

Adapting the Seasonal Adjustment of Local Area Unemployment Statistics to the COVID-19 Pandemic

Richard Tiller, Jennifer Oh, Lizhi Liu

Bureau of Labor Statistics, 2 Massachusetts Ave. NE, Washington, DC 20212

Abstract

The Covid-19 pandemic delivered an instantaneous shock to the U.S. labor market in March/April 2020. This crisis presents a challenge to seasonal adjustment of labor force data. In this paper we explore various options for seasonally adjusting series during the pandemic using as examples 421 series from the Bureau of Labor Statistics Local Area Unemployment Statistics program. The basic issue is how to prevent distortions in seasonal factor estimation from outliers generated by the pandemic while efficiently using new information generated during the pandemic period. This task is complicated because at the onset of a pandemic there is little data available to estimate its duration and dynamics. Since a large number of series, must be seasonally adjusted, an automated approach is necessary. We explored several options in terms of the sequence and mix of outlier types allowed in the automated modeling process and used information criteria to select the parsimonious model.

Key Words: Automatic Outlier adjustment, AICC, Parsimonious models

1. Introduction

The Local Area Unemployment Statistics (LAUS) program in the Bureau of Labor Statistics (BLS) produces and seasonally adjusts monthly 1,000 labor force series comprising states and metro areas using X-13A-S (U.S. Census Bureau). The normal policy is to use concurrent seasonal adjustment where all relevant data is used as it becomes available. Both theoretical and numerous empirical studies have shown that the revisions of current seasonally adjusted values due to filter changes can be reduced substantially if concurrent seasonal factors are used instead of year-ahead forecasted seasonal factors. In the LAUS program several restrictions are imposed. The ARIMA model form, parameters, and outliers are fixed during the current year and the concurrent seasonal adjustment filter is used for the entire year. Revisions to the ARIMA model and filter are made at the end of the year.

The Covid-19 pandemic delivered an instantaneous shock of unparalleled magnitude and scope to the labor market which required immediate modifications to our normal policy and even raised the issue as to the efficacy of concurrent adjustment during the early part of the pandemic period. We considered two options; use forecasted seasonal factors based on pre-pandemic data; continue with concurrent adjustment accompanied with real time outlier detection and correction. We decided to use the second option despite the pandemic disruptions since it does not arbitrarily discard new information from the pandemic period while still being capable of preventing major distortions to the estimated seasonal factors.

2. Outlier Modeling and Selection

While there are many approaches to modeling time series outliers, the one used most often by statistical agencies is based on the Box-Tiao intervention model that represents outliers as deterministic deviations from normal behavior as described by a non-stationary stochastic time series model. In X-13A-S the outlier effects are estimated by fitting a fixed coefficient regression model with ARIMA time series errors (RegArima). The regressors are appropriately coded dummy variables where the outlier effects are estimated from the coefficients and removed from the data before seasonal adjustment.

2.1 Type of Outliers

X-13A-S provides several pre-specified outliers listed in Table 1. The first three are the most relevant for this study. A single outlier can be used in combination with other types as well as in a sequence with the same type to model complex behavior. The last column of Table 1 shows how X-13A-S allocates outliers to the various components of the time series decomposition. The AO and TC are classified as part of the irregular independent of the trend-cycle. This classification may be plausible if either the AO or TC is viewed as an isolated effect, but this interpretation is questionable during the pandemic period when they are combined with LS's to approximate the deviation of the trend-cycle from normal.

Type	Description	Component
Additive (AO)	Abrupt change affects only one observation	Irregular
Level Shift (LS)	Abrupt change to a new level	Trend-cycle
Temporary change (TC)	Abrupt change followed by exponential decay to normal	Irregular
Temporary Level Shift	Sequence of level shifts summing to zero	Trend-cycle
Ramp (linear or quadratic)	Constant or quadratic change to new level	Trend-cycle
Seasonal Outlier	Abrupt change in seasonality	Seasonal

The goal is to find a parsimonious fit that avoids under or over fitting. At the peak of the pandemic underfitting is the greater risk since it may result in over smoothing shifts in level and inflating seasonal factors that propagate pandemic shifts into future years as seasonal effects. During the recovery period overfitting may become a serious risk especially with large deterministic LS's, based on a few extreme observations, that may now be too inflexible.

