Identifying Seasonality

Tucker S. McElroy and Kathleen McDonald-Johnson

Research and Methodology Directorate

and Economic Statistical Methods Division

Identifying Seasonality

This presentation is a tutorial based upon “Identifying Seasonality,” a report (January 21, 2022) of the Interagency Seasonal Adjustment Team:

- William R. Bell (U.S. Census Bureau)
- Baoline Chen (U.S. Bureau of Economic Analysis)
- Kathleen M. McDonald-Johnson (U.S. Census Bureau)
- Tucker S. McElroy (U.S. Census Bureau)
- Brian C. Monsell (U.S. Bureau of Labor Statistics)
- Osbert Pang (U.S. Census Bureau)

United States Census Bureau

Introduction

Seasonal adjustment, the process of removing seasonal patterns from time series data, generally involves three steps:

1. Determine if the time series is seasonal. If not, stop.
2. If the time series is seasonal, perform seasonal adjustment.
3. Determine if the seasonal adjustment performed is adequate. If so, stop. If not, return to Step 2, modify the seasonal adjustment procedure (for example, choose different program options), and continue this process until obtaining an adequate seasonal adjustment.

Introduction

If no seasonal adjustment is acceptable, either:

a) Decide not to adjust the series

b) Or accept an inadequate adjustment

Introduction

Graphs, hypothesis tests, and other diagnostics can assist with this procedure (steps 1 through 3).

In Step 3: we assess whether *residual seasonality* (RS) is present in the seasonally adjusted series. Our goal is to have no RS.

Introduction

This presentation:

- Documents seasonality diagnostics available in the X-13ARIMA-SEATS seasonal adjustment software.
- Illustrates their application to time series data.
- Examines testing for seasonality in an unadjusted time series, or *pretesting.*
- Examines checking for RS in a seasonally adjusted time series, or *posttesting.*

Introduction

Remark: seasonality diagnostics available for the unadjusted series (Step 1) are essentially the same diagnostics available for detecting RS in a seasonally adjusted series (Step 3), as well as in the model residuals, or the estimated irregular.

- The only difference might be the data span.
- This is common practice, even though determining whether an unadjusted series is seasonal is a different problem from determining whether a seasonally adjusted series contains RS.

Outline

1. Discussion of the direct and indirect seasonal adjustment methods for aggregate series.
2. Description of common software for seasonal adjustment.
3. Description of seasonality diagnostics currently available in X-13ARIMA-SEATS.
4. Discussion of examples of time series and their associated diagnostics.

Direct and Indirect Seasonal Adjustment

Most published series are components of some greater aggregate; so it is natural to study seasonal adjustment adequacy for batches of related series.

- Direct Adjustment: an aggregate time series is seasonally adjusted without respecting its relation to its component series. “Sum, then Adjust.”
- Indirect Adjustment: an aggregate time series is seasonally adjusted by aggregating its individual seasonally adjusted components. “Adjust, then Sum.”
- Aggregation: typically is summation across hierarchy or sampling frequency.

United States® Census Bureau

Direct and Indirect Seasonal Adjustment

It can happen that many components exhibit no detectable seasonality according to pretesting, and yet the aggregation of them is seasonal! Hence indirect seasonal adjustment yields RS. What to do?

- Do direct seasonal adjustment. (But aggregation relations will be violated.)
- Modify the models for each component time series, so as to still yield adequacy for components, and such that the aggregate has no seasonality either. (This is tricky.)
- Modify the component seasonal adjustments as little as possible, but such that their aggregate has no seasonality. (Requires numerical optimization.)

United States® Census Bureau

Seasonal Adjustment Software

X-11 Method as Implemented in X-13ARIMA-SEATS

- X-11 is a seasonal adjustment procedure that the U.S. Census Bureau developed in the 1950s and 1960s.
- Statistics Canada enhanced the method with the addition of ARIMA modeling for forecast extension.
- The X-11 procedure separates a time series into a trend-cycle component, a seasonal component, and an irregular component by iteratively filtering the original (or transformed) time series, using moving averages.
- Users (or automatic provisions in the software) choose the moving average filters from a set of fixed (precoded) options.
- X-13ARIMA-SEATS has an X11 specification that implements the X-11 seasonal adjustment method.

