

Optimal Seasonal Filtering

JONATHAN H. WRIGHT

March 9, 2017

Abstract

In this paper, I compare model-based and moving average-based approaches to seasonal adjustment. I propose an optimal moving average filter, designed to minimize the distance from a model-based filter. In simulations, I consider the accuracy of the alternative methods in estimating seasonal factors in a simple model. I find that the most reliable results are obtained either by the model-based approach, or by the optimal moving average filter. I apply the proposed methods to seasonal adjustment of employment data.

1 Introduction

Seasonal adjustment is an enormously important element in the construction of macroeconomic data. In their raw form, before seasonal adjustment, employment and GDP data exhibit regular within-year fluctuations that are comparable in magnitude to the business cycle. Because these within-year fluctuations are short-lived and owe to factors such as climate and the timing of holidays, macroeconomic stabilization policy does not aim to smooth them out. The macroeconomic series that receive almost all of attention from policymakers, academics, financial markets and the press are seasonally adjusted, and aim to remove any regular within-year variation.

There are two broad methods for seasonal adjustment. One is based on estimating an explicit time series model. The dominant program for implementing this method is known as TRAMO-SEATS, and is used in European official statistics. It contains two components: TRAMO, which removes outliers and deterministic effects, and SEATS, which fits an ARIMA model and then extracts the implied seasonal factors. The other method, used in U.S. official statistics, is based on applying moving average filters without explicit estimation of a model. Several variants of this method have been operationalized over more than fifty years (Shishkin, 1957; Shishkin, Young, and Musgrave, 1967; Findley, Monsell, Bell, Otto, and Chen, 1998; Ladiray and Quenneville, 1989), and the current version is known as the X-13 filter, maintained by the U.S. Census Bureau. The current X-13 program incorporates both methods for seasonal adjustment in the sense that the SEATS seasonal adjustment is now available as an option within the program. However, US official statistics continue to be based on the filtering approach.

There is a strong argument for employing the model-based formulation. The degree of time variation in seasonal factors is explicitly estimated from the data. However, the filter-

based method has the advantage that it avoids the need to specify a particular parametric model. Moreover, it is the standard in U.S. official statistics, and incorporates many widely used diagnostics, and so it seems worthwhile to think about how to design that filter optimally. That is the goal of the current paper. A critical part of the X-13 process involves estimating seasonal factors using weighted moving averages of data in the same month or quarter of different years. There are 6 possible seasonal moving average schemes, and the default settings in the X-13 select among three of these based on the relative variability of seasonal and irregular components in a pilot initial decomposition of the series. The idea is that if the seasonal component is variable, then it should be estimated using a short window of data, whereas if the seasonal component is stable, then a longer window is to be preferred. While that is a reasonable principle, there is no formal rationale for the specific rule and cutoffs used in the default settings in the X-13 that I am aware of. In this paper, I propose picking the seasonal moving average filter from among the choices in the X-13, selecting the filter so as to minimize the distance between the resulting seasonal factors and those from a model-based procedure.

The plan for the remainder of this paper is as follows. In section 2, I describe the proposed method of optimally selecting the seasonal moving average filter. In section 3, I consider Monte-Carlo simulations in which alternative methods of seasonal adjustment are compared, and are assessed based on their proximity to the true seasonally adjusted data. In section 4, I apply alternative methods of seasonal adjustment to the nonfarm payrolls data in the Current Employment Statistics (CES) produced by the Bureau of Labor Statistics (BLS), which is perhaps the most widely watched macroeconomic data release, and for which seasonal effects are very important. Section 5 concludes.

2 Methodology

The X-13 program involves a number of steps. First, a regression ARIMA model of the form:

$$\phi(L)\Phi(L^S)(1-L)^d(1-L^{12})^D(y_t - \beta'x_t) = \theta(L)\Theta(L^S)\varepsilon_t \quad (2.1)$$

is fitted to the time series y_t , where x_t are user-chosen regressors, L denotes the lag operator, $\phi(\cdot)$, $\Phi(\cdot)$, $\theta(\cdot)$ and $\Theta(\cdot)$ are polynomials of orders p , P , q and Q respectively, d and D are integer difference operators, ε_t is an iid error term and S is the number of periods per year—12 for monthly data. This is referred to as a seasonal ARIMA(p,d,q)x(P,D,Q) model. The model is estimated by maximum likelihood, and used to forecast and backcast the time series. The seasonal adjustment is then applied to the series $y_t - \hat{\beta}'x_t$ where $\hat{\beta}$ is the maximum likelihood estimate of the coefficient β . The X-13 program gives the user the choice of seasonally adjusting this series by SEATS or by a sequence of moving average filters, known as the X-11, because this was the entirety of an earlier vintage of the program.