2.2 Automated Selection Procedures

The classic intervention model approach assumes we know the start, end dates and type of all the outliers that occurred in the series over a prespecified interval. Rarely is all this information known a priori. To fill in this required information an outlier methodology is implemented in the outlier spec of X-13A-S. For the pandemic we do know the approximate time of its occurrence which narrows down the number of observations to search through, but the response patterns are unknown and only unfold gradually as additional data accumulate. The problem is how to choose the best approximating mix of outliers for a large number of series with only a limited amount of data from the pandemic period. Two types of automated procedures are available in X-13-A-S (U.S. Census Bureau, 2020) which can be targeted specifically to the pandemic period.

The Iterative stepwise approach searches over candidate regression variables consisting of up to 3 types of outliers (AO, TC, LS) at each time point within the designated time interval. Since the total number of

regressors to be tested is up to 3 times the number of observations searched this process leads to a downward bias in the nominal alpha levels (the probability of falsely identifying outlier effects). To reduce this bias, t-test critical values are increased with the number of observations tested (X-13A-S default, based on Ljung, 1993).

The Iterative stepwise approach begins with a forward step which adds one outlier at a time with the most significant t-stat and repeats until no significant t-stats remain. Pruning of variables that have become insignificant during the forward addition is implemented through a backward deletion of 1 outlier at a time until only significant outliers remain. This intensive search is effective for identifying a mix of different outlier types towards the middle of series, but less useful towards the end of a series where data on the outlier effect is sparse.

The second type of automated procedure is the sequence of successive outliers of the same type. Introduced by Lytras and Bell, this approach is implemented in the regression spec of X-13A-S. It adds a sequence of outliers of the same type over a specified time interval to the regression. Either a sequence of additive outliers (AOS) or level shifts (LSS) may be selected. Since it requires much less intensive testing, it is useful towards the end of the series.

If t_o marks the start of the outlier effect and r the number of periods since the occurrence of the first outlier, the sequence of potential outliers is given by,

$$O_{t_o+r} = \sum_{i=0}^r \beta_i x_{i,t_o+i} = \begin{cases} \beta_r & \text{if AOS} \\ \sum_{i=0}^r \beta_i & \text{if LSS} \end{cases}$$

where the $x_{i,t}$ are the appropriate zero-one coded regression variables. At each time point the rule is to test the entire sequence of regressors and retain those with significant t-values.

The AOS generates non-overlapping effects that do not accumulate over time. At time t_o+r the effect is given by the estimate of β_r if it is significant; otherwise, no outlier effect is in the model at that time point. Note removal of an AO outlier effect from a series is equivalent to substitution of the ARIMA predicted value for the original observation. If all regressors in the test period are significant, AOS is equivalent to the factor projection method.

In contrast LSS represents overlapping effects that accumulate over time. For LSS, even though there may not be a new LS at the end of the series, there is still the possibility of an outlier effect that persists from previous periods if at least one of the previously estimated coefficients has a significant t-value.

2.3 Minimum AICC criterion for Outlier Model Selection

Given different options available for determining the outlier mix, we propose to use the minimum AICC criterion to select the most parsimonious one as described below. The usefulness of this approach depends on having estimated the parameters of the ARIMA model and outlier effects prior to the known occurrence of the pandemic. This is discussed below in the computation of the AICC as shown below,

$$AICC_N = -2L_N + 2n_p + \frac{2n_p(n_p + 1)}{N - n_p - 1}$$

where,

L_N = estimated maximum value of the exact log likelihood function

n = number of estimated parameters in the model
 p

N = effective number of observations.

The first term is the “information loss” from fitting the RegArima model, where the ARIMA part and historic outliers from the pre-pandemic period are fixed across all pandemic outlier options to allow for a focus on just the pandemic outlier specification. The second term is the penalty for complexity which includes the additional parameters required in the outlier specification during the pandemic period. The third term is a small sample correction. The minimum AICC picks the most parsimonious outlier combination in the sense that it balances minimizing loss in fit with increased model size due to the pandemic outlier specification.

Note this is the reverse of the approach for using the AICC criteria to compare ARIMA models described in the X-13A-S user guide, where it is recommended the outlier specification be held fixed across models (U.S. Census Bureau, 5.51). This may be misinterpreted by users as a warning against using AICC for outlier model selection.

3. Empirical Examples

We test alternative options for automatic outlier selection for unemployment series in 421 metro areas. The test period covers the time from March to December of 2020.