United States® Census Bureau

Seasonal Adjustment Software

SEATS Method as Implemented in X-13ARIMA-SEATS

- SEATS (Signal Extraction in ARIMA Time Series) is the seasonal adjustment part of the program TRAMO-SEATS that Agustín Maravall and Victor Gómez developed while at the Bank of Spain.
- TRAMO (Time series Regression with ARIMA noise, Missing observations, and Outliers) fits regARIMA models (short for regression models with ARIMA errors).
- SEATS uses the fitted model to estimate the components for seasonal adjustment, using the canonical ARIMA model-based approach of Hillmer and Tiao to decompose a time series into a trend-cycle component, a seasonal component, and an irregular component, each of which follows its own underlying ARIMA model.

United States® Census Bureau

Diagnostics for Identifying Seasonality

First: graph the time series!

- Graphing the series across consecutive time points, as well as year over year, will help in determining whether the series has a seasonal pattern.
- Graphs can illuminate additional patterns of series behavior, as well as unusual points or subspans within a series.
- Graph the outlier-adjusted (or prior-adjusted) series to discern whether large outliers might obscure the series' patterns.

Diagnostics for Identifying Seasonality

Second: examine the autocorrelation function (ACF) of the original and first differenced series.

- This will assess the dependence over time in the data, including seasonal dependence.
- For a seasonal monthly series, one would expect large, positive ACF values at lag 12, and multiples of 12 (24, 36, etc.).
- For a seasonal quarterly series, one would expect large, positive ACF values at the seasonal lags 4, 8, 12, etc.
- Large seasonal lag autocorrelations are necessary, but not sufficient, to indicate the presence of seasonality; small seasonal lag autocorrelations indicate a lack of evidence of seasonality in the series.

United States® Census Bureau

Diagnostics for Identifying Seasonality

Three specific seasonal adjustment diagnostics:

1. Model-based F Test
2. Maravall’s QS
3. Peaks at Seasonal Frequencies in Spectral Plots

United States® Census Bureau

Model-based F Test

The model-based F test checks for evidence of a stable seasonal pattern in the original series.

- Designed for use in pretesting.
- Testing for RS can be done, but we recommend using a subspan.
- Requires one to build an ARIMA model with fixed seasonal regressors.
- If a SARIMA(pdq)(011) has been already fitted, one can try fitting an ARIMA(pdq) with fixed seasonal regressors.

Model-based F Test

How it works: the null hypothesis is that the coefficients of the fixed seasonal regressors are zero.

- Rejections of the null correspond to non-zero coefficients, indicating the existence of a fixed seasonal pattern called *stable seasonality*.
- Results depend upon the parameters of the fitted ARIMA model, and the subspan.

United States® Census Bureau

Model-based F Test

Example: monthly U.S. Retail Sales of Office Supply, Stationery, and Gift Stores, from 2002 through 2017 (Source: Monthly Retail Trade and Food Services, U.S. Census Bureau).

- Test based on the full span of the original series.
- Table below shows X-13A-S output; the null is rejected.

Regressor	DF	F Statistic	P-Value
Seasonal	11, 170	369.04	0.00

United States® Census Bureau

Model-based F Test

Example: the seasonal pattern is very apparent.