2.1 SEATS

With SEATS, the seasonal factors are recovered directly from the estimated seasonal ARIMA model, using the canonical decomposition of Hillmer and Tiao (1982). If the seasonal ARIMA model in equation (2.1) satisfies certain parameter restrictions, then it is observationally equivalent to a model in which the data are the sum of three orthogonal components:

1. A trend T_t such that $\phi(L)\Phi(L^S)(1-L)^{d+D}T_t = \theta_T(L)\varepsilon_{Tt}$,
2. A seasonal component, S_t , such that $(1 + L\cdots + L^{S-1})S_t = \theta_S(L)\varepsilon_{St}$ and,
3. An irregular white noise process $N_t = \varepsilon_{Nt}$,

where the time series ε_{Tt} , ε_{St} and ε_{Nt} are iid over time with mean zero and variances $\sigma_{\varepsilon_T}^2$, $\sigma_{\varepsilon_S}^2$ and $\sigma_{\varepsilon_N}^2$, and are mutually independent. If the parameter restrictions allow for such a

decomposition, the seasonal ARIMA model is said to be *admissible*. The decomposition into these three components is not unique, but if one maximizes the variance of the white noise component, then this yields a unique decomposition, called the canonical decomposition.

Given the parameters of the seasonal ARIMA model, one can work out the parameters in the lag polynomials $\theta_T(\cdot)$, $\theta_S(\cdot)$ and also the variances $\sigma_{\varepsilon_T}^2$, $\sigma_{\varepsilon_S}^2$ and $\sigma_{\varepsilon_N}^2$, by the algorithm of Hillmer and Tiao (1982)¹. Given the time series, $\{y_t\}_{t=1}^T$, the Kalman smoother then gives the implied seasonal component, S_t .

The SEATS algorithm takes the estimated seasonal ARIMA model parameters and computes the implied seasonal components from this canonical decomposition. If the estimated seasonal ARIMA model is not admissible, then it selects an alternative specification until one is found such that the estimated model is admissible.

2.2 X-11 moving average filters

With the X-11 moving average filters—the approach used in U.S. official statistics—the series that is to be decomposed, y_t , is modeled as either:

$$y_t = T_t + S_t + I_t \tag{2.2}$$

or:

$$y_t = T_t S_t I_t \tag{2.3}$$

where T_t , S_t and I_t are the trend/cycle, seasonal and irregular components. These are known as the additive and multiplicative decompositions, respectively. The trend/cycle is recovered from the data by applying a trend filter of the form $T_t = \sum_{i=-(m-1)/2}^{(m-1)/2} \theta_i y_{t+i}$ with symmetric filter weights:

$$\theta_i \propto [(p-1)^2 - i^2][p^2 - i^2][(p+1)^2 - i^2][3p^2 - 11i^2 - 16], \tag{2.4}$$

¹The canonical decomposition was implemented using the SSMMatlab toolbox written by Víctor Gómez.

where $p = \frac{m+3}{2}$ and the constant of proportionality ensures that the weights sum to unity². The seasonal factors are then estimated by taking weighted moving averages of $y_t - T_t$ or $\frac{y_t}{T_t}$ (for an additive or multiplicative decomposition) in the same period of different years. The seasonal moving average weighting scheme is known as a $3 \times n$ filter (for n an odd integer)³, and the weight given to data in the same period j years either before or after the current year is:

$$\begin{aligned}
 w_j &= \frac{1}{3}, \quad j = 0, 1, 2, & n = 1 & \quad (2.5) \\
 w_j &= \frac{\max(\frac{3+n}{2} - j, 3)}{3n}, \quad j = 0, \dots, \frac{3+n}{2} - 1, & n \neq 1 &
 \end{aligned}$$

The X-11 process allows the seasonal moving average filter to be a 3×1 , 3×3 , 3×5 , 3×9 , 3×15 or the stable filter, that is instead simply the equal-weighted average over the sample period in all years of the sample⁴. The algorithm involves running iterations of the trend filter in equation (2.4) and the seasonal moving average filter in equation (2.5), and it includes other adjustments, for example for outliers and the number of trading days in a month, and also the use of asymmetric counterparts of equations (2.4) and (2.5) at the start and end of the sample. The algorithm is thus quite involved, and is described in considerable detail in Ladiray and Quenneville (1989). But the crucial parameters are m (the number of terms in equation (2.4)) and the choice of n in the $3 \times n$ seasonal moving average filter in equation (2.5). The default settings in the X-11 process within the X-13 program select among $m = 9$ and $m = 13$ alone, based on the relative variability of T_t and I_t and select among the 3×3 ,

²This trend filter is known as a Henderson filter (Henderson (1916)). The filter minimizes the variance of the third differences of the input series. Bell and Monsell (1992) plot the gain functions for alternative Henderson filters—the larger is m , the more high-frequency cycles are eliminated by the filter.

³The origin of the terminology is that the weights can be thought of as 3 period moving averages of n period moving averages.

⁴The 3×15 filter is only available with a span of at least 20 years of data—otherwise it is defined to be equivalent to the stable filter.

