\psi(\lambda) < 1$ and $g^2_\psi(\lambda) > 1$, respectively. Regarding seasonal adjustment, both input $\mathbf{x}$ and output $\hat{\mathbf{s}}$ of respective filters $\mathcal{W}^{(t)}$ are typically nonstationary. However, the principle outlined above also applies in a more general way to this kind of time series, see Section 2 of Findley and Martin (2006) for a brief discussion.

To handle nonstationarity, seasonal ARIMA models are usually employed even though the motivation is quite different for both seasonal adjustment approaches considered here. A time series $\{x_t\}$ is called seasonal ARIMA process if there exists a white noise $\{\varepsilon_t\}$ such that

$$\phi(B)\Phi(B^\tau)\nabla^d\nabla^D_\tau(\{x_t\}) = \theta(B)\Theta(B^\tau)(\{\varepsilon_t\}), \tag{5}$$

where B is the backshift operator, i.e. $B^k x_t = x_{t-k}$ for any $k \in \mathbb{Z}$, τ is the seasonal period, e.g. $\tau = 12$ for monthly data, $\nabla = 1 - B$, $\nabla_\tau = 1 - B^\tau$ and $(d, D) \in \mathbb{N}^2$. Furthermore, $\phi(B) = 1 - \phi_1 B - \cdots - \phi_p B^p$, where $\phi_p \neq 0$, and $\Phi(B^\tau) = 1 - \Phi_1 B^\tau - \cdots - \phi_P B^{\tau P}$, where $\Phi_P \neq 0$, denote the nonseasonal and seasonal AR polynomial. Analogously, $\theta(B) = 1 - \theta_1 B - \cdots \theta_q B^q$ with $\theta_q \neq 0$ and $\Theta(B^\tau) = 1 - \Theta_1 B^\tau - \cdots \Theta_Q B^{\tau Q}$ with $\Theta_Q \neq 0$ are the nonseasonal and seasonal MA polynomial. Thereby, $k \in \mathbb{N} \cup \{\infty\}$ for each $k \in \{p, q, P, Q\}$. We additionally assume that all roots of each of these polynomials lie outside the unit circle, i.e. both AR polynomials are stationary and both MA polynomials are invertible. To indicate a seasonal ARIMA model of type (5), we shall use the standard notation $f\,(pdq)(PDQ)_\tau$ in what follows, where f is any particular Box-Cox transformation that might be applied to the values of $\{x_t\}$ prior to model fitting.

To account for outliers and calendar effects prior to seasonal adjustment of $\{x_t\}$, a time-varying mean function is typically included in (5). Since this model coincides with a linear regression model with (seasonal) ARIMA errors, it is usually referred to as a regARIMA model. This modelling capability is common to both seasonal adjustment approaches implemented in X-13ARIMA-SEATS. We shall thus assume for convenience that calendar effects have already removed from $\{x_t\}$, as opposed to (nonseasonal) outliers which should generally remain visible in seasonally adjusted figures.

2.2 X-11 approach

Within the X-11 framework, the signal of model (1) is composed of a trend-cyclical component $\{t_t\}$ and an irregular component $\{i_t\}$. The former conceptually accounts for long-term movements as well as periodic fluctuations around it whose cycles last longer than one year, while the latter captures nonseasonal short-term random shocks. Assuming a multiplicative decomposition, model (1) can thus be rewritten as

$$x_t = t_t \cdot n_t \cdot i_t. \tag{6}$$

Estimation of all unobserved components occurring in (6) is achieved through an iterated application of various predefined linear filters. Each iteration step comprises several subroutines which basically follow a set pattern. The trend-cyclical component is extracted first by smoothing the subroutine's "input" series with either a simple moving average or a Henderson filter, depending on the particular subroutine. The latter type of filter can be chosen automatically according to the I/C ratio or manually by the user. After removal of these trend-cyclical movements, an estimate of the seasonal component is obtained from smoothing the corrected "input" series (also called seasonal-irregular component) for each period with a weighted $3 \times k$ moving average, where $k \in \{1, 3, 5, 9, 15\}$. As for Henderson filters, seasonal filters can be chosen automatically based on the (global) I/S ratio or manually. In the latter case, different seasonal filters may be selected for different periods. To reduce the size of distortions potentially caused by extreme values in the seasonal-irregular component, the X-11 routine is also equipped with an automatic detection and down-weighting procedure of such extremes. Eventually, removal of estimated seasonal fluctuations from the "input" series yields an estimate of its seasonally adjusted version. Throughout each subroutine, symmetric trend and seasonal filters are considered whenever possible. In case of boundary problems, asymmetric versions are automatically applied.

Final X-11 seasonal adjustment filters, whose weights are stored in $\mathcal{W}$ of (2), thus result from the convolution of all trend and seasonal filters chosen during each iteration step. Several programs, such as X-12-ARIMA and X-13ARIMA-SEATS, provide additional features to enhance the basic X-11 principle outlined above. Most prominently, regARIMA forecasts of unadjusted figures can be used to mitigate boundary problems by enabling application of less asymmetric X-11 filters, especially at the current end of the time series under review. Further details are provided by Findley et al. (1998), Section 4 of Ghysels and Osborn (2001) and Ladiray and Quenneville (2001), among others.