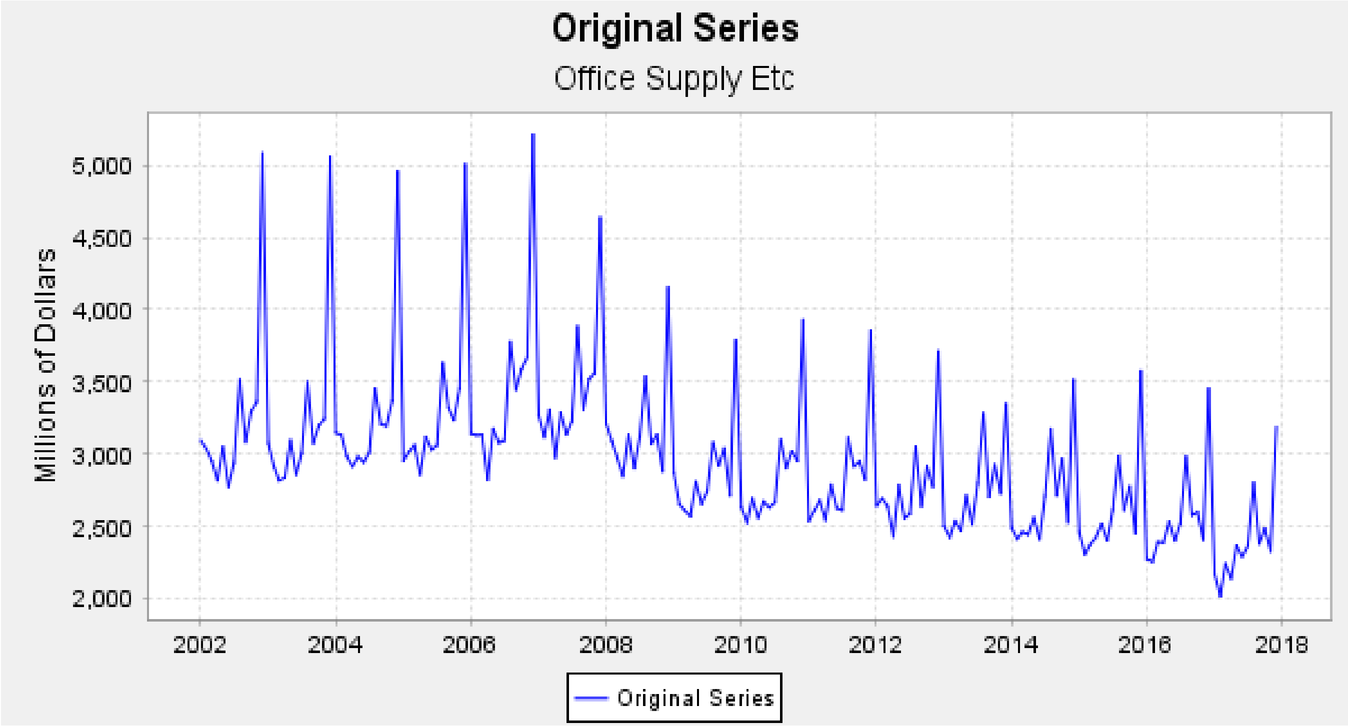

Model-based F Test

Example: the seasonal pattern appears to be stable.

