Examining the Performance of Seasonality Diagnostics for Detecting Residual Seasonality

Osbert Pang
U.S. Census Bureau

April 26, 2018

Acknowledgements

Joint work with Brian Monsell and William Bell.

Introduction

- There are multiple diagnostics that can be used to detect the presence of seasonality in a given time series.
- These diagnostics should be adequate when used on raw series.
- This may not necessarily be true when applying same diagnostics to seasonally adjusted series.
- Same may also hold when applying diagnostics to a quarterly seasonal adjustment obtained by aggregating a monthly seasonally adjusted series (which is frequently done when raw data is not available for an unadjusted quarterly series).
- Simulation study was conducted to look at some diagnostics.

Simulation study

- To start, let θ_1 take one of the values 0.3, 0.5, 0.7, and 0.9, and let the underlying model be a (nonseasonal) ARIMA (0,1,1).

- Simulate 5 000 monthly series of length 50 years for each of the above values of θ_1 (arbitrary starting point of Jan 1965).

- General idea: Using a few different spans for these series, check the model-based F-test on a fixed seasonal regressor and the QS and QSS statistics (QSS is the QS, but applied only to the last 8 years of the series).

- One point: these are nonseasonal series being forced through a seasonal adjustment. In practice, these would not be adjusted, but we check what happens anyway.

Simulation study (2)

- For the F-test, look at
  - Fitting correct model form with seasonal regressor to original monthly series, whether using estimate $\widehat{\theta}_1$ or true model parameter θ_1
  - Fitting correct model form with seasonal regressor to quarterly aggregate of original monthly series
  - Fitting (0 1 1) with seasonal regressor to a seasonal adjustment of original series - adjustment done using x11 and automatic model identification
  - Fitting (0 1 1) with seasonal regressor to quarterly aggregate of previous monthly seasonal adjustment

Simulation study (3)

- For QS/QSS, look at statistic for
  - Original monthly series
  - Quarterly aggregate of original monthly series
  - Seasonal adjustment of original monthly series (again, adjustments done using x11 and automatic model identification)
  - Quarterly aggregate of monthly seasonal adjustment
  - Seasonal adjustment of quarterly aggregate of original

Simulation study (4)

	θ_1			
	0.3	0.5	0.7	0.9
fitting (0 1 1) with estimated $\widehat{\theta}_1$, proportion for which p-value ≤ 0.05				
50 years	0.0578	0.0520	0.0470	0.0486
15 years	0.0484	0.0506	0.0504	0.0554
8 years	0.0612	0.0594	0.0538	0.0586
fitting (0 1 1) to quarterly aggregate of original monthly series, proportion for which p-value ≤ 0.05				
50 years	0.0550	0.0524	0.0486	0.0524
15 years	0.0566	0.0564	0.0578	0.0576
8 years	0.0992	0.0928	0.0688	0.0732
fitting (0 1 1) to monthly seasonal adjustment, proportion for which p-value ≤ 0.05				
50 years	0.0000	0.0000	0.0000	0.0000
8 year subspan of 50 years	0.0000	0.0000	0.0000	0.0000
15 years	0.0000	0.0000	0.0000	0.0000
8 year subspan of 15 years	0.0000	0.0000	0.0000	0.0000
fitting (0 1 1) to quarterly aggregate of monthly seasonal adjustment, proportion for which p-value ≤ 0.05				
50 years	0.0000	0.0000	0.0002	0.0000
8 year subspan of 50 years	0.0182	0.0064	0.0010	0.0006
15 years	0.0014	0.0006	0.0010	0.0000
8 year subspan of 15 years	0.0192	0.0070	0.0012	0.0010

Table 1: Simulation results for model-based F-test, comparing against an $\alpha = 0.05$ level.

Simulation study (5)

- F-test on the seasonal regressor is about the right size when applied to the original monthly series and is not sensitive to underlying θ_1 (known result); slightly higher for shorter spans. Also, since (0 1 1) monthly aggregates to (0 1 1) quarterly, the size is again about right for the quarterly aggregate of the original series; again, shorter spans are more affected.

- Since additive seasonal adjustment is supposed to eliminate fixed seasonal effects, unsurprising that the F-test to fixed seasonal regressor over the full span of monthly seasonal adjustment is not significant. But also holds for 8 year subspan.

- Similar finding for quarterly aggregate of monthly adjustment. Using just the last 8 year subspan does find some series have seasonality present, but at lower than 5% levels, and with some dependence on the underlying θ_1.

- Results are similar when comparing against an α of 0.01.

Simulation study (6)

	θ_1			
	0.3	0.5	0.7	0.9
Original monthly series, proportion for which p-value $\leq$ 0.05				
QS, 50 years	0.0332	0.0428	0.0518	0.0540
QS, 15 years	0.0318	0.0408	0.0508	0.0524
Quarterly aggregate of original series, proportion for which p-value $\leq$ 0.05				
QS, 50 years	0.0230	0.0216	0.0342	0.0510
QS, 15 years	0.0188	0.0240	0.0266	0.0440
Seasonal adjustment of original monthly series, proportion for which p-value $\leq$ 0.05				
QS, 50 years	0.0000	0.0000	0.0000	0.0000
QS, 15 years	0.0000	0.0000	0.0000	0.0000
Quarterly aggregate of monthly seasonally adjustment, proportion for which p-value $\leq$ 0.05				
QS, 50 years	0.0000	0.0000	0.0000	0.0000
QS, 15 years	0.0000	0.0000	0.0000	0.0000
Seasonal adjustment of quarterly aggregate of original series, proportion for which p-value $\leq$ 0.05				
QS, 50 years	0.0000	0.0000	0.0000	0.0000
QS, 15 years	0.0000	0.0000	0.0000	0.0000

Table 2: Simulation results for QS testing of series, for $\alpha = 0.05$.

Simulation study (7)

- QSS numbers are similar to QS numbers when available (not for the last 8 year subspans of 15 year quarterly series).