3x5 and 3x9 filter alone, based on the relative variability of S_t and I_t .⁵

2.3 Proposed optimal filters

In this paper, I consider seasonally adjusting the data by the model-based SEATS seasonal adjustment, and using all 6 seasonal moving average filters in the X-11 (with the default trend filter). Let $S_M(t)$ denote the model-based estimate of the seasonal factor and let $S_{SMA,i}(t)$ denote the seasonal factor at time t estimated from the i th of the 6 seasonal moving average filters. Among these alternatives, I propose selecting the optimal seasonal filter as:

$$i_1^* = \arg \min_i T^{-1} \sum_{t=1}^T (S_{SMA,i}(t) - S_M(t))^2 \quad (2.6)$$

The data can then be seasonally adjusted by the X-11 filter using this choice of the seasonal moving average. A variant on this is to consider seasonally adjusting the data by different permutations of the trend filter *and* the seasonal moving average filter. I take 6 possible values of m in equation (2.4)— $m=\{7,9,13,17,23,33\}$. In conjunction with the seasonal moving averages, this gives a total of 36 different seasonal adjustment procedures. Let $S_{TSMA,i}(t)$ denote the seasonal factor at time t estimated from the i th of these 36 methods. In this approach, I propose selecting the optimal seasonal filter as:

$$i_2^* = \arg \min_i T^{-1} \sum_{t=1}^T (S_{TSMA,i}(t) - S_M(t))^2 \quad (2.7)$$

I refer to these two alternatives as the proposed optimal seasonal moving average (SMA)

⁵Specifically, the I/C ratio (the average absolute irregular component divided by the average absolute trend-cycle) and the I/S ratio (the average absolute irregular component divided by the average absolute seasonal component) are both estimated in an initial seasonal adjustment pass. If the I/C ratio is below 1, then $p = 9$, otherwise $p = 13$. If the I/S ratio is below 2.5, between 3.5 and 5.5, or above 6.5, then the 3x3, 3x5 or 3x9 filter is used, respectively. If it does not fall into any of these three regions, then the last year of data is deleted and the procedure is re-run. The algorithm is iterated until one of the three filters is selected or five years' data have been dropped, whichever comes sooner. If in the end, no filter has been selected, the 3x5 filter is employed.

and optimal SMA and trend filters, respectively⁶.

3 Monte-Carlo Simulations

In this section, I report the results of a Monte-Carlo simulation. Artificial data are simulated from the model with pseudo-monthly data:

$$(1 - L)(1 - L^{12})y_t = (1 + \theta L)(1 + \Theta L^{12})\varepsilon_t \quad (3.1)$$

where ε_t is iid $N(0,1)$ noise and the sample size is $T = 120$. This is the “airline” model of Box and Jenkins (1986), and is a seasonal ARIMA(0,1,1)x(0,1,1) specification. As shown in Hillmer and Tiao (1982), if we assume that $\Theta < 0$ and $-1 < \theta < 1$ then this model is admissible. I take the true parameter values and work out the parameters in the lag polynomials $\theta_T(\cdot)$, $\theta_S(\cdot)$ and also the variances $\sigma_{\varepsilon_T}^2$, $\sigma_{\varepsilon_S}^2$ and $\sigma_{\varepsilon_N}^2$, by the Hillmer-Tiao algorithm. Given a draw of the time series, $\{y_t\}_{t=1}^T$, the Kalman smoother then gives the implied seasonal component, S_t . I interpret this as the *true* seasonal component, though no researcher can extract it in practice because it depends on the unknown parameters θ and Θ .

The researcher can, however, apply any of the seasonal filters⁷ to the observed time series $\{y_t\}_{t=1}^T$, giving an estimated seasonal factor, \hat{S}_t . For each of these estimated seasonal factors, I compute the root mean square deviation from the true seasonal component, obtained from the canonical decomposition based on the true parameters:

$$RMSE = \sqrt{T^{-1} \sum_{t=1}^T (\hat{S}_t - S_t)^2} \quad (3.2)$$

⁶This is assuming that the X-13 is using an additive decomposition (equation (2.2)). If it is instead multiplicative (equation (2.3)), the corresponding parametric model should be in logs, and then the seasonal factors in equations (2.6) and (2.7) should be logged.

⁷The X-13 was implemented throughout this paper using the X-13 toolbox for Matlab written by Yvan Lengwiler.

I then average these root mean square errors over 1,000 Monte Carlo replications. The seasonal filters that are considered are SEATS, the X-11 default, and the two optimal moving average filters proposed in this paper⁸. I also consider using the X-11 default trend filter in conjunction with fixing the seasonal moving average filter at the 3x1, 3x3, 3x5, 3x9 or stable alternatives.

The X-13 seasonal adjustment starts with by fitting a seasonal ARIMA model regardless of whether the subsequent seasonal adjustment is done by X-11 or SEATS. Table 1 reports the average RMSE (equation (3.2)) from the simulations where the seasonal ARIMA model is an ARIMA(0,1,1)x(0,1,1) specification, and considering an additive X-11 seasonal adjustment. SEATS gives much smaller RMSE than the X-11 default in all cases considered. For example, with $\theta = -0.3$ and $\Theta = -0.9$, the average RMSE for SEATS is 0.06 whereas for the X-11 default it is 0.23. The X-11 with the optimal SMA or the optimal SMA and trend both give lower RMSE than the X-11 default in almost all cases, with the improvement being quite big in some cases. For example $\theta = -0.3$ and $\Theta = -0.9$, the average RMSE for the optimal SMA and optimal SMA and trend filters are 0.13 and 0.11, respectively. However, the optimal filters never get as low a RMSE as SEATS, with the settings in Table 1. The optimal SMA and trend filter (equation (2.7)) gives lower RMSE than the optimal SMA filter (equation (2.6)), but the differences are small.