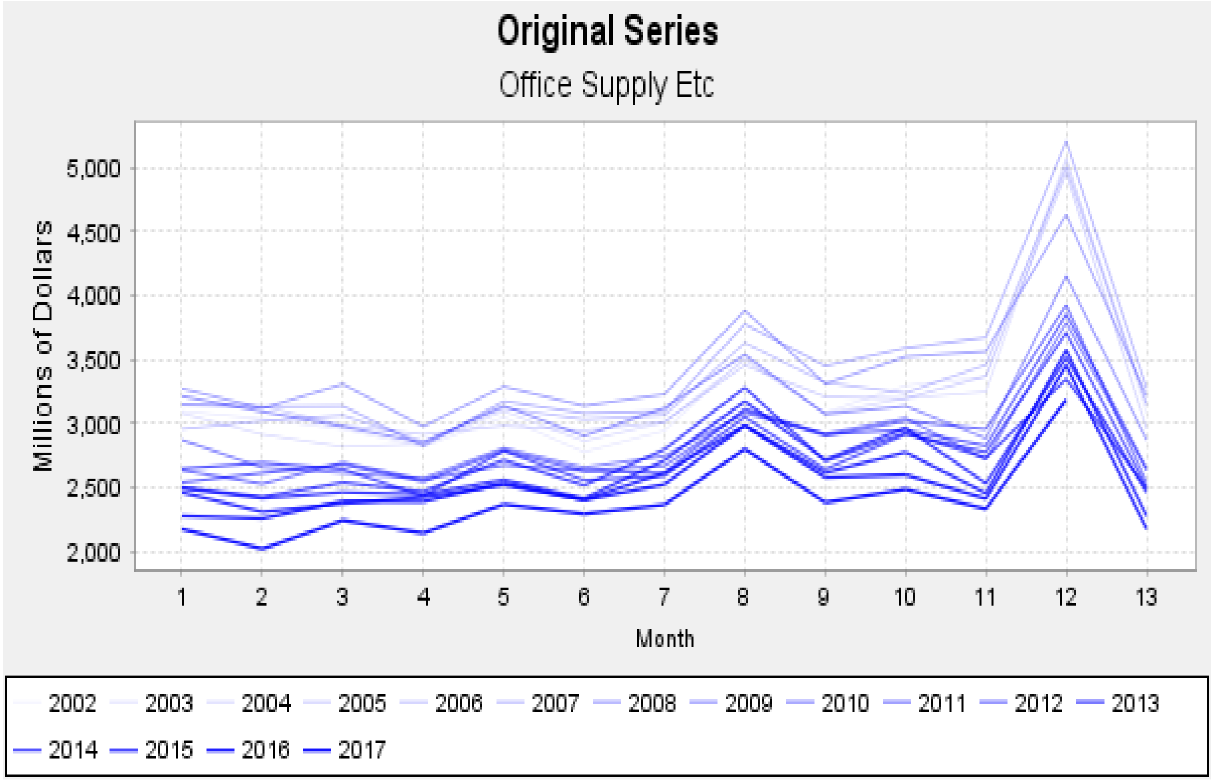

Maravall's QS

The QS diagnostic examines the first two seasonal lag autocorrelations, with large positive values indicating seasonality may be present.

- The lags are 12 and 24 for monthly data, 4 and 8 for quarterly data.
- The null hypothesis is that these first two seasonal autocorrelations are zero, corresponding to the series not being seasonal.
- The rationale is that a series with seasonality (or RS) should exhibit substantial positive autocorrelation at seasonal lags.

Maravall’s QS

How it works:

- The series is differenced to remove trend effects, and the sample autocorrelations are computed.
- If the first seasonal autocorrelation is zero or negative, QS is set to zero.
- Otherwise, QS equals a function of the first and second seasonal autocorrelation, and a heuristic chi-square distribution with two degrees of freedom is used to obtain p-values.

Maravall’s QS

Caveats:

- Any small degree of seasonal lag autocorrelation will be flagged as seasonal, so even very mild degrees of seasonality can lead to rejection of the null hypothesis.
- Therefore, QS is better suited for detecting RS. For pretesting, series with very mild degrees of seasonality need not be seasonally adjusted.
- QS can be significant for non-seasonal time series that happen to have high seasonal lag autocorrelation.

Maravall’s QS

Types of applications:

- Original series, with or without extreme value adjustment (pretest)
- RegARIMA model residuals
- Seasonally adjusted series, with or without extreme value adjustment (posttest)
- Irregular component, with or without extreme value adjustment (posttest)
- Indirect seasonal adjustments

Maravall's QS

Other remarks:

- Run QS on a subspan of the seasonal adjustment or irregular if there are concerns about sudden changes in seasonality, especially near the current values of the series.
- QS diagnostics on various components can yield differing results; the seasonally adjusted component is the most vital.
- QS can also be applied to a quarterly version of monthly stock time series.

United States® Census Bureau

Maravall's QS

Example: monthly U.S. Retail Sales of Office Supply, Stationery, and Gift Stores, from 2002 through 2017 (Source: Monthly Retail Trade and Food Services, U.S. Census Bureau). Tables show X-13A-S output.

- Test based on monthly and quarterly flow time series.
- Null of no seasonality is rejected for the original series, indicating the presence of seasonality.
- Null is not rejected for the regARIMA model residuals, seasonally adjusted series, and the irregular component; there is no RS.