2.3 AMB approach

Within the AMB framework, nonseasonal intra-year variability in unadjusted figures is further specified as the irregular component of model (6) is decomposed according to

$$i_t = r_t \circ w_t.$$

Thereby, $\{w_t\}$ is white noise, which is primarily assumed to facilitate testing and statistical interpreting. To also capture fluctuations not persistent enough to be considered trend-cyclical but still too persistent to reflect white noise like behaviour, the transitory component $\{r_t\}$ is introduced as a rather technical necessity. Assuming an additive decomposition, model (1) thus turns into

$$x_t = t_t + n_t + r_t + w_t. \tag{7}$$

AMB theory's key assumption states that each unobserved component of (7) is driven by an individual ARIMA process. For the sake of brevity, however, we shall stick to model (1) instead of (7) in what follows and assume that ARIMA models for signal and noise are given by

$$\phi^{(k)}(B)(\{k_t\}) = \theta^{(k)}(B)\left(\left\{\varepsilon_t^{(k)}\right\}\right),$$

where $k \in \{s, n\}$ and $\phi^{(k)}$ and $\theta^{(k)}$ are AR and MA polynomials of orders $p^{(k)}$ and $q^{(k)}$, respectively. Their roots are assumed to lie on or outside the unit circle, however, they do not share common roots for each k. Similarly, both AR (MA) polynomials do not have common (unit) roots. Eventually, both innovation sequences are uncorrelated Gaussian white noise processes. By construction, $\{x_t\}$ then admits the ARIMA representation

$$\phi^{(x)}(B)(\{x_t\}) = \theta^{(x)}(B)(\{\varepsilon_t\}),$$

where $\phi^{(x)}(B) = \phi^{(s)}(B)\phi^{(n)}(B)$ and

$$\theta^{(x)}(B)(\{\varepsilon_t\}) = \phi^{(n)}(B)\theta^{(s)}(B)\left(\left\{\varepsilon_t^{(s)}\right\}\right) + \phi^{(s)}(B)\theta^{(n)}(B)\left(\left\{\varepsilon_t^{(n)}\right\}\right). \tag{8}$$

In theory, the minimum mean squared error (MMSE) estimator of the signal, i.e. the estimator $\{\hat{s}_t\}$ that minimises $\mathbb{E}\left[(s_t - \hat{s}_t)^2 \middle| \mathbf{x}\right]$, can be easily obtained if $\mathbf{x}$ is infinite and both signal and noise are stationary. As shown by Whittle (1963), $\{\hat{s}_t\} = \nu(B)(\{x_t\})$ in this case where ν is the Wiener-Kolmogorov (WK) filter given by

$$\nu(B) = \frac{\sigma^2_{\varepsilon^{(s)}}}{\sigma^2_{\varepsilon}} \frac{\theta^{(s)}(B)\phi^{(n)}(B)}{\theta^{(x)}(B)} \frac{\theta^{(s)}(B^{-1})\phi^{(n)}(B^{-1})}{\theta^{(x)}(B^{-1})}. \tag{9}$$

Thereby, variances belong to the innovation sequences $\left(\left\{\varepsilon^{(s)}_t\right\}\right)$ and $\{\varepsilon_t\}$ occurring in (8). Several authors, including Bell (1984) and Maravall (1988), demonstrate that the WK filter still yields the MMSE estimator if $\mathbf{x}$ is finite, as in (2), and both signal and noise are nonstationary.

In practice, however, the quantities forming the numerator of (9) are unknown and have to be derived from the ARIMA model fitted to the observed time series. To this end, the decomposition algorithm developed originally by Burman (1980) and improved by Bell and Hillmer (1984) and Hillmer and Tiao (1982) is employed. Basically, this is a two-step procedure. By analogy with (5), $\hat{\phi}^{(x)}$ is factorised first. Using the identity

$$\hat{\phi}(B)\hat{\Phi}(B^\tau)\nabla^d\nabla^D_\tau = \hat{\phi}^{(s)}\hat{\phi}^{(n)},$$

(unit) roots of all polynomials occurring on the left hand side are assigned to either signal or noise according to their associated frequencies. For example, the nonseasonal unit root of ∇ is associated with the trend frequency $\lambda = 0$ and thus assigned to the signal. In a second step, unknown MA polynomials and variances appearing in the right hand side of (8) are derived from a partial fraction expansion. At this point, the canonical assumption has to be made to achieve a unique decomposition. The latter basically states that the variance of the white noise component is maximised. If this assumption cannot be met, which is likely to happen for rather complex ARIMA models, several approximations can be considered. (A detailed documentation of these approximations is currently work in progress at the Bank of Spain.)

Eventually, unknown quantities of (9) are replaced with estimates obtained from the decomposition sketched above to yield the estimated WK filter. Its weights thus directly constitute the weighting matrix $\mathcal{W}$ of (2) for the final SEATS seasonal adjustment filter. Further details are given by Gómez and Maravall (2001) and Maravall (1995), inter alia.

3. Decision tree

We suggest the three-step procedure illustrated in Figure 1 to choose between the X-11 and AMB approach for any observed time series. For obvious reason, the latter should be examined for presence of stable seasonality in the first step to ensure that it actually is in need for seasonal adjustment.

If so, conceptual differences between both methods can be regarded in a second step, which is labelled “first empirical considerations” in Figure 1. Its basic idea is that the X-11 approach has advantages whenever the AMB method’s key assumption of the data generating process (DGP) coinciding with an ARIMA model (with time-constant coefficients) is likely to be too restrictive to adequately account for all dynamics of the observed time series. This might be the case for various phenomena but we shall only focus on two of them here. First, seasonal heteroskedasticity cannot be modelled explicitly within the AMB framework and has to be captured as good as possible by seasonal ARIMA models instead. In contrast, GARCH-type effects may be taken into account, at least up to some degree, within X-11’s built-in extreme value detection procedure by considering period-specific variances of the irregular component. If their importance, however, is not sufficiently high, reliance on the X-11 or AMB approach usually does not really make a difference as far as seasonal heteroskedasticity is concerned, at least according to our experience. In addition,