Within the possible fixed seasonal moving averages in the X-11, a short window is optimal if Θ is close to zero, whereas a longer window is optimal if Θ is more negative, which corresponds to a relatively stable seasonal component, and it is intuitive that a stable seasonal should call for a long window. But of course the researcher does not know the true data generating process in practice and so will not necessarily know which filter to pick. The

⁸With a sample size of 120, the 3x15 and stable filters are identical and so the former are not considered in this simulation.

appeal of SEATS or the optimal filters is that they adapt automatically to the data generating process.

Table 2 repeats the results from Table 1, except where the seasonal ARIMA model is selected by the automatic model selection procedure of Gómez and Maravall (1996, 2013). The RMSEs are little changed for the X-11 default or any of the fixed seasonal moving average filters. The RMSEs for SEATS are larger in this case than in Table 1. This is not surprising, since the data generating process is an $ARIMA(0,1,1) \times (0,1,1)$, so using the automatic model selection method introduces the possibility of misspecification, and SEATS is directly based on the chosen seasonal ARIMA model. The increase in the RMSE is greatest when $\Theta = -0.9$. Nonetheless, the RMSE of SEATS is again smaller than that of the X-11 default in all cases considered in this table. If $\Theta = -0.9$, the optimal seasonal filters actually give smaller RMSEs than either SEATS or the X-11 default.

I am not aware of any existing papers that simulate data from ARIMA models, compute the canonical decomposition and then assess different seasonal filters based on their proximity to the seasonal factor from that decomposition, nor am I aware of any work evaluating the optimal moving average filters proposed in subsection 2.3. But there are several existing studies comparing the empirical properties of different seasonal filters in related but somewhat different ways. Hood, Ashley, and Findley (2000) conducted Monte-Carlo simulations where they started from actual macroeconomic seasonally adjusted data, added in pseudo-random irregular components, and then assessed the quality of competing seasonal filters in terms of their ability to recover the original seasonally adjusted data. When the irregular component had high variance, they found that SEATS did better than the moving average filters, which is consistent with my findings. Hood and Findley (1999), Hood (2002) and Tiller, Chow, and Scott (2007) use diagnostics to compare the quality of adjustments

from SEATS and moving average filters. Depoutot and Planas (1998) compute the filter weights from the canonical decomposition for different choices of the true parameters θ and Θ , and then provide a look-up table telling us which X-11 filter from among the 3x3, 3x5, 3x9 and 3x15 alternatives gives the minimum mean square difference between the X-11 and canonical decomposition filter weights. When Θ is close to -1, the 3x15 filter is selected. This exercise does not involve any estimation, but again points to the intuitive idea that when the seasonal component is known to be stable relative to the irregular, a long filter is to be preferred, as was found in Table 1. Tiller, Chow, and Scott (2007) and Wright (2013) noted that parameter estimates that are obtained from fitting an airline model to macroeconomic data would, in conjunction with the look-up table of Depoutot and Planas (1998), imply that a 3x9 or 3x15 filter is often optimal, even though the 3x9 filter is seldom chosen by the X-13 default and the 3x15 is not even an option in the X-13 default settings.

4 Application to Nonfarm Payrolls Data

In this section, I consider the seasonal adjustment of nonfarm payrolls data from the BLS current employment statistics (CES) survey (the “establishment” survey) that is the most widely-followed monthly economic indicator. Seasonal adjustment for these series is conducted *indirectly* meaning that 150 disaggregates are seasonally adjusted separately, and then the seasonally adjusted data are aggregated to yield seasonally adjusted nonfarm payrolls data. The BLS makes available all the specification files to enable this process to be replicated⁹, and I implemented this replication. I then redid the seasonal adjustment keeping all the specifications unchanged, except for (i) doing the seasonal adjustment using SEATS rather than a moving average filter and (ii) doing the seasonal adjustment using the proposed

⁹BLS only publishes rounded unadjusted data and so the replication cannot be perfect.

optimal filters. The seasonal adjustment as implemented by BLS uses a ten-year rolling window of data. As a consequence, the 3x15 and stable filters are identical and so the former are not considered in the optimal filters for this exercise.

Table 3 reports the number of disaggregate series for which the optimal seasonal filter in equation (2.6) picked each of the possible alternative seasonal moving average filters. It also reports the number of disaggregate series for which the optimal seasonal filter in equation (2.7) picked each seasonal moving average and trend filter. For 60% of the series, a long moving average was selected—either the 3x9 or the 3x15. This is consistent with Wright (2013) who used other approaches to argue for a long seasonal moving average for many series. At the same time, there are some series for which quite short moving averages are found to be optimal. For example, the 3x1 filter is found to be optimal for the employment of couriers and messengers (such as Fedex and UPS) where the rise of online holiday shopping causes a seasonal spurt in employment in December that is soared in the last five years. But very short filters are also selected for other series. Interestingly, for 5 of the 6 categories of government employment, the very short 3x1 filter is selected. Table 3 also shows the filters chosen by the default X-11—for 111 out of the 150 series, this is either the 3x5, and for another 38 series it is the 3x3. The optimal seasonal filter that selects the trend filter in equation (2.4) tends to pick a smaller value of m than in the X-11 default.