United States Census Bureau

Maravall’s QS

Series	Span	QS	p-value
Original Series	Full Series	307.31	0.0000
Original Series (extreme value adjusted)	Full Series	322.95	0.0000
Residuals	Full Series	0.00	1.0000
Seasonally Adjusted Series	Full Series	0.00	0.9995
Seasonally Adjusted Series (extreme value adjusted)	Full Series	0.00	1.0000
Irregular Series	Full Series	0.00	1.0000
Irregular Series (extreme value adjusted)	Full Series	0.00	1.0000
Original Series	Subspan	132.56	0.0004
Original Series (extreme value adjusted)	Subspan	149.70	0.0000
Residuals	Subspan	0.00	1.0000
Seasonally Adjusted Series	Subspan	0.17	0.9170
Seasonally Adjusted Series (extreme value adjusted)	Subspan	0.00	1.0000
Irregular Series	Subspan	0.00	1.0000
Irregular Series (extreme value adjusted)	Subspan	0.00	1.0000

United States® Census Bureau

Maravall’s QS

Series	Span	QS	p-value
Original Series	Full Series	99.78	0.0000
Original Series (extreme value adjusted)	Full Series	103.47	0.0000
Seasonally Adjusted Series	Full Series	0.00	1.0000
Seasonally Adjusted Series (extreme value adjusted)	Full Series	0.00	1.0000
Original Series	Subspan	42.33	0.0000
Original Series (extreme value adjusted)	Subspan	45.44	0.0000
Seasonally Adjusted Series	Subspan	0.00	1.0000
Seasonally Adjusted Series (extreme value adjusted)	Subspan	0.00	1.0000

United States® Census Bureau

Spectral Peaks at Seasonal Frequencies

Spectral plots are a graphical technique for identifying periodic patterns in a time series.

- For computing the spectrum reliably, at least eight years of data should be used.
- Spectral plots highlight the computed spectral estimates at the seasonal frequencies 1/12, 2/12, 3/12, 4/12, 5/12, and 6/12. These frequencies correspond to seasonal effects recurring several times per year, according to the numerator.
- Large, or *visually significant,* peaks at the seasonal frequencies indicate evidence of seasonality.

United States® Census Bureau

Spectral Peaks at Seasonal Frequencies

Thresholds for visual significance of a peak are a height that is

- Six *stars* above the taller of the two nearest-neighbor frequencies on the plot;
- Above the median height of all the plotted frequencies.

A *star* is a unit of measure equal to 1/52 of the spectral range (maximum minus minimum). The seasonal frequencies are marked by an S.

United States® Census Bureau

Spectral Peaks at Seasonal Frequencies

Spectral plots also indicate the trading day frequencies.

- For monthly series, 0.348 and 0.432 cycles per month are the trading day frequencies, marked with a T on the plot.
- Visually significant trading day peaks indicate that trading day regressors should be included in the model (or their specification should be fixed if they are already present).

Spectral Peaks at Seasonal Frequencies

Types of applications:

- Original series (pretest)
- RegARIMA model residuals
- Modified (for extremes) seasonally adjusted series (posttest)
- Modified (for extremes) irregular series (posttest)
- Indirect seasonal adjustment and irregular (posttest)

United States® Census Bureau

Spectral Peaks at Seasonal Frequencies

How it works:

- If any of the first four seasonal frequencies is visually significant, the possible presence of seasonality is indicated.
- The fifth and sixth seasonal frequencies are typically ignored.
- Some users require at least two peaks to be visually significant in order to flag seasonality.
- Can be used as pretest or posttest, although seasonally adjusted series often have a trough, rather than a peak, at seasonal frequencies. This is consistent with no RS.

United States® Census Bureau

Spectral Peaks at Seasonal Frequencies

Example: monthly U.S. Retail Sales of Office Supply, Stationery, and Gift Stores, from 2002 through 2017 (Source: Monthly Retail Trade and Food Services, U.S. Census Bureau). Figures show X-13A-S output.

- First figure is spectral plot from the X-13A-S main output file, for the original series (pretest).
- Second figure omits the stars, but includes S for visually significant spectral peaks. It is also for the original series (pretest).
- Third figure is for the seasonally adjusted series (posttest).

Spectral Peaks at Seasonal Frequencies

G 0 10*LOG(SPECTRUM) of the differenced, transformed Prior Adjusted Series (Table B1)

Spectrum estimated from 2010.Jan to 2017.Dec.