- Simulated size seems to be affected by underlying value of θ_1, unlike the model-based F.

- Quarterly aggregates of original monthly series seem to have lower levels than the monthly series.

- As with the model-based F, QS applied to seasonal adjustment (of nonseasonal data) does not really flag anything.

United States Census Bureau

Simulation study (8)

	θ_1			
	0.3	0.5	0.7	0.9
Original monthly series, proportion for which p-value $\leq$ 0.01				
QS, 50 years	0.0112	0.0164	0.0190	0.0208
QS, 15 years	0.0108	0.0110	0.0190	0.0204
Quarterly aggregate of original series, proportion for which p-value $\leq$ 0.01				
QS, 50 years	0.0048	0.0050	0.0108	0.0196
QS, 15 years	0.0044	0.0084	0.0098	0.0180

Table 3: Subset of simulation results for QS testing of series, for $\alpha = 0.01$.

- Note that simulated size appears to be affected by the underlying value of θ_1, but not consistently so.

- About right for larger θ_1 and too low for smaller θ_1 when looking at $\alpha = 0.05$.

- But about right for smaller θ_1 and too high for larger θ_1 when looking at $\alpha = 0.01$.

Simulation study (9)

- Since QS statistic is calculated using the first 2 seasonal lag autocorrelations, might be worth taking a look at the acf.

- Take the first-differenced monthly seasonal adjustments (for both 50- and 15-year spans) and look at the autocorrelations at lags 12 and 24.

- Then, take the first difference of the quarterly aggregates of those monthly adjustments and look at the autocorrelations at lags 4 and 8.

- Means (and standard deviations) can be computed over the simulated series for each of the values of θ_1 .

Simulation study (10)

	θ_1			
	0.3	0.5	0.7	0.9
Autocorrelations for first differenced monthly seasonally adjusted series				
50 years, average lag 12	−0.1809	−0.1802	−0.1560	−0.1528
50 years, average lag 24	−0.1132	−0.1180	−0.1175	−0.1183
15 years, average lag 12	−0.1765	−0.1762	−0.1671	−0.1616
15 years, average lag 24	−0.1111	−0.1143	−0.1156	−0.1158
Autocorrelations for first differenced quarterly aggregates of monthly SA				
50 years, average lag 4	−0.1533	−0.1568	−0.1480	−0.1533
50 years, average lag 8	−0.1056	−0.1089	−0.1112	−0.1164
15 years, average lag 4	−0.1682	−0.1688	−0.1664	−0.1619
15 years, average lag 8	−0.1152	−0.1149	−0.1115	−0.1136

Table 4: Simulation results for seasonal lag autocorrelations of first differenced series, means.

Simulation study (11)

- So, average autocorrelation at seasonal lag corresponding to 1 year back is in the -0.2 to -0.15 range, approximately

- Some slight differences for time spans and values of θ_1, but not significantly so.

- Average autocorrelation at seasonal lag corresponding to 2 years back seems to center right around -0.11.

- Not shown, but standard deviations are fairly consistent in magnitude across θ_1, slightly higher for shorter spans, and higher for quarterly aggregation than for monthly.

Simulation study (12)

	θ_1			
	0.3	0.5	0.7	0.9
Autocorrelations for first differenced monthly seasonally adjusted series, proportion positive				
50 years, lag 12	0.0002	0.0004	0.0012	0.0020
50 years, lag 24	0.0092	0.0060	0.0094	0.0094
15 years, lag 12	0.0054	0.0044	0.0144	0.0198
15 years, lag 24	0.0796	0.0846	0.0878	0.0952
Autocorrelations for first differenced quarterly aggregates of monthly SA, proportion positive				
50 years, lag 4	0.0108	0.0112	0.0250	0.0330
50 years, lag 8	0.0750	0.0692	0.0790	0.0852
15 years, lag 4	0.0774	0.0660	0.0864	0.1080
15 years, lag 8	0.1834	0.1804	0.1960	0.2080

Table 5: Simulation results for seasonal lag autocorrelations, proportion of autocorrelations that are positive.

Simulation study (13)

- For monthly adjustments, very few positive autocorrelations at seasonal lags observed when looking at the acf over 50 years, but lag 24 seems noticeably higher for the acf over 15 years.

- More positive autocorrelations at seasonal lags for the acf of first differenced quarterly aggregates, especially for shorter time span.

- Some dependence on the underlying θ_1 for number of positives.

- But for the most part, the autocorrelations at seasonal lags (for a seasonal adjustment) seem to be concentrated below 0.

Simulation study (14)