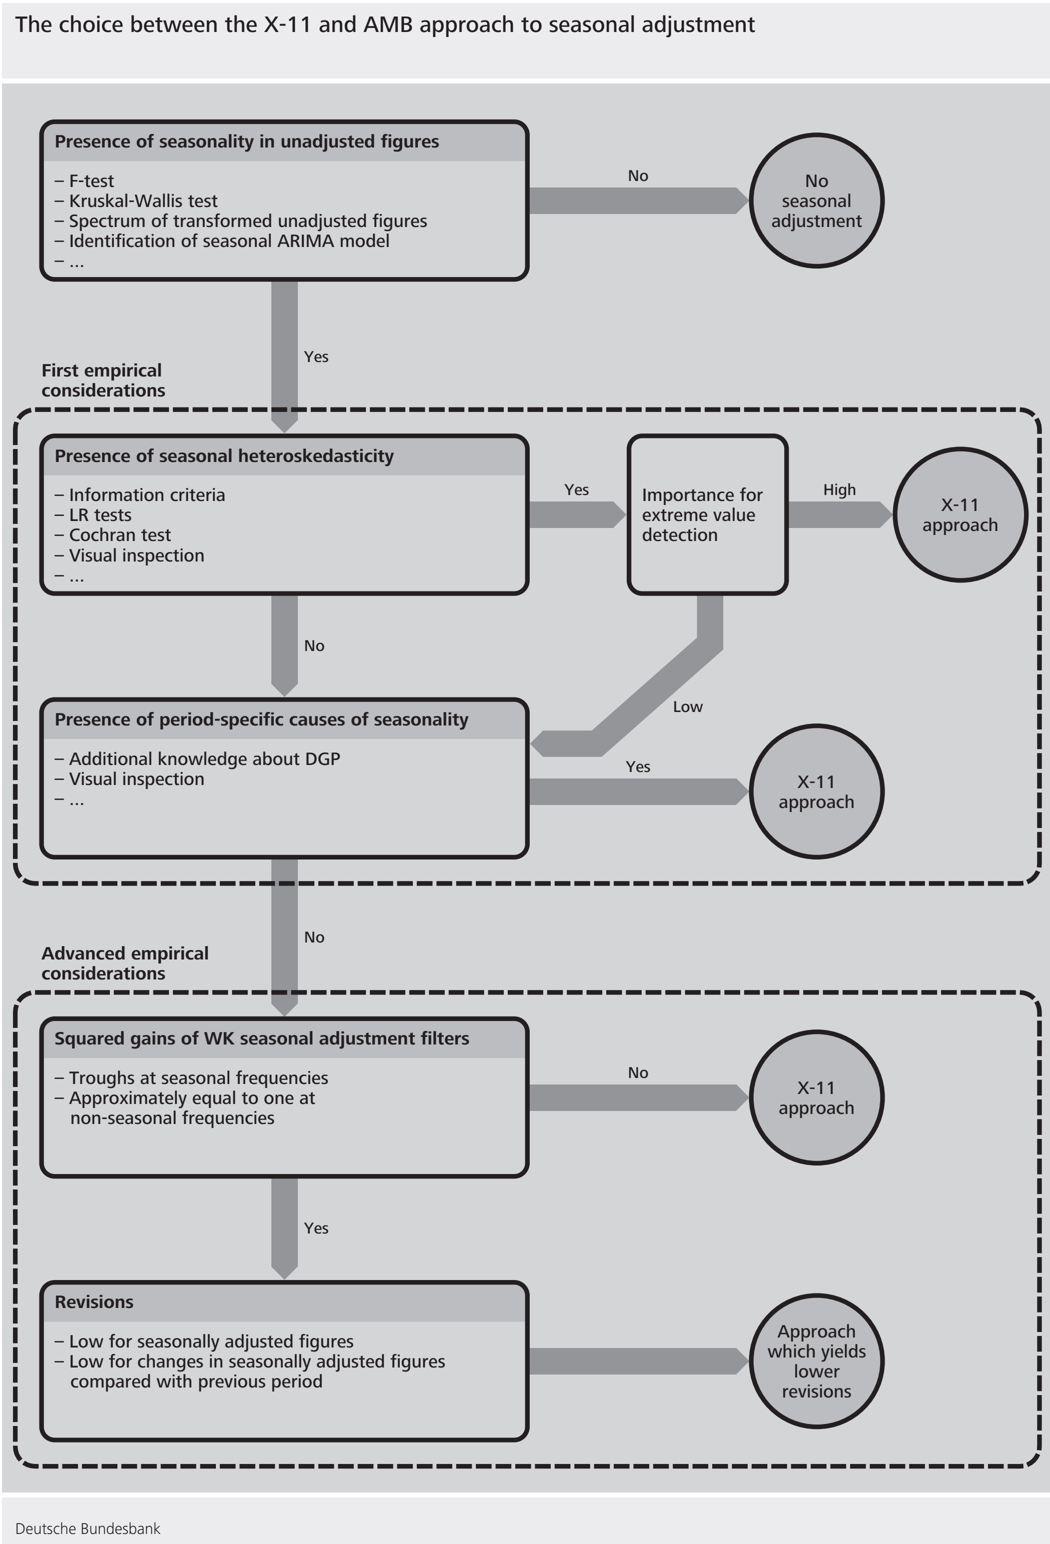

Figure 1: Decision tree.

it should be kept in mind that this phenomenon is somehow related to sample size as it is less likely to be apparent in short and moderate-length time series. Consequently, the AMB approach may look more viable a priori for those kind of series.

The second situation where the X-11 method might be more preferable emerges when major causes of (moving) seasonality in observed unadjusted figures are substantially different across periods. A prime example is given by German retail sales. Their usual peak in December attributable to the Christmas business has decreased faster than seasonal movements in other months in the recent past. A widely accepted explanation is that Christmas bonuses have been paid by fewer firms or have been cut steadily compared to regular salaries. Unadjusted retail sales for December should thus be treated somewhat individually by, for example, application of shorter seasonal filters in X-11 as compared to other months which are unlikely to be affected by this change. As for seasonal heteroskedasticity, the AMB approach is intentionally not capable of modelling such effects explicitly.