Figure 1 plots the difference between aggregate employment (summing over the 150 categories) using SEATS and using the default X-11 filter. The differences can be over 100,000 in either direction. In the years from 2007 to 2012, SEATS tends to give lower employment levels in the winter and early spring months and higher employment levels in the rest of the year, with the pattern being strongest closest to the Great Recession. The interpretation is that job losses during the Great Recession were most severe in winter and

early spring. This caused the seasonal factors in the X-11 to move a great deal because of the small window size that is chosen, as discussed at some length in Wright (2013) and papers cited therein. The seasonal factors in SEATS were less affected. Consequently seasonally adjusted data using SEATS are lower in the winter and early spring and higher the rest of the year. From 2012 to the end of the sample, differences between SEATS and X-13 seasonally adjusted data are smaller, but the pattern flips—SEATS tends to give higher employment levels in winter and early spring. SEATS, with its longer effective filter weights, was less influenced by the Great Recession than the X-11 default, but the flip-side of this is that the effect lasted for longer.

Figures 2 and 3 repeat the same exercise, except giving the difference between the data as seasonally filtered with the two optimal filters, less the data using the default trend and seasonal moving average X-11 filters, as implemented by the BLS. The optimal filters differ from their default counterparts in the same direction as SEATS, which is not surprising given that they are constructed to approximate SEATS. However, the default X-11 filter is generally closer to the optimal filters than to the SEATS filter, and so the differences reported in Figures 2 and 3 are of somewhat smaller magnitude than those shown in Figure 1.

One could argue that the differences between alternative seasonally adjusted data are not usually that big. The BLS estimates the sampling standard error in the monthly level of employment to be about 56,000. However, the Fed, financial markets and the press analyze employment numbers with great—and perhaps excessive—precision, and the differences that I find are clearly ones that would get the attention of analysts and market participants.

The seasonal adjustment implemented by the BLS uses only data from 2007:01 to 2017:01. Especially once one allows for the possibility of using a 3x15 or stable filter, a longer rolling

window of data may give quite different results. I accordingly redid the CES seasonal adjustment but using data from 1990:01 to 2017:01—a span of nearly 30 years¹⁰. Note that the length of the sample permits me to consider the 3x15 filter separately from the stable filter, which was not possible in the previous shorter sample.

Table 4 reports the number of disaggregate series for which the optimal seasonal filters and default X-11 picked each of the possible seasonal moving average and trend filters. Again, the default X-11 usually selects the 3x5 filter, whereas the optimal filters are longer for the majority of the series.

Figures 4-6 give the counterparts of Figures 1-3 but using the longer sample period. The SEATS and optimal filters give different seasonally adjusted data from the X-11 default. But the differences are smaller in absolute magnitude in the years 2007-2010, probably because this longer sample period is less influenced by the extreme shock of the Great Recession. Still one does see other examples of how the different seasonal adjustment methods respond to shocks, such as the timing of business cycle downturns¹¹. For example, the worst of the 1990-1991 recession came in the winter of 1990-1991. The seasonally adjusted employment data using SEATS showed a lower level of employment in that winter and for the next three winters because the seasonal factor was more stable than using the X-11 default.

Since any of the methods for seasonal adjustment allow for seasonal factors to vary over time (with the exception of the stable filter), it is hard to see a rationale for completely discarding all data outside a ten year rolling window.

¹⁰For many disaggregates, the specification files used by the BLS contain additive outliers—I incorporate these outliers for the period since 2007, but the specification files do not give outliers from before 2007 and so no additive outliers can be included for these earlier years. Needless to say this applies equally to the X-11 default, SEATS and the proposed optimal seasonal adjustment.

¹¹Canova and Ghysels (1994) raised the concern about the timing of cyclical fluctuations distorting seasonal factors.

5 Conclusions

In this paper, I have compared model-based and moving average-based approaches to seasonal adjustment and proposed optimal moving average filters, designed to minimize the distance from a model-based filter. I assess different filters by their ability to recover the true seasonal factors in Monte-Carlo simulations with a known data generating process. The best results are obtained either by the model-based approach, or by the optimal moving average filters. The current default setting of the X-11 does materially worse for some parameter values and does not do materially better for any parameter values in the simulation that I consider. I compare the model-based and optimal moving average filters with the current default settings of the X-11 in the problem of seasonally adjusting employment data, and find that the model-based and optimal filters tend to use longer windows to estimate seasonal factors. It can make a material difference to the aggregate seasonally adjusted payroll employment. My substantive recommendations would be to use SEATS or either of the proposed optimal filters for seasonally adjusting employment data. In any case, a longer rolling window of data should be used no matter what seasonal adjustment technology is chosen.