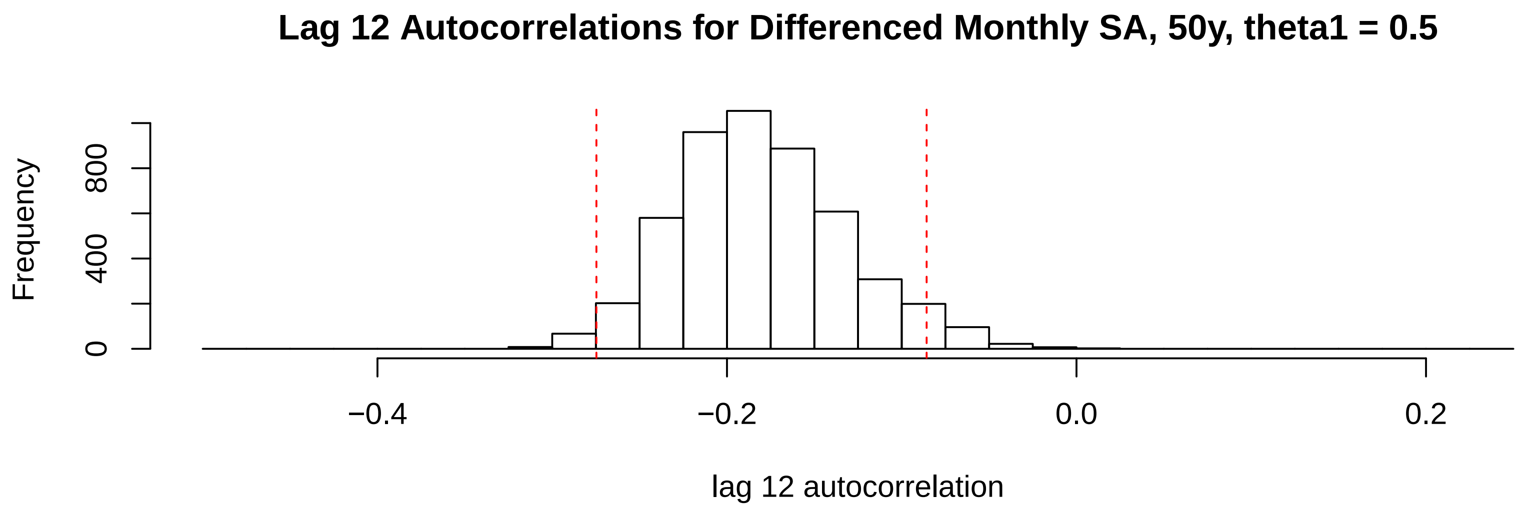

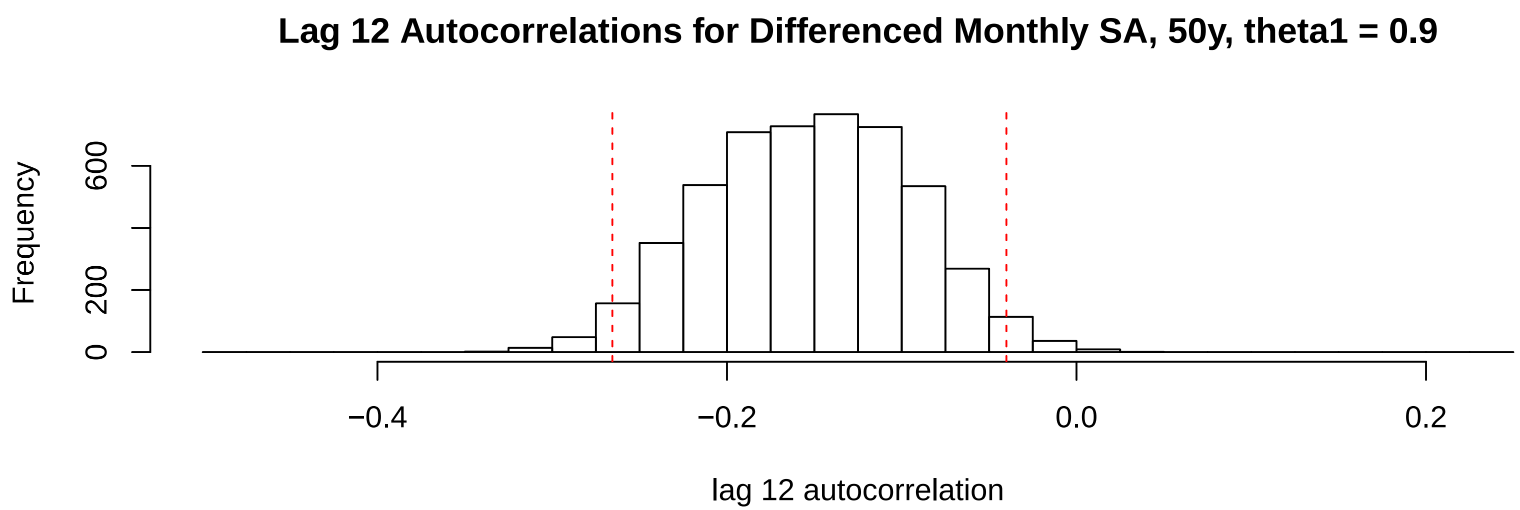

United States Census Bureau

Simulation study (15)