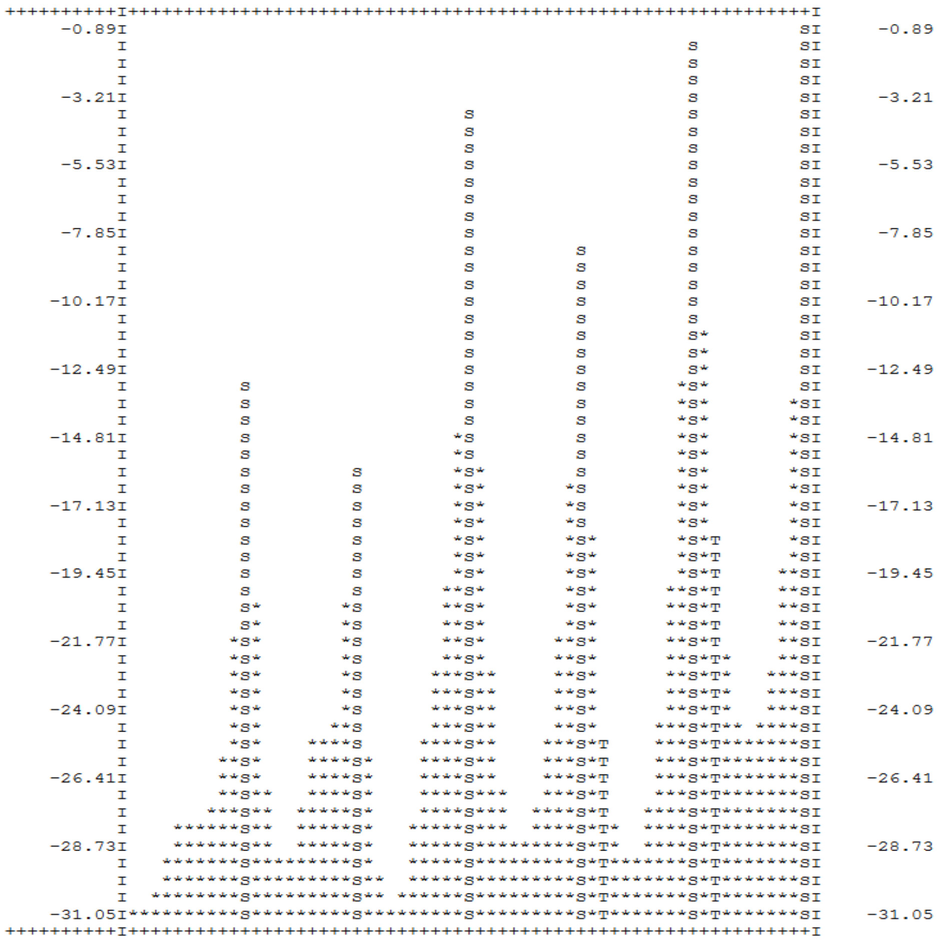

Spectral Peaks at Seasonal Frequencies

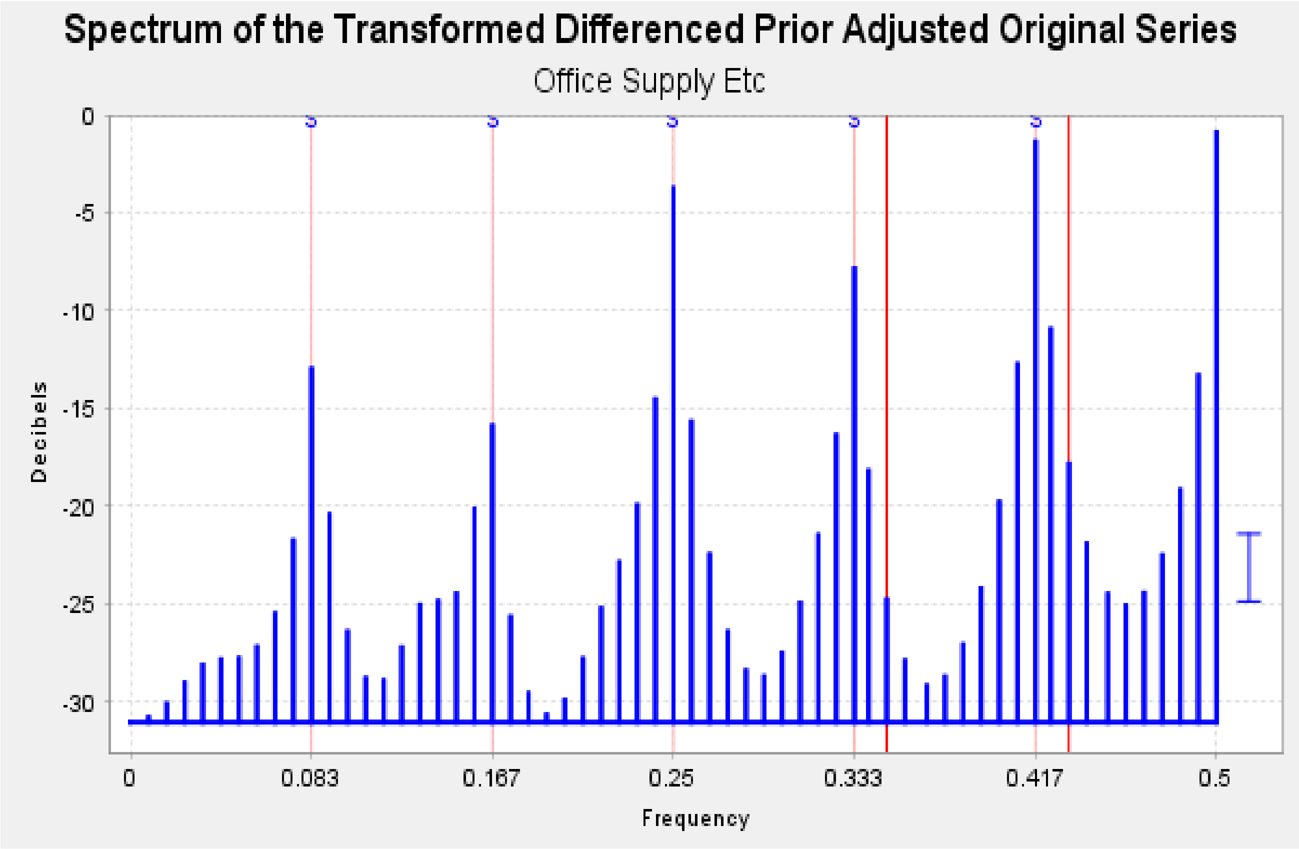

United States® Census Bureau

Spectral Peaks at Seasonal Frequencies

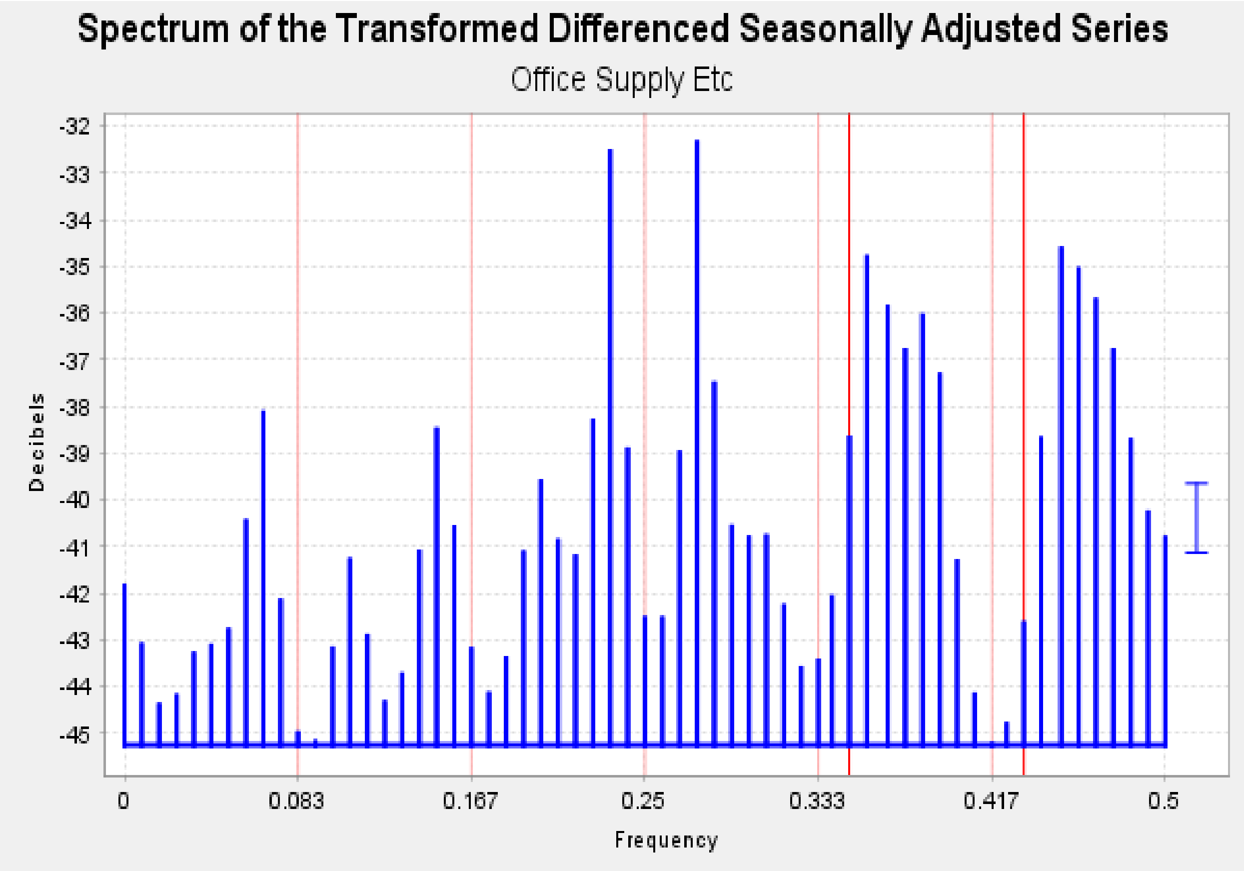

United States® Census Bureau

Spectral Peaks at Seasonal Frequencies

Example: we also show the table output from X-13A-S.

- Whereas the original series has seasonality, the adjusted series does not.
- The model residuals have a (visually) non-significant peak.

Series Name	Sig Ori Peaks	Resid Peaks	Sig SAdj Peaks	Sig Irr Peaks	Nonsig Seasonal Peaks	Nonsig TD Peaks
Office Supply Etc	s1 s2 s3 s4 s5				rsd [1.1]	

Remarks on Pretesting and Posttesting

The model-based F test examines a different form of seasonality from the QS and Spectral diagnostics.