In practice, many observed time series behave rather less exceptionally and, hence, the considerations outlined so far do usually not suffice to select an appropriate seasonal adjustment approach for all of them. Accordingly, further criteria summarised as “advanced empirical considerations” in Figure 1 should be employed in a third step. At first, squared gains of (selected) final X-11 and SEATS seasonal adjustment filters are compared. The key idea is that any “acceptable” seasonal adjustment filter should completely eliminate seasonal fluctuations without altering movements associated with nonseasonal frequencies. Assuming stable seasonality, the squared gain of an ideal seasonal adjustment filter would thus be equal to one for all frequencies except the seasonal ones where it immediately drops to zero. However, virtually all seasonal adjustment filters derived from observed time series display squared gains that deviate from this ideal, mainly due to the finite length of and the amount of moving seasonality inherent in the latter. Therefore, the minimum requirement any “acceptable” seasonal adjustment filter should meet is that its squared gain stays close to one at all nonseasonal frequencies and exhibits dips at all seasonal ones. Since final X-11 seasonal adjustment filters pass this spectral inspection by construction, it has to be checked only for their WK counterparts. If both approaches still perform equally well, we suggest further comparison with respect to revisions they generate in seasonally adjusted figures and their period-on-period changes. The main reason for including the latter is that statements on economic developments typically focus on them rather than seasonally adjusted figures themselves. Eventually, the method which yields lower revisions on average is chosen, which according to our experience should be in line with demands of many data users.

4. Illustration

We run X-13ARIMA-SEATS with default options for both the X-11 and the AMB method, using its Windows companion Win X-13 (version 1.0 build 150). In addition, R (version 2.14.1) is employed to calculate squared gains of X-11 seasonal adjustment filters as these are not provided by any Win X-13 output.

4.1 Data

We consider monthly turnover of industry (IND), producers of capital goods (CAP) and producers of consumer goods (CON) at current prices as available in May 2013.[2] These time series cover the period as of January 1991 up to December 2012 with their annual

[2]The latter two items are defined in Regulation (EC) No 656/2007 and “industry” is defined as their aggregate plus producers of intermediate goods.

Table 1: Results of selected tests.

Series	F-test	Kruskal-Wallis-test	Cochran-test
	Stable seasonality		Homoskedasticity
IND	424.440	246.908	0.137
CAP	291.641	240.373	0.153
CON	256.609	219.077	0.175

averages of 2010 being set to 100. Also, they have already been adjusted for calendar effects using linear regression techniques.

As reported in Table 1, standard tests provide strong evidence of the presence of stable seasonality for each series as their test statistics are larger than respective critical values at any conventional level of significance. A similar conclusion can be drawn from spectral analysis. Estimated spectra of month-on-month changes of the linearised series universally exhibit visible peaks at all seasonal frequencies, even though they are somewhat less pronounced at the first one, see Figure 2. In addition, the DGP of each series is found to be well approximated by a seasonal ARIMA model, see Table 2 below. Seasonal adjustment of these three series thus makes perfect sense.

4.2 First empirical considerations

The last column of Table 1 shows results of (one-sided) tests for equality of month-specific variances of the irregular component (as given in output table B13 of X-13ARIMA-SEATS). Since the critical value is 0.15 for a series of 264 monthly observations at a $5\ \%$ level of significance, this test suggested originally by Cochran (1941) does not indicate presence of seasonal heteroskedasticity in IND but provides some evidence for CAP and CON. This evidence, however, is rather weak and extreme value detection is barely affected by it for both series. The AMB approach is thus not ruled out on these grounds.

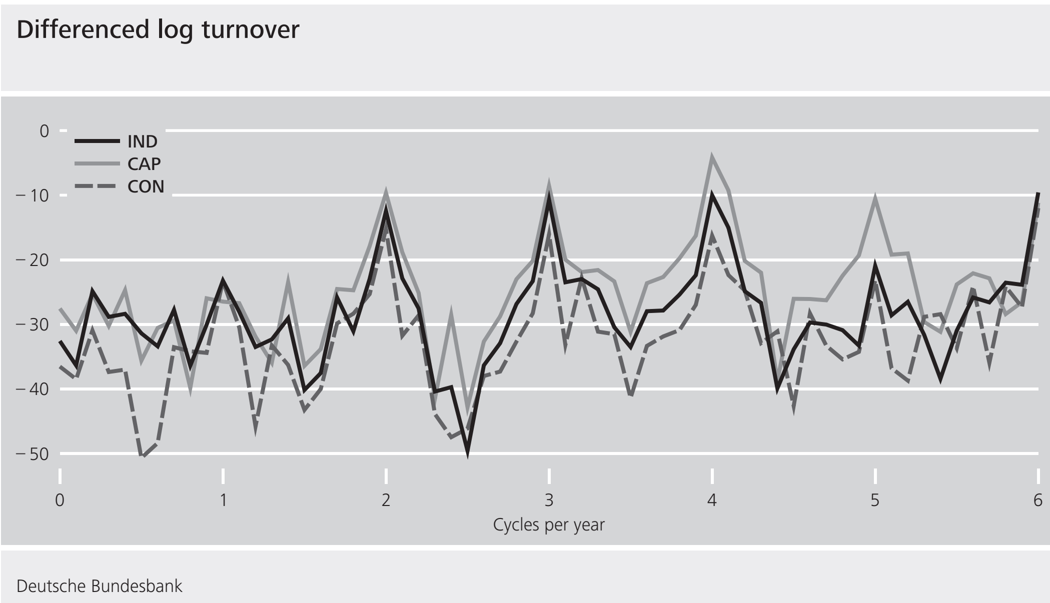

Figure 2: Periodograms (in decibel).

Even though additional knowledge about each turnover indicator's DGP includes detailed information on how data are regularly collected, it does not give statistically sufficient hints of the need for application of different month-specific seasonal filters. We hence cannot decide on an appropriate seasonal adjustment method for any turnover series at this early stage. Accordingly, advanced empirical considerations have to be employed for that purpose.