Lag 24 Autocorrelations for Differenced Monthly SA, 50y, theta1 = 0.5

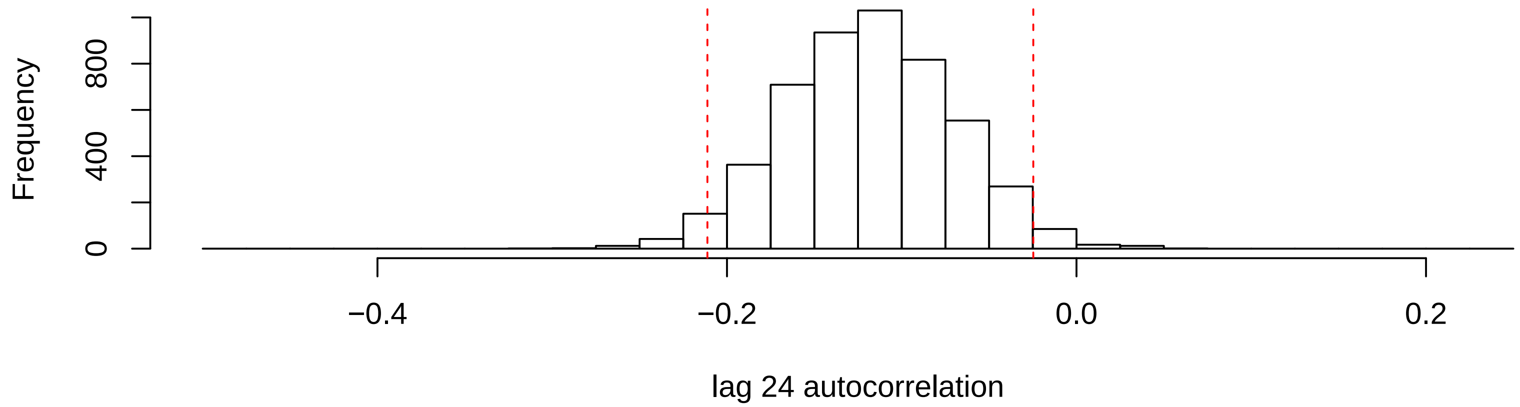

Lag 24 Autocorrelations for Differenced Monthly SA, 50y, theta1 = 0.9

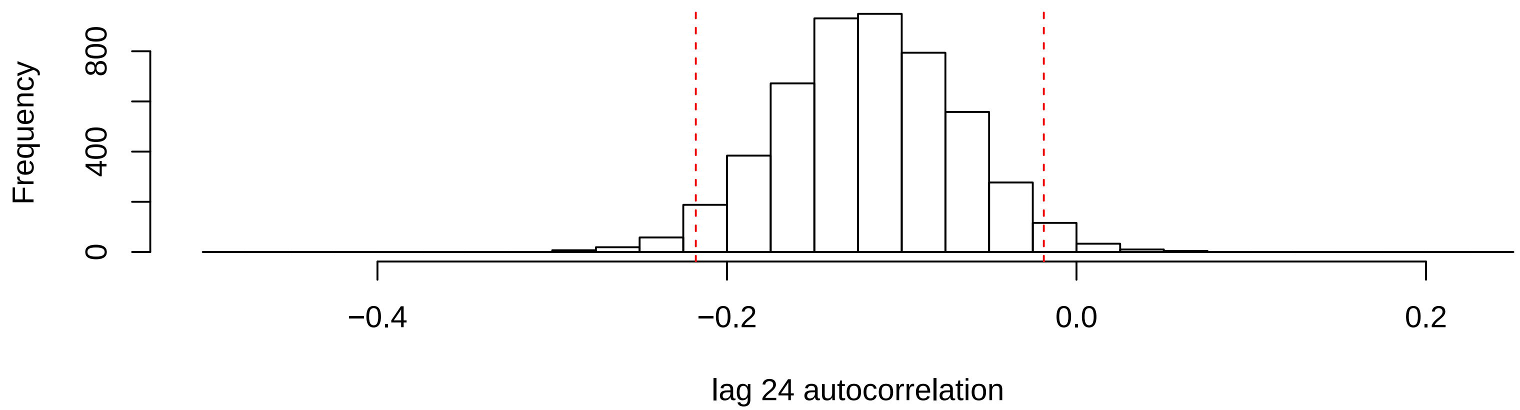

Simulation study (16)