- The model-based F test is testing for the presence of a fixed seasonal pattern.
- Even if seasonality is evolving over time, there will typically be a fixed pattern underlying the evolving seasonal component.
- Detection of a fixed seasonal pattern is a strong indication of seasonality.

United States® Census Bureau

Remarks on Pretesting and Posttesting

On the other hand:

- QS and Spectral diagnostics test for seasonal dependence in the form of positive seasonal autocorrelations (QS) or peaks in the spectrum at seasonal frequencies (Spectral).
- The presence of fixed seasonal effects will produce large estimated seasonal autocorrelations and strong seasonal peaks in the estimated spectrum, and this will generate significant QS and Spectral diagnostics.
- QS and the Spectral diagnostic can also detect moderate (or even mild) seasonal autocorrelation that would not necessarily produce discernible seasonal patterns in the data, and thus may not suggest seasonal adjustment.

United States® Census Bureau

Remarks on Pretesting and Posttesting

- For pretesting, QS and the Spectral diagnostic are less useful than the model-based F test.
- For posttesting, the model-based F test is not useful because direct adjustment removes all fixed seasonal effects. But QS and the Spectral diagnostic are useful, since mild to moderate positive seasonal autocorrelation can be taken as indicating residual seasonality.
- A significant posttest result might indicate a need to modify the regARIMA model (if using the SEATS method) or modify the seasonal adjustment options (if using the X-11 method) to find settings that more thoroughly remove the seasonal effects.

United States® Census Bureau