A useful overview of the behavior of unemployment is provided by the plot in Figure 1 of the seasonally adjusted (SA) and not seasonally adjusted (NSA) U.S. rate from the Current Population Survey (CPS). Figure 1 shows a strong pandemic effect in April 2020 when the rate rose from below 4% to over 15%. This eclipsed the previous peak in 2010 during the aftermath of the Great Recession. Following the April peak, the rate declined rapidly, but at the end of the year was still about twice as high as the pre-pandemic level in February 2020.

Visually it appears that a single TC declining at an exponential rate would fit much of the year although at year’s end another more persistent outlier effect may be needed. This suggests that a TC may be an important component in the outlier mix when modeling state and metro areas. Of course, in real time we did not have this information. At the beginning of the pandemic the sample is undersized for testing anything other than an AO. As several new observations accumulated, we chose to use LSS. Our reasoning is that AOS is likely to discount too much of the pandemic data, whereas in the interim period as new data accumulate, LSS may make better use of the pandemic data if the level or rate of change in the series stabilizes or experiences discrete shifts.

At the end of the year, with the benefit of hindsight, we compared 4 alternative specifications shown in Table 2 for each of the 421 unemployment time series, using the minimum AICC criterion as discussed above to select the most parsimonious combination of outliers. The base period for identifying and estimating the parameters of the RegArima models is the period from 2000 to 2019. The span for outlier detection is March through December 2020. The t-test critical value is set at 3.16 for 10 observations to be tested.

AOS	Regression spec for sequence of AO’s
LSS	Regression spec for sequence of LS’s
LS_TC_AO	Automatic detection of AO’s, LS’s & TC’s
LS_AO	Automatic detection of LS & AO’s

Figure 2 shows for each option the percentage of the 421 areas with the minimum AICC. The LS_TC_AO performed best in over half of the areas; if the TC is excluded, the percentage drops to 19% (LS_AO). That the TC plays an important role in attaining a parsimonious specification is not surprising given the shape of the national unemployment rate during the pandemic as shown in Figure 1. LSS is the second but distant best in 23% of the areas. As expected, it performs much better than AOS which is preferred in only 6% of the areas.

During real time outlier detection, the AOS is the only viable option for at least the first few observations where the sample of pandemic observations is too small to effectively estimate alternative outlier effects. AOS provides crucial protection at the April peak of the pandemic where the magnitude of the distortions to seasonal adjustment would be the largest. Moreover, continuing to compute AOS for new observations has diagnostic value since the coefficients for each of the AO dummy variables equal the prediction error for the corresponding data point and may reveal a pattern suggestive of a more parsimonious specification.

Figure 3 displays the minimum, median and maximum number of outliers per area for each option. Not surprisingly, the median number of regressors is smallest for LS_TC_AO and highest for AOS. The LSS has a slightly lower median than AOS.

Figure 4 shows the distribution of the total number of outliers for each option across each month of the pandemic period. For both the LSS and AOS options the upper limit to the total number of outliers per month is the total number of series (421) since the two options cannot have more than one outlier in the same series for a given month. The other two options may have multiple outliers at the same time point. The April shock is picked up by all the options; immediately afterwards all but AOS have much fewer regressors in later months. Between May and December LS_TC_AO has the fewest number of regressors.

Following the April peak AOS looks increasingly over parameterized compared to the other options. At the end of the series, however, AOS appears to be underfitting. By December the AOS option has no outlier effects in 183 (43%) areas. Since the AO is either on or off at this observation this implies many of the areas have recovered from the pandemic even though at the national level the recovery is far from complete. Since an AO estimate is effectively based on only one observation per area, the decreasing but non-zero magnitude of the outlier effects may not be as accurately estimated as with the other options that use more data. As for the other options that contain LS, their effects will continue indefinitely at a fixed level if no additional outliers are added. As a result, they must be continuously monitored for potential over estimation of the level of the series in the post-pandemic period.

4. Conclusions

After accumulating 10 observations in the pandemic period, the AOS method, which heavily discounts the pandemic data, is the least favored of 4 options for all but 6% of the 421 unemployment series in this study. The overall pattern is an eventual return towards normal or new normal where pandemic data provide useful information for seasonal adjustment. LS's in combination with TC's and AO's provides the best fit overall. Care must be taken to ultimately turn off pandemic LS's that outlive their usefulness.