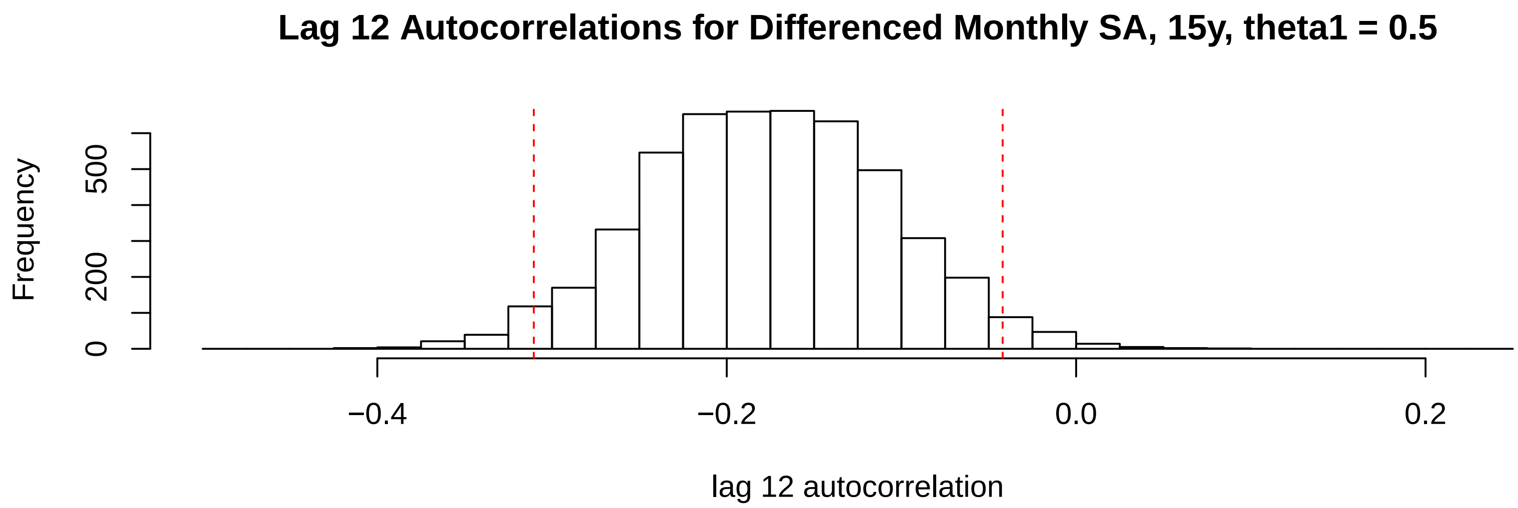

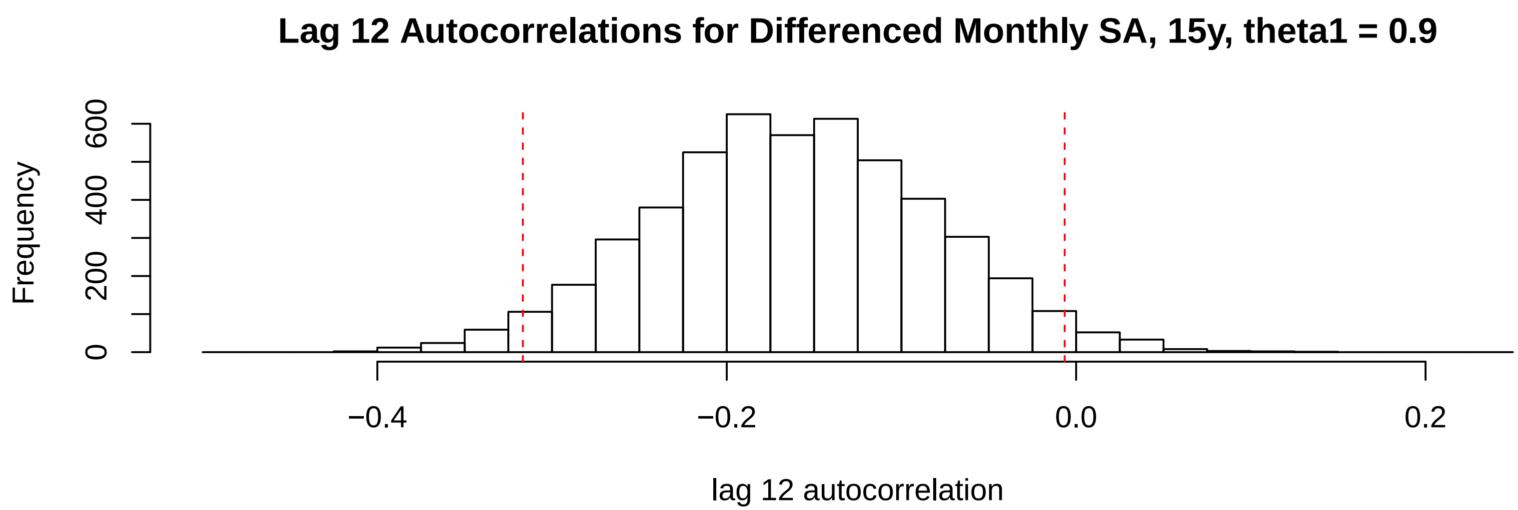

Simulation study (17)