4.3 Advanced empirical considerations

In the first step, we fit a regARIMA model to the log of each indicator to take outliers into account and to obtain two years of forecasts. The log transformation is preferred by the automatic log-level test implemented in X-13ARIMA-SEATS. To account for the strong economic downturn starting in 2008, we introduce a series of level shifts over the crisis' period. In addition, we make use of automatic outlier detection based on default critical values. Regarding final ARIMA model choices, we similarly rely on the program's automatic model selection procedure. Respective results are summarised in Table 2, where parameter estimates are reported in the order of their appearance in the regARIMA model. To exemplify, the regARIMA model fit to turnover of industry, $\{x_t\}$, is (approximately) given by

$$(1 - 0.37B - 0.08B^2 - 0.26B^3)\nabla\nabla_{12}(\{\tilde{x}_t\}) = (1 - 0.64B)(1 - 0.62B^{12})(\{\varepsilon_t\}), \quad (10)$$

where $\{\tilde{x}_t\}$ is the outlier adjusted version of $\{\log x_t\}$.

In a second step, we compare squared gains of both central and concurrent seasonal adjustment filters obtained from X-11 and SEATS. To derive the weighting matrix $\mathcal{W}$ occurring in (2) according to the impulse response method for X-11, we note that the 13-term Henderson filter is chosen automatically for IND and CAP, whereas the 23-term Henderson filter is selected for CON. In addition, the 3×5 seasonal moving average is automatically applied to all three series. Squared gains of respective X-11 seasonal adjustment filters are then calculated according to (4). Regarding both WK seasonal adjustment filters, no further calculations are necessary since their squared gains are directly provided by X-13ARIMA-SEATS.

Figure 3 shows that both central and concurrent X-11 filters perform reasonable well as their squared gains stay close to one at nonseasonal frequencies and decrease rapidly to

Table 2: RegARIMA modelling summary.

		Parameter estimates (standard errors in parentheses)		
Series	ARIMA model	Nonseasonal	Seasonal	Residual variance
IND	$\log(311)(011)_{12}$	0.3744 (0.15140)	0.6197 (0.05149)	0.0002 (0.00002)
		0.0798 (0.07031)		
		0.2619 (0.06144)		
		0.6407 (0.15043)		
CAP	$\log(011)(011)_{12}$	0.4587 (0.05539)	0.5959 (0.05212)	0.0008 (0.00007)
CON	$\log(311)(011)_{12}$	−1.2608 (0.67744)	0.7947 (0.04012)	0.0002 (0.00002)
		−0.7593 (0.44801)		
		−0.2276 (0.23323)		
		−0.5910 (0.68214)		

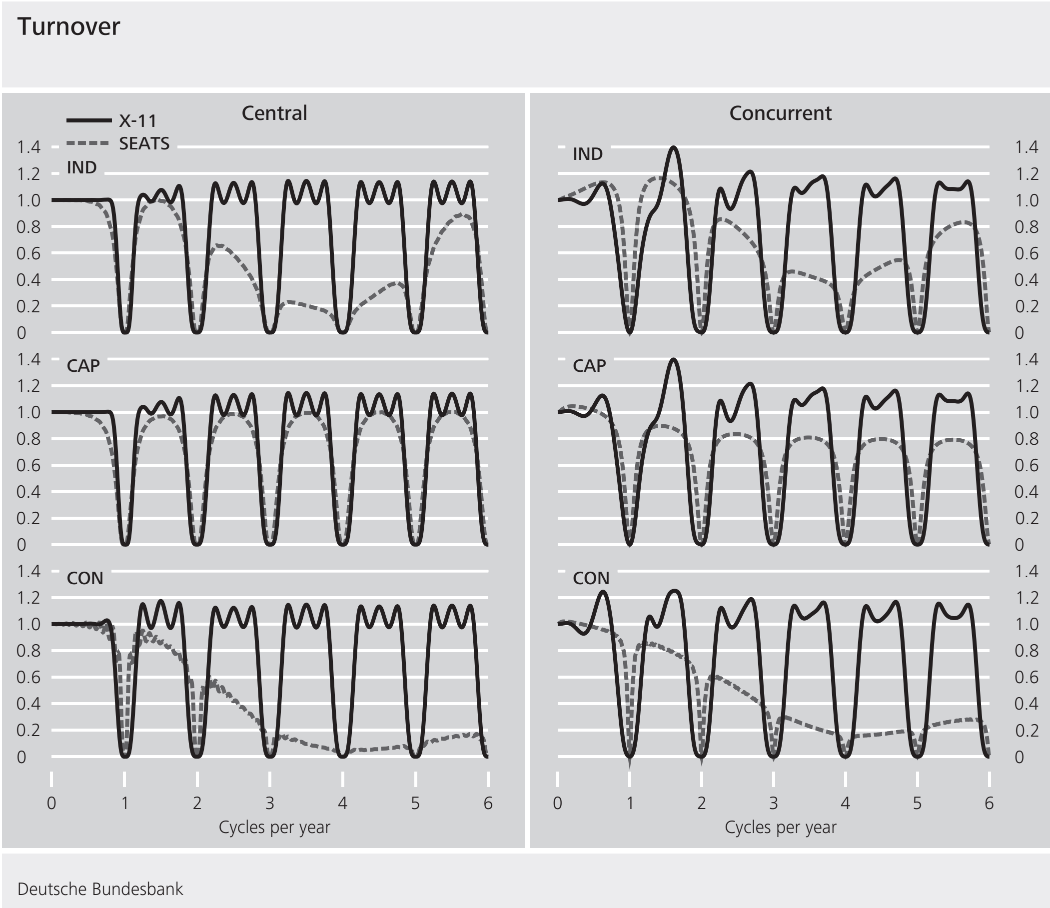

Figure 3: Squared gains of seasonal adjustment filters.

zero in the vicinity of seasonal frequencies, regardless of the series considered. For central X-11 filters, the latter display oscillatory behaviour between seasonal frequencies that is typical of finite filters, signaling potential amplification of fluctuations associated with those frequencies in seasonally adjusted figures. Amplification is also reflected in squared gains of concurrent X-11 filters which stay above one at many nonseasonal frequencies.