Acknowledgment

The authors thank Brian Monsell for modifying X-11A-S in a special build that extends the LSS option to include a LS at the endpoint. This modification simplifies real time identification of LS's when there are many series to process. This feature is now available in the latest public release (Build 58) of X-13A-S.

References

- Box, G. E. P. and G. C. Tiao (1975). Intervention analysis with applications to economic and environmental problems. *Journal of the American Statistical Association* 70, 70-79.
- Lytras, D., and Bell, W.R. (2013) "Modeling Recession Effects and the Consequences on Seasonal Adjustment," Proceedings of the Business and Economic Section of the 2013 Joint Statistical Meetings.
- Ljung, G. M. (1993). On outlier detection in time series. *Journal of Royal Statistical Society B* 55, 559-567.
- U.S. Census Bureau (2020) "X-13ARIMA Reference Manual, Version 1.1," Time Series Research Staff, Center for Statistical Research and Methodology, Washington, DC, available at <https://www.census.gov/data/software/x13as.html>.

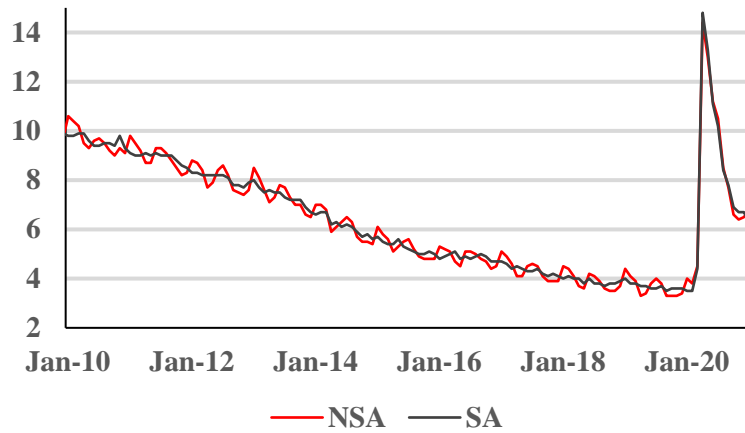


Figure 1: U.S. CPS Unemployment Rate

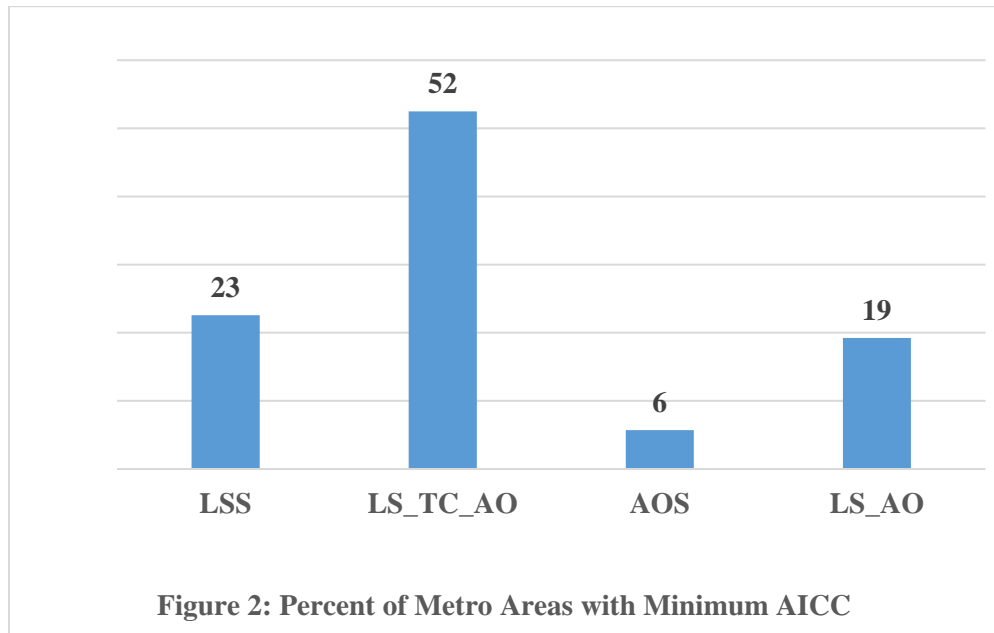


Figure 2: Percent of Metro Areas with Minimum AICC