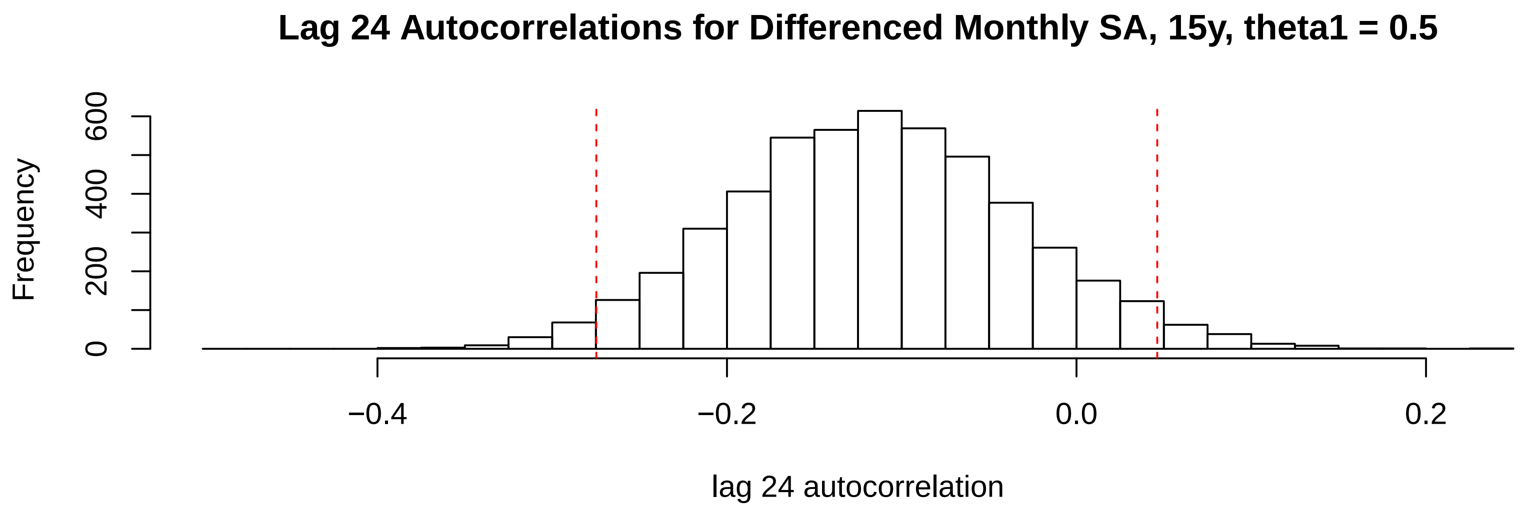

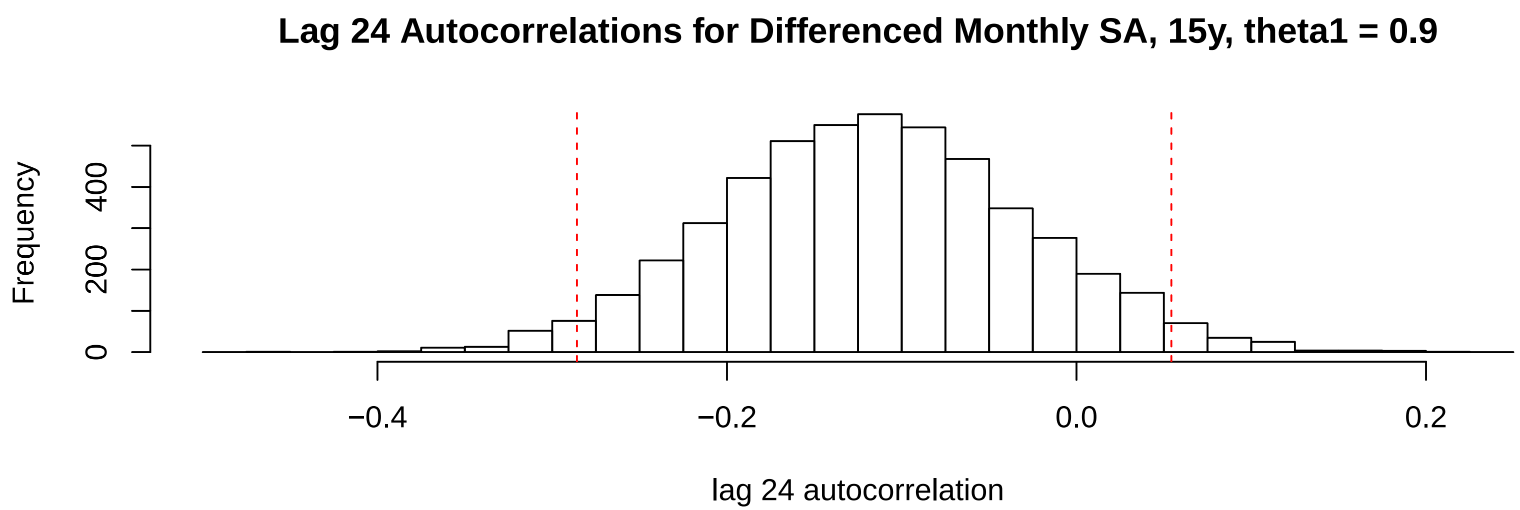

Simulation study (18)

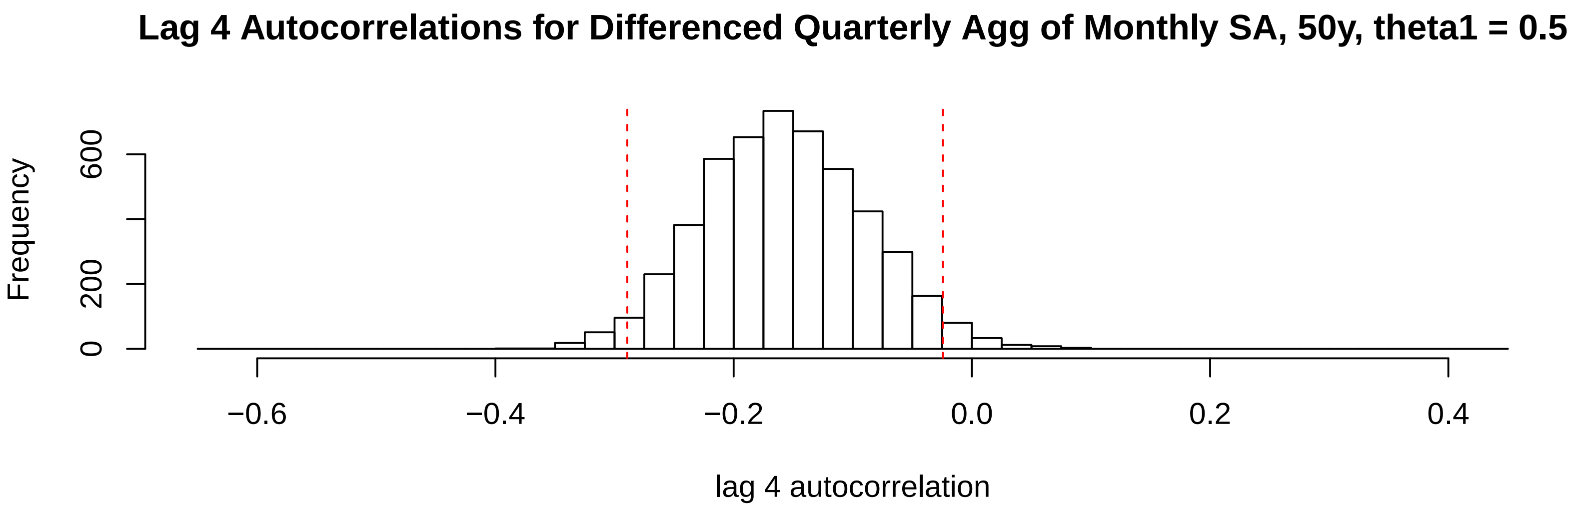

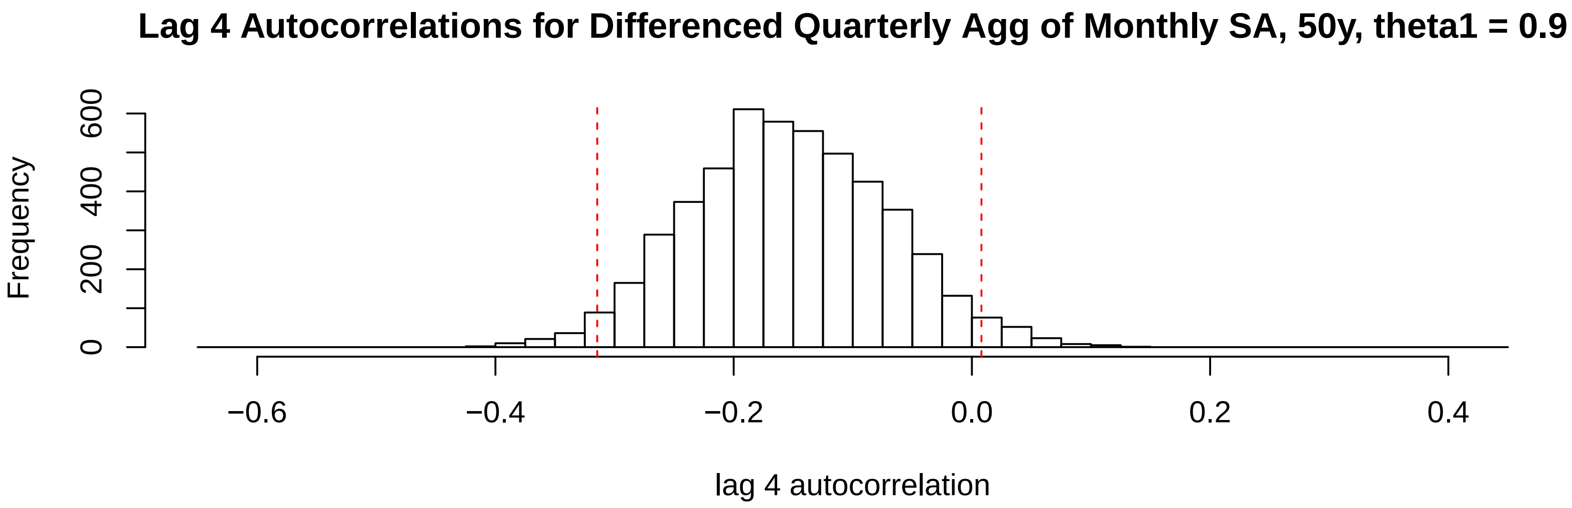

United States Census Bureau

Simulation study (19)