Regarding the AMB approach, WK filters perform somewhat poorly for IND and CON, even though they still remove all seasonality contained in unadjusted figures. For IND, squared gains of both central and concurrent WK filter display an awkward U-shaped curvature between the second and fifth seasonal frequency (excluding seasonal frequencies) as they drop to 0.2 and 0.4, respectively, near the fourth seasonal frequency. For CON, squared gains of both WK filters look even worse as they start to decrease immediately behind the first seasonal frequency and do not manage to exceed 0.2 (central) and 0.3 (concurrent) behind the third seasonal frequency. Hence, WK filters heavily suppress many fluctuations in IND and CON not associated with seasonal behaviour, resulting in seasonally adjusted figures that are much smoother than those obtained from X-11, see Figure 5 below for the latter series. The X-11 approach is preferred for both turnover indicators as its overall performance is closer to what a satisfactory seasonal adjustment procedure should achieve.

When looking at CAP, things are completely different. WK seasonal adjustment filters now also yield acceptable results as their squared gains exhibit a smooth curvature that is somewhat close to the ideal described in Section 3. This is particularly true for the central

Table 3: Revision analysis.

		Seasonally adjusted figures (as a percentage)			Changes in seasonally adjusted figures compared with previous period (in percentage points)		
Series	Core	MR	MAR	STD	MR	MAR	STD
IND	X-11	0.0412	0.3668	0.4747	0.0142	0.4443	0.5792
	SEATS*	0.0060	0.3478	0.4300	0.0062	0.4136	0.5266
CAP	X-11	0.0432	0.6192	0.7711	0.0006	0.7683	0.9951
	SEATS	−0.0259	0.5546	0.7041	−0.0009	0.6531	0.8445
CON	X-11	0.0448	0.3302	0.4281	−0.0049	0.5235	0.6546
	X-11*	0.0376	0.3296	0.4252	−0.0070	0.5211	0.6510
	SEATS*	0.0273	0.2881	0.3807	0.0015	0.4837	0.7014

WK filter whose squared gain does not signal amplification of any frequency due to oscillatory behaviour (which is the case for the central X-11 filter). Even though the concurrent WK filter's squared gain indicates minor suppression of nonseasonal intra-year cycles by dropping to 0.8 at respective frequencies, the AMB method seems to be advantageous for CAP. Since this advantage is rather slight than huge, we shall also compare both approaches with respect to revisions before making a final decision.

For that purpose, we calculate diverse revision measures for CAP in a third step. Since we are primarily interested in the way revisions are induced by application of a particular seasonal adjustment method, we do not study them in real-time, excluding corrections of unadjusted figures from our analysis. Instead, we seasonally adjust truncated versions of CAP to generate revisions in its seasonally adjusted figures and their month-on-month changes. During this procedure, the regARIMA model is reestimated for each truncated series. Revisions are then calculated over the time span as of January 2000 up to December 2007, allowing for a "burn-in" period of nine years and excluding an atypical revision pattern that may be caused by the economic downturn starting in 2008. Table 3 reports mean revisions (MR), mean absolute revisions (MAR) and revisions' standard deviations (STD) for CAP that result from both approaches.[3] Since these measures tend to be slightly smaller when seasonal adjustment is carried out with SEATS, the AMB approach is finally chosen for CAP. It should be noted, however, that these seasonally adjusted figures do not substantially differ from the ones obtained from X-11. For example, mean difference and mean absolute difference between both seasonally adjusted CAP series amount to 0.02 and 0.26 index point, respectively.

5. Cautionary remarks

As shown in the previous section, seasonal adjustment according to the AMB approach turns out to be problematic for IND and CON, two series not fit by an airline model. In both cases, squared gains of WK seasonal adjustment filters exhibit an awkward curvature that signals noticeable suppression of nonseasonal movements whose periods are shorter than six months. One may thus suppose that some transitory effects have been falsely

[3]Table 3 also shows these measures for IND and CON even though the X-11 approach has already been selected for both series. This is done in preparation for our discussion of modified X-13ARIMA-SEATS runs in Section 5. Rows corresponding to such modifications are marked with an asterisk.