Table 1: Root Mean Square Errors of Alternative Seasonal Factor Estimates

θ	-0.3	-0.3	-0.3	-0.6	-0.6	-0.6	-0.9	-0.9	-0.9
Θ	-0.3	-0.6	-0.9	-0.3	-0.6	-0.9	-0.3	-0.6	-0.9
SEATS	0.07	0.07	0.06	0.05	0.06	0.05	0.06	0.07	0.05
X11-Default	0.31	0.17	0.23	0.29	0.17	0.24	0.34	0.19	0.27
Optimal SMA	0.22	0.18	0.13	0.21	0.18	0.13	0.24	0.20	0.14
Optimal SMA and trend	0.20	0.16	0.11	0.19	0.16	0.11	0.23	0.19	0.13
<i>Fixed Seasonal Moving Average Filters</i>									
3x1	0.22	0.23	0.35	0.20	0.24	0.36	0.24	0.27	0.41
3x3	0.25	0.19	0.28	0.23	0.19	0.30	0.27	0.22	0.34
3x5	0.32	0.17	0.22	0.29	0.17	0.24	0.34	0.19	0.27
3x9	0.44	0.20	0.15	0.41	0.19	0.17	0.46	0.21	0.19
Stable	0.70	0.34	0.11	0.67	0.33	0.11	0.75	0.37	0.13

Notes: This table reports the simulated root mean square errors in equation (3.2) when simulating data from the model $(1-L)(1-L^{12})y_t = (1-\theta L)(1-\Theta L^{12})\varepsilon_t$ where ε_t is standard normal, the sample size is $T = 120$. An additive seasonal decomposition is conducted with an ARIMA(0,1,1)x(0,1,1) specification, using SEATS, the default X-11 and the two proposed optimal X-11 filters. Results with the X-11 using the default trend filter and using the 3x1, 3x3, 3x5, 3x9 and stable filters are also included. All results are averaged over 1,000 Monte Carlo replications.

**Table 2: Root Mean Square Errors of Alternative Seasonal Factor Estimates:
Automatic Model Selection**

θ	-0.3	-0.3	-0.3	-0.6	-0.6	-0.6	-0.9	-0.9	-0.9
Θ	-0.3	-0.6	-0.9	-0.3	-0.6	-0.9	-0.3	-0.6	-0.9
SEATS	0.09	0.09	0.21	0.07	0.08	0.20	0.09	0.11	0.21
X11-Default	0.31	0.17	0.23	0.29	0.17	0.24	0.34	0.19	0.27
Optimal SMA	0.22	0.18	0.14	0.21	0.18	0.15	0.25	0.21	0.17
Optimal SMA and trend	0.20	0.17	0.14	0.19	0.17	0.14	0.23	0.20	0.17
<i>Fixed Seasonal Moving Average Filters</i>									
3x1	0.22	0.23	0.35	0.20	0.24	0.36	0.24	0.27	0.41
3x3	0.25	0.19	0.28	0.23	0.19	0.30	0.27	0.22	0.34
3x5	0.32	0.17	0.22	0.29	0.17	0.24	0.34	0.19	0.27
3x9	0.44	0.20	0.16	0.41	0.19	0.17	0.46	0.21	0.19
Stable	0.70	0.34	0.13	0.67	0.33	0.12	0.76	0.37	0.14

Notes: As for Table 1, except that the ARIMA model is chosen by the automatic model selection procedure of Gómez and Maravall (1996, 2013).

Table 3: Number of series for which optimal filter selected each filter option

		Selecting Optimal SMA	Selecting Optimal SMA and Trend	X-11 Default
SMA	3x1	21	20	
	3x3	14	16	38
	3x5	24	24	111
	3x9	46	46	1
	Stable	45	44	
Trend	$m = 7$		78	
	$m = 9$		26	128
	$m = 13$		15	22
	$m = 17$		12	
	$m = 23$		9	
	$m = 33$		10	