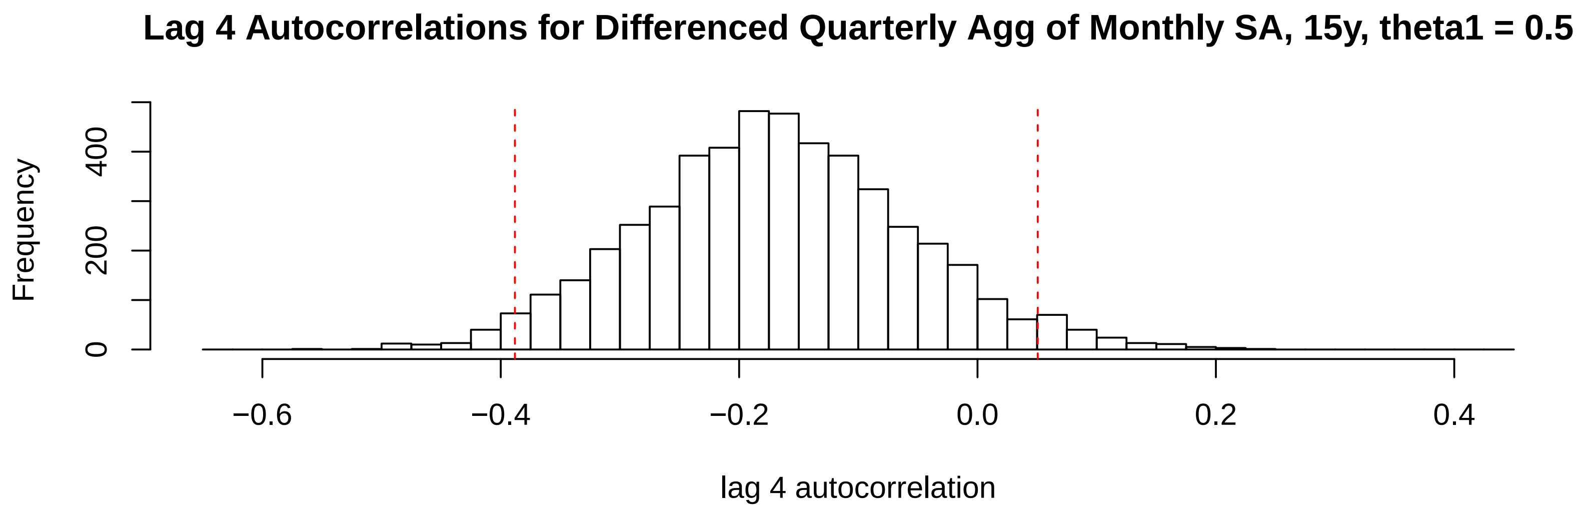

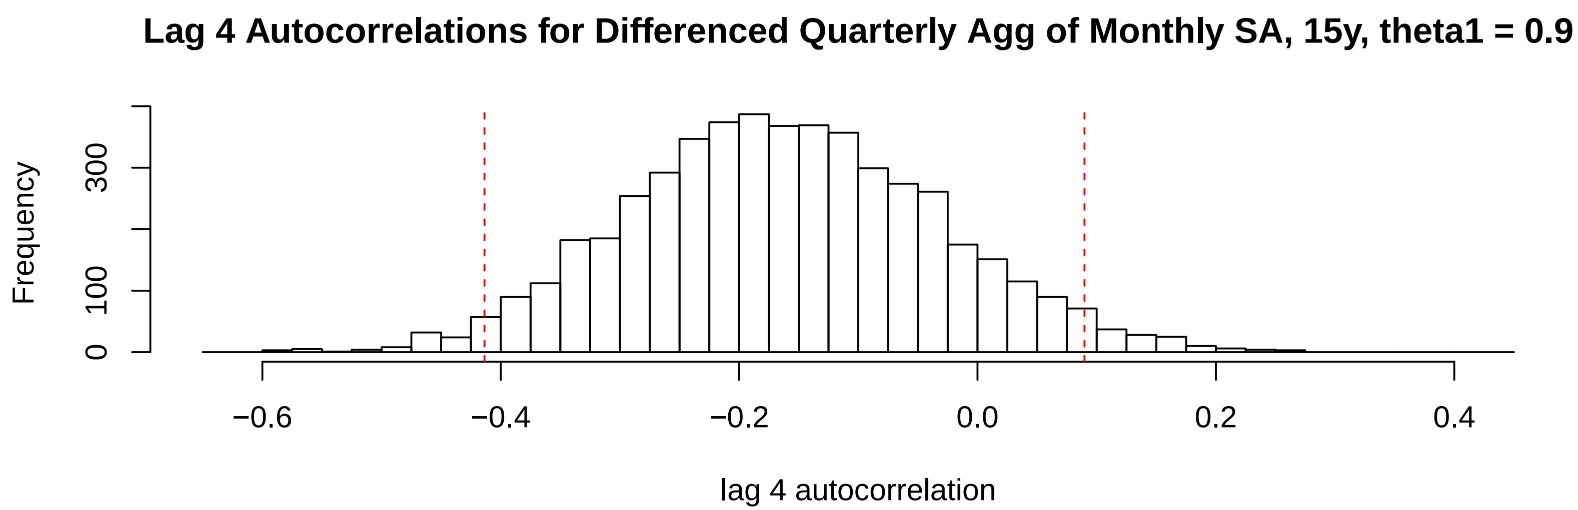

United States Census Bureau

Simulation study (20)