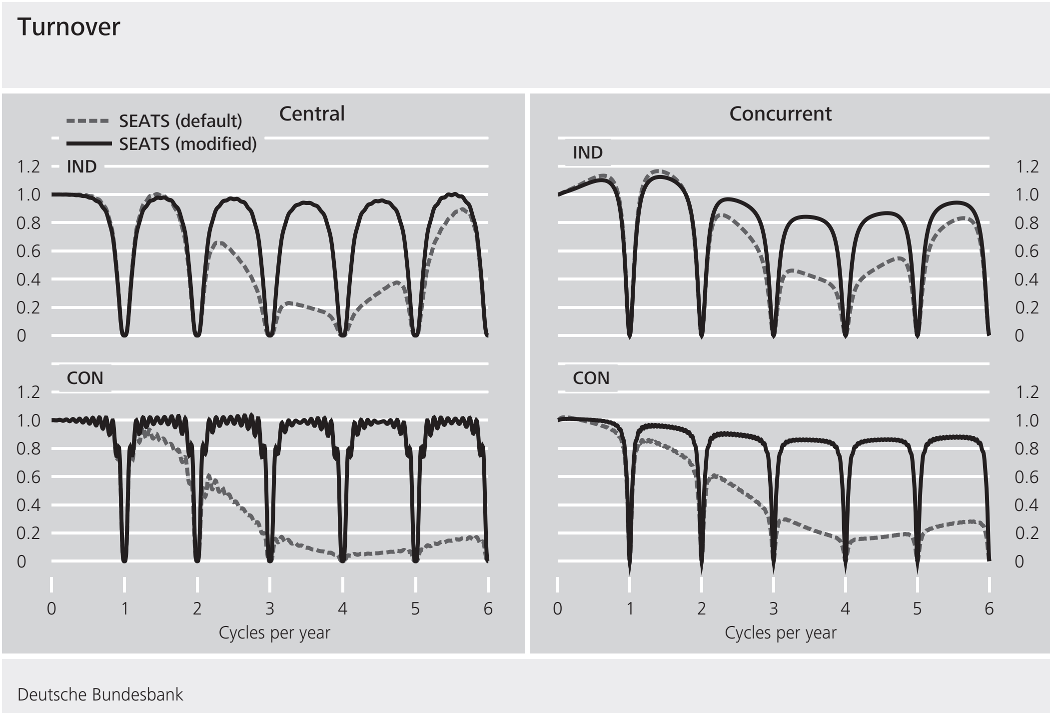

Figure 4: Squared gains of WK seasonal adjustment filters.

attributed to the seasonal component in both cases. To verify this conjecture, we now examine the stationary AR polynomials of both series' regARIMA models. Bearing in mind that we have run X-13ARIMA-SEATS with default options so far, rerunning the SEATS algorithm with appropriate modifications is thus likely to improve the results obtained under the former setting.

5.1 Industry

The stationary AR polynomial of (10) factorises as

$$(1 - 0.84B)(1 + 0.47B + 0.31B^2). \tag{11}$$

In fact, the SEATS decomposition algorithm correctly assigns the real positive root of (11) to the trend-cyclical component since its (reciprocal) modulus is larger than the default threshold 0.5. However, the complex root of (11) is assigned to the seasonal component since its associated frequency is approximately $0.64\,\pi$ which, under default settings, is sufficiently close to the fourth seasonal frequency, $2\pi/3$. We thus reduce this range of "near-seasonal" frequencies by lowering the `epsphi` argument of the `SEATS` spec to enforce introduction of an additional transitory component that accounts for the complex root of (11). As a result, the seasonal component captures only "true" seasonal fluctuations as its assigned AR polynomial now coincides with the annual aggregation operator that is part of ∇_{12} in (10). As shown in the top row of Figure 4, WK seasonal adjustment filters perform much better under this modification than under default settings as their squared gains somehow retain their overall curvature but stay significantly closer to one at virtually all nonseasonal frequencies.

What is more, the X-11 and AMB approach now yield very similar results when comparing squared gains shown in Figure 3 and Figure 4. A further study of revisions generated for both methods seems thus sensible. As reported in Table 3, the latter are slightly

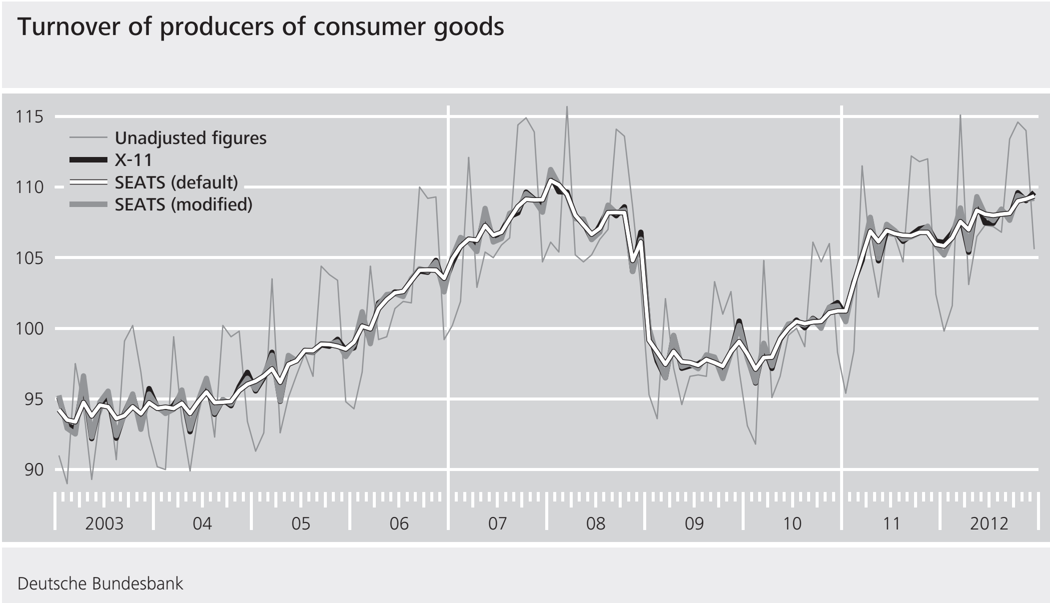

Figure 5: Seasonally adjusted figures.

smaller on average for both seasonally adjusted figures and their month-on-month changes in case SEATS is used. However, this advantage might be considered insignificant by some users. We thus conclude that (based on our decision tree) both methods seem to be equally appropriate for IND after certain changes of default options controlling the SEATS decomposition.

5.2 Consumer goods

Beside false assignment of transitory effects, investigation of the regARIMA model fit to CON reveals an additional effect that may cause severe problems. As shown in Table 2, all nonseasonal parameter estimates combine with exceptionally high standard errors. Cancellation of a common nonseasonal AR and MA factor might be held responsible for this kind of instability. In fact, the stationary AR polynomial factorises as

$$(1 + 0.63B)(1 + 0.63B + 0.36B^2). \tag{12}$$

The real negative root of (12) is associated with the sixth seasonal frequency, π, as is the nonseasonal MA polynomial $(1 + 0.59B)$. Since both roots have approximately the same modulus, they nearly cancel and we thus drop them from the model. Reestimation of a reduced version of type $\log(210)(011)_{12}$ for the CON series, $\{x_t\}$, leaves all parameter estimates basically unchanged as it yields

$$(1 + 0.66B + 0.36B^2)\nabla\nabla_{12}(\{\tilde{x}_t\}) = (1 - 0.79B^{12})(\{\varepsilon_t\}), \tag{13}$$

where $\{\tilde{x}_t\}$ is the outlier adjusted version of $\{\log x_t\}$. However, in model (13), standard errors associated with both nonseasonal parameter estimates are greatly reduced now as they are given by 0.05914 and 0.05796, respectively. Still, the stationary complex root of (13) is assigned to the seasonal component since its associated frequency, $0.69\,\pi$, is sufficiently close to $2\pi/3$ under default settings. We therefore apply the same modification as for the IND case. Eventually, this complex root is captured by a transitory component additionally introduced and, hence, seasonal adjustment according to the AMB approach provides