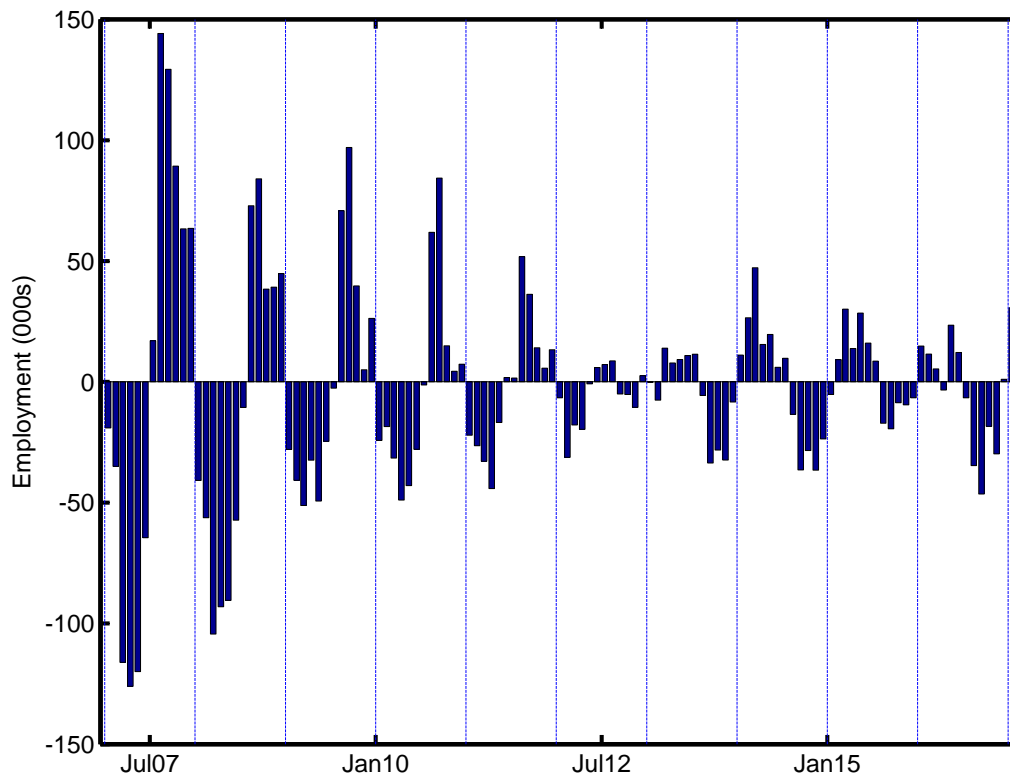
Notes: For the 150 CES disaggregates, this table reports the number of disaggregates for which the optimal filter in equation (2.6) picked each of the possible alternative seasonal moving average filters (in the Selecting Optimal SMA column). It also reports the number of disaggregates for which the optimal seasonal filter in equation (2.7) picked each of the possible seasonal moving average and trend filters (in the selecting optimal SMA and trend column). Results for the X-11 default choices are shown in the right-most column. For the default seasonal moving average, the possible choices are 3x3, 3x5 and 3x9 alone. For the default trend, the possible choices are 9 and 13 alone.

**Table 4: Number of series for which optimal filter selected each filter option.
Sample: January 1990-January 2017.**

		Selecting Optimal SMA	Selecting Optimal SMA and Trend	X-11 Default
SMA	3x1	12	13	
	3x3	11	17	60
	3x5	36	27	90
	3x9	40	44	0
	3x15	34	31	
	Stable	17	18	
Trend	$m = 7$		104	
	$m = 9$		24	137
	$m = 13$		7	13
	$m = 17$		5	
	$m = 23$		5	
	$m = 33$		5	


Notes: As for Table 3, except that the sample for seasonal adjustment is January 1990-January 2017 instead of the January 2007-January 2017 window used by the BLS.

Figure 1: Aggregate seasonally adjusted employment:
Difference between SEATS and default X-11.




NOTE: This figure plots the difference between aggregate monthly employment, summing over the 150 disaggregates when the seasonal adjustment is done using SEATS and using the default X-11 filter in the X-13 process (SEATS seasonally adjusted data less default X-11 seasonally adjusted data). Vertical lines represent January of each year.

Figure 2: Aggregate seasonally adjusted employment:
Difference between optimal SMA filter and default X-11.




NOTE: This figure plots the difference between aggregate monthly employment, summing over the 150 disaggregates when the seasonal adjustment is done using the proposed optimal SMA filter and using the default X-11 filter in the X-13 process (optimal seasonally adjusted data less default X-11 seasonally adjusted data). Vertical lines represent January of each year.

**Figure 3: Aggregate seasonally adjusted employment:
Difference between optimal SMA and trend filter and default X-11.**



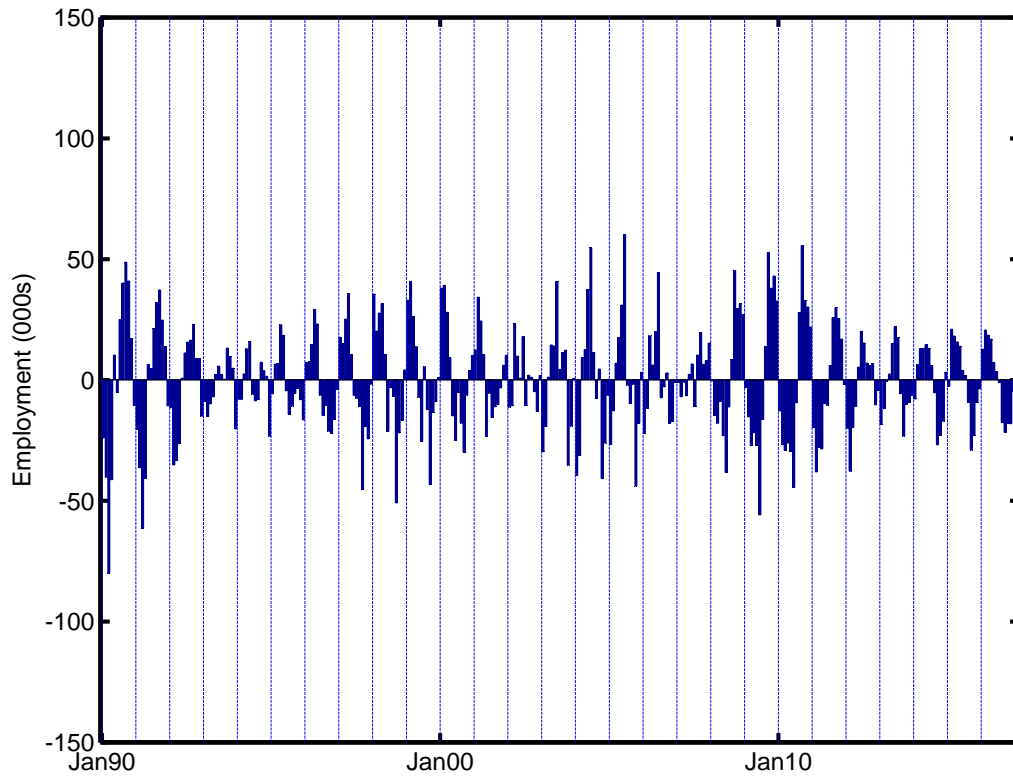
NOTE: This figure plots the difference between aggregate monthly employment, summing over the 150 disaggregates when the seasonal adjustment is done using the proposed optimal SMA and trend filter and using the X-11 filter in the X-13 process (optimal seasonally adjusted data less default X-11 seasonally adjusted data). Vertical lines represent January of each year.