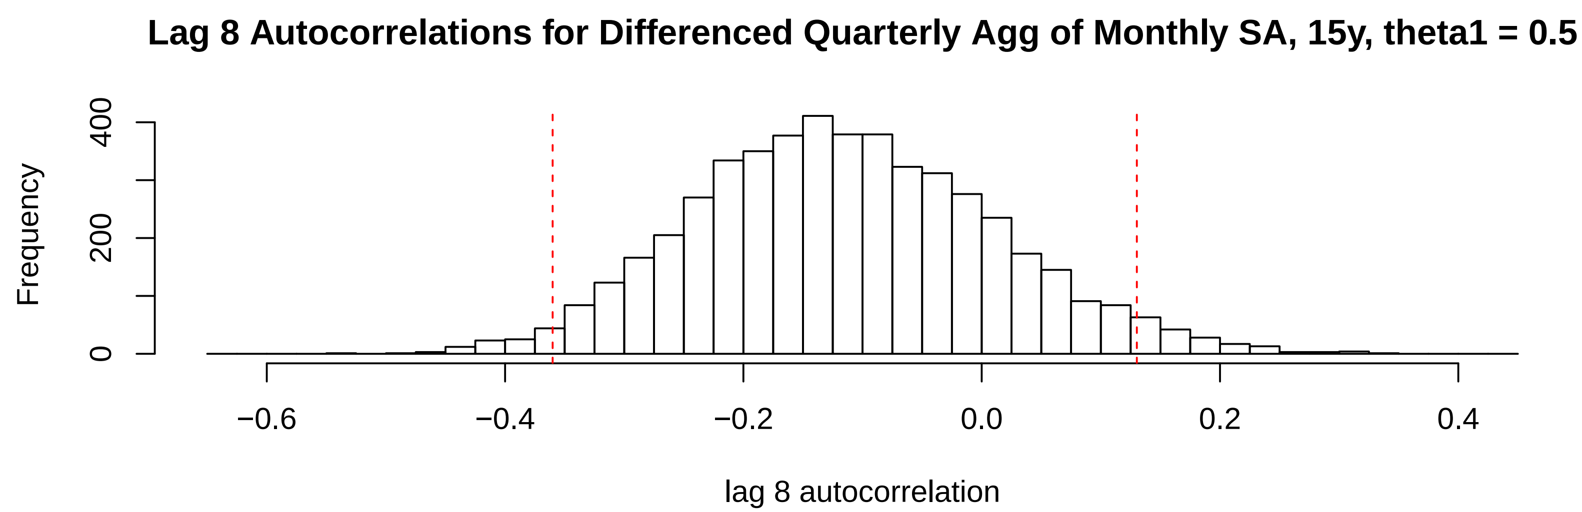

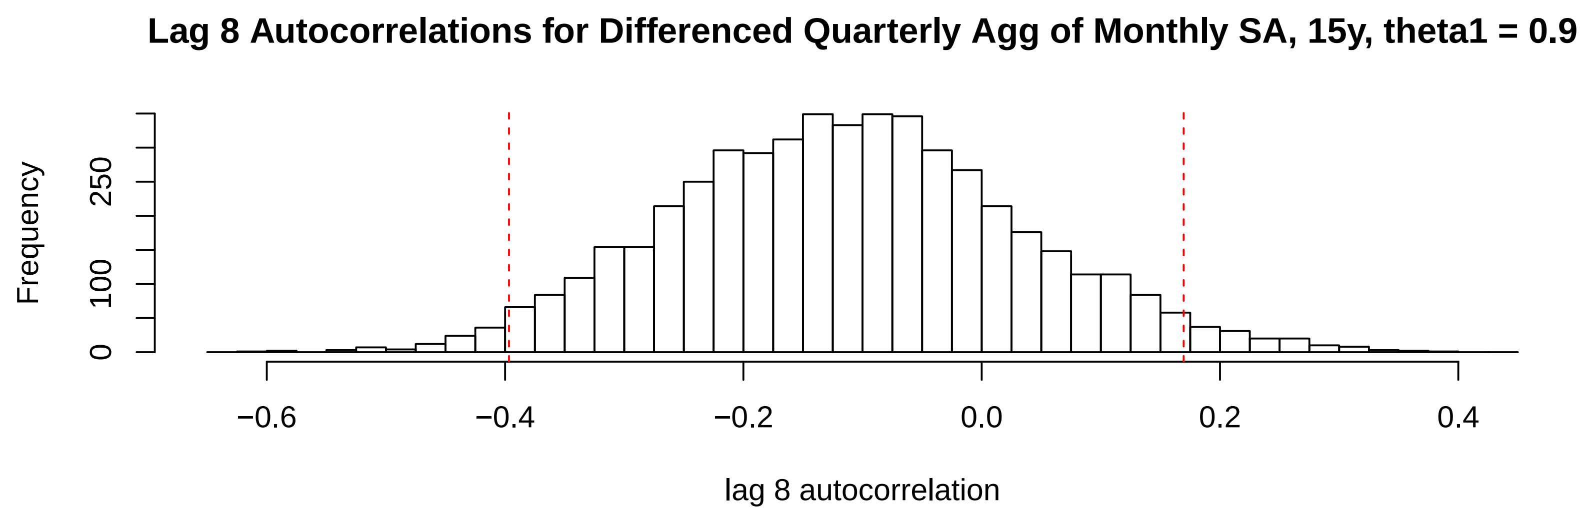

Thoughts and future directions

- Diagnostics that seem right for testing unadjusted data for seasonality do not work as well when testing adjusted data for seasonality.
- But thus far, we have primarily looked at nonseasonal models, so diagnostics that generally do not find seasonality in either the original series or a seasonal adjustment thereof may not be a bad thing.
- So it's worth checking the diagnostics for seasonal models and adjustments (both good and bad) of those.

Contact: osbert.c.pang@census.gov

United States Census Bureau