Figure 6: Changes in seasonally adjusted figures compared with previous period.

results that are much more sensible than under default settings. As for IND, both WK seasonal adjustment filters display squared gains that look now much more like the ideal described in Section 3 even though it stays slightly below 0.9 behind the third seasonal frequency for the concurrent filter, see the bottom row of Figure 4. Nevertheless, a comparison with Figure 3 reveals that the huge advantage formerly gained by X-11 becomes negligible as differences between squared gains of X-11 and WK seasonal adjustment filters become significantly smaller. Dips at seasonal frequencies are even narrower for the latter. Consequently, seasonally adjusted figures obtained from SEATS are less smooth than under default settings. As illustrated exemplarily in Figure 5 for the span as of January 2003 up to December 2012, this effect remains visible throughout the whole period covered but is especially apparent in times of high volatility, such as the year 2003. As a result, seasonally adjusted CON figures are barely distinguishable when comparing X-11 and SEATS. Slight differences, however, may be recognised if their month-on-month changes are considered instead, see Figure 6 for the observation period's current end.

As for IND, revisions should thus be studied in addition to complement our spectral analysis. Since the regARIMA model has been changed for CON, they have to be (re)calculated for both approaches considered. Table 3 shows that this has virtually no effect on revisions obtained from the X-11 method. More importantly, it also uncovers that revisions tend to be generally smaller on average for the AMB approach. The only exception, however, is revisions' volatility of month-on-month changes. In sum, seasonal adjustment with X-11 may still be preferred by some users for CON. The AMB approach, however, seems to perform equally well when run with appropriate modifications of default settings.

6. Summary

We suggest a decision tree for choosing between the X-11 and SEATS seasonal adjustment core for any particular observed time series. Running X-13ARIMA-SEATS with default options for either core, the X-11 approach is preferred for turnover of industry and producers of consumer goods, while the AMB method is recommended for turnover of producers of capital goods. Based primarily on visual inspection of squared gains, we further demonstrate that SEATS, when run in default mode, may suffer from transitory effects being falsely assigned to the seasonal component when the regARIMA model fit to the observed time series does not belong to the class of airline models. Appropriate modifications of default settings may thus prevent SEATS from yielding results that are somewhat misleading. When employing such modifications, the AMB approach may also be selected for both turnover series seasonally adjusted with X-11 so far.

Future research should thus focus on further improvements of our decision tree. First of all, the recursive structure of our advanced empirical considerations could be elaborated especially with regard to the SEATS core. Similarly, the observed time series' length could be included more explicitly as it directly relates to the seasonal heteroskedasticity issue that is part of our first empirical considerations. In practice, more attention should be paid to short and moderate-length time series as the rather problematic $\log(311)(011)_{12}$ model found for two out of three turnover indicators studied here appears to be some sort of last resort for longer time series. Further criteria, such as time shifts of final seasonal adjustment filters, could be additionally incorporated into the decision-making process. Apart from our decision tree, future research may also focus on issues not covered here. For example, the SEATS algorithm automatically makes use of certain approximations to ARIMA models which do not admit an admissible decomposition to enforce admissibility. In this case, one may ask if these approximations still yield acceptable representations of the true DGP.

REFERENCES

Bell, W. R. (1984), "Signal Extraction for Nonstationary Time Series," *The Annals of Statistics*, 12, 646–664.

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

Bell, W. R., Holan, S. H., and McElroy, T. S. (eds.) (2012), *Economic Time Series – Modeling and Seasonality*, Boca Raton: CRC Press.

Burman, J. P. (1980), "Seasonal Adjustment by Signal Extraction," *Journal of the Royal Statistical Society*, Ser. A, 143, 321–337.

Cochran, W. G. (1941). "The Distribution of the Largest of a Set of Estimated Variances as a Fraction of Their Total," *Annals of Human Genetics*, 11, 47–52.

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

Ghysels, E., and Osborn, D. R. (2001), *The Econometric Analysis of Seasonal Time Series*, Cambridge, UK: Cambridge University Press.

Gómez, V., and Maravall, A. (1996), "Programs TRAMO and SEATS – Instructions for the User," *Working Paper No 9628*, Servicio de Estudios, Banco de España, Madrid.

Gómez, V., and Maravall, A. (2001), "Seasonal Adjustment and Signal Extraction in Economic Time Series," in *A Course in Time Series Analysis*, eds. D. Peña, G. C. Tiao and R. S. Tsay, New York: Wiley, pp. 202–247.

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

Hylleberg, S. (ed.) (1992), *Modelling Seasonality*, Oxford, UK: Oxford University Press.

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

Maravall, A. (1988), "A Note on Minimum Mean Squared Error Estimation of Signals with Unit Roots," *Journal of Economic Dynamics & Control*, 12, 589–593.

Maravall, A. (1995), "Unobserved Components in Economic Time Series," in *The Handbook of Applied Econometrics Vol. 1*, eds. M. H. Pesaran and M. R. Wickens, Oxford, UK: Blackwell Publishing, pp. 12–72.

Shiskin, J., Young, A. H., and Musgrave, J. C. (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.

Whittle, P. (1963), *Prediction and Regulation by Linear Least-Square Methods*, London: English University Press.