**Figure 4: Aggregate seasonally adjusted employment:
Difference between SEATS and default X-11.
Sample: January 1990-January 2017.**




NOTES: As for Figure 1, except that the sample for seasonal adjustment is January 1990-January 2017 instead of the January 2007-January 2017 window used by the BLS.

**Figure 5: Aggregate seasonally adjusted employment:
Difference between optimal SMA filter and default X-11.
Sample: January 1990-January 2017.**



NOTES: As for Figure 2, except that the sample for seasonal adjustment is January 1990-January 2017 instead of the January 2007-January 2017 window used by the BLS.

**Figure 6: Aggregate seasonally adjusted employment:
Difference between optimal SMA and trend filter and default X-11.
Sample: January 1990-January 2017.**



NOTES: As for Figure 3, except that the sample for seasonal adjustment is January 1990-January 2017 instead of the January 2007-January 2017 window used by the BLS.

References

- BELL, W. R., AND B. C. MONSELL (1992): “X-11 symmetric linear filters and their transfer functions,” SRD Research Report No. RR-92/15, U.S. Census Bureau.
- BOX, G. E. P., AND G. M. JENKINS (1986): *Time Series Analysis: Forecasting and Control (revised edition)*. Holden-Day.
- CANOVA, F., AND E. GHYSELS (1994): “Changes in Seasonal Patterns: Are they Cyclical?,” *Journal of Economic Dynamics and Control*, 18, 1143–1171.
- DEPOUTOT, R., AND C. PLANAS (1998): “Comparing seasonal adjustment extraction filters with application to a model-based selection of X-11 linear filters,” Eurostat Working Paper.
- FINDLEY, D. F., B. C. MONSELL, W. R. BELL, M. C. OTTO, AND B.-C. CHEN (1998): “New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program,” *Journal of Business and Economic Statistics*, 16, 127–152.
- GÓMEZ, V., AND A. MARAVALL (1996): “Programs TRAMO and SEATS. Instructions for the User,” Working Paper 9628, Servicio de Estudios, Banco de España.
- (2013): “Automatic modeling methods for univariate time series,” in *A Course in Time Series Analysis*, ed. by D. Peña, G. C. Tiao, and R. S. Tsay. J. Wiley and Sons.
- HILLMER, S. C., AND G. C. TIAO (1982): “An ARIMA-model-based approach to seasonal adjustment,” *Journal of the American Statistical Association*, 77, 63–70.
- HOOD, C. E. (2002): “Comparison of time series characteristics for seasonal adjustments from SEATS and X12-ARIMA,” *American Statistical Association Proceedings of the Business Economics Section*.
- HOOD, C. E., J. D. ASHLEY, AND D. F. FINDLEY (2000): “An empirical evaluation of the performance of TRAMO-SEATS on simulated series,” *American Statistical Association Proceedings of the Business Economics Section*.
- HOOD, C. E., AND D. F. FINDLEY (1999): “An evaluation of TRAMO/SEATS and comparison with X12-ARIMA,” *American Statistical Association Proceedings of the Business Economics Section*.
- LADIRAY, D., AND B. QUENNEVILLE (1989): *Seasonal Adjustment with the X-11 Method*. Springer.
- SHISHKIN, J. (1957): “Electronic Computers and Business Indicators,” *Journal of Business*, 30, 219–267.

SHISHKIN, J., A. H. YOUNG, AND J. C. MUSGRAVE (1967): “The X-11 variant of the Census method II Seasonal Adjustment Program,” Technical Paper 15, Bureau of the Census.

TILLER, R. T., D. CHOW, AND S. SCOTT (2007): “Empirical Evaluation of X-11 and Model-Based Seasonal Adjustment Methods,” Working Paper, Bureau of Labor Statistics.

WRIGHT, J. H. (2013): “Unseasonal Seasonals?,” *Brookings Papers on Economic Activity*, 2, 65–110.