A Diagnostic for Seasonality Based Upon Autoregressive Roots

Tucker McElroy
U.S. Census Bureau

Seasonal Adjustment Practitioners Workshop, November 20, 2019

Disclaimer

These slides are released to inform interested parties of research and to encourage discussion. The views expressed on statistical issues are those of the author and not necessarily those of the U.S. Census Bureau.

Outline

1. Background on Seasonality

2. Criteria for a Diagnostic of Seasonality

3. Persistent Oscillations and Seasonality Testing

4. Simulation Evidence

5. Data Applications

Background on Seasonality

Seasonality in Official Time Series. Many official time series – such as gross domestic product (GDP) and unemployment rate data – have an enormous impact on public policy, and the seasonal patterns often obscure the long-run and mid-range dynamics.

What is Seasonality? Persistency in a time series over seasonal periods that is not explainable by intervening time periods.

- Requires persistence year to year

- Non-seasonal trending series have persistence, which comes through intervening seasons – we must screen out such cases

Background on Seasonality

The Seasonal Adjustment Task. Given a raw time series:

1. Does it have seasonality? If so, seasonally adjust.

2. Does the seasonal adjustment have seasonality? If not, publish.

Both these tasks require a seasonality diagnostic, although the properties of a time series before and after seasonal adjustment can be quite different.

Background on Seasonality

Pre-Testing. Testing for seasonality in a raw series, where the seasonality could be deterministic (stable), moving and stationary (dynamic), or moving and non-stationary (unit root). These categories are not mutually exclusive, e.g., we could have both unit root and deterministic seasonality.

Post-Testing. Testing for seasonality in a seasonally adjusted series, where the seasonality typically will only be dynamic. However, forecast-extension used in filtering introduces local non-stationarity to the beginning and end of the series.

Criteria for a Diagnostic of Seasonality

1. Rigorous statistical theory
2. Precise correspondence between seasonal dynamics and the diagnostic
3. Applicable to diverse sampling frequencies
4. Applicable to multiple frequencies of seasonal phenomena
5. Ability to assess over- and under-adjustment

Criteria for a Diagnostic of Seasonality

Correspondence. Diagnostic takes on a high value if and only if a high degree of seasonality is present (at a frequency of interest). We don't want high values occuring when seasonality is not present (spurious flagging of seasonality), or low values when seasonality is present (failure to detect).

Q_s**.** A diagnostic is based on autocorrelation at seasonal lags. These can take on high values for a non-seasonal AR(1) process, generating spurious indications of seasonality.

Criteria for a Diagnostic of Seasonality

Daily Time Series. A non-standard sampling frequency for official statistics. Also, there may be weekly seasonality (frequencies $2\pi j/7$ for $j = 1, 2, 3$) – corresponding to "trading day" – and annual seasonality (frequency $2\pi/365.25$ and integer multiples).

Autocorrelations. Only available at integer lags, hence not helpful for fractional periods like 365.25.

Criteria for a Diagnostic of Seasonality

Over- and Under-Adjustment. Too much seasonality removed (over-) versus too little (under-), described by Nerlove (1964) and others.

- Over-adjustment generates dips in the spectral density, corresponding to oscillatory effects in the inverse autocorrelations.

- Under-adjustment leaves peaks in the spectral density, corresponding to oscillatory effects in the autocorrelations.

Persistent Oscillations

Main Contribution. Associate persistence in a stationary time series to the presence of strong roots in its autoregressive polynomial.

1. Distribution theory for ARMA processes

2. High values of ρ (persistence) correspond to oscillatory effects in the Wold coefficients, and hence the autocorrelations

3. Adapts to any sampling or seasonal frequency (non-integer periods are fine)

4. Over-adjustment assessed through presence of strong roots in the moving average polynomial

Persistent Oscillations

The AR(∞) representation of a stationary process $\{X_t\}$ is

$$\pi(B)X_t = \epsilon_t \sim \text{WN}(0, \sigma^2),$$

and there are oscillatory effects of frequency ω and persistence ρ if

$$\pi(\rho^{-1}e^{i\omega}) = 0. \tag{1}$$

So a ρ-persistent seasonal effect is present if (1) holds and $\omega = \pi j/s$ for some $1 \leq j \leq s/2$ (the seasonal frequencies).

Persistent Oscillations

Estimate $\pi(B)$ via fitting an ARMA model (or an AR model, using OLS), and test statistic based on sample size T is

$$T \left| \widehat{\pi}(\rho^{-1} e^{i\omega}) \right|^2,$$

which has known asymptotic distribution (a squared Gaussian) under null hypothesis (1).

Testing: large values reject presence of seasonality of persistence ρ. BUT: a greater or lesser persistence of seasonality could still be present, so we should test across many values of ρ. We consider testing over all $\rho \in [.98, 1)$, with rejection indicating no seasonality of any persistency in this range.

Anti-Persistent Oscillations

The MA(∞) representation of a stationary process $\{X_t\}$ is

$$X_t = \psi(B)\epsilon_t \sim \text{WN}(0, \sigma^2),$$

and there are oscillatory effects of frequency ω and anti-persistence ρ if

$$\psi(\rho^{-1} e^{i\omega}) = 0. \tag{2}$$

A ρ-persistent anti-seasonal effect is present if (2) holds and $\omega = \pi j / s$ for some $1 \leq j \leq s/2$ (the seasonal frequencies).

Anti-Persistent Oscillations

Estimate $f(z) = \sum_{|h| \leq q} \gamma_h \, z^h$ with $z = \rho^{-1} e^{i\omega}$ for some q by using sample autocovariances. The test statistic based on sample size T is

$$T \, \left| \widehat{f}(\rho^{-1} e^{i\omega}) \right|^2,$$

which has known asymptotic distribution (a squared Gaussian) under null hypothesis (2).

Testing: large values reject presence of anti-seasonality of persistence ρ. We consider testing over all $\rho \in [.98, 1)$, with rejection indicating no anti-seasonality of any persistency in this range.

Seasonality Hypothesis Testing

ARMA Summary. So for an invertible ARMA process with $\psi(z) = \theta(z)/\phi(z)$ and $\pi(z) = \phi(z)/\theta(z)$, the AR roots govern oscillations in $\{\psi_j\}$ and autocorrelations, whereas the MA roots govern oscillations in $\{\pi_j\}$ and the inverse autocorrelations.

Implementation Notes.

- May be easier to fit a high order AR model (sieve approach) using OLS
- Identify AR order using BIC (we found that AIC leads to mis-sized results)
- Use estimated parameters to obtain null limit distribution

Seasonality Hypothesis Testing

Testing Procedure. Focus on post-test (for seasonally adjusted data):

1. Remove the first and last few years of data, so as to remove local non-stationarity

2. Fit an invertible ARIMA model, and obtain the AR(∞) representation of the differenced process as $\pi(z) = \phi(z)/\theta(z)$

3. For any given ω, test $H_0(\rho)$ for all $\rho \in (0,1)$ at level α

4. Obtain interval $C(\alpha)$ consisting of all ρ for which we failed to reject

Seasonality Hypothesis Testing

Is it Seasonal? An interval $C(\alpha)$ is obtained for each ω of interest. Seasonality exists if for at least one ω corresponding to a seasonal frequency, an interval contains $\rho = .98$ (this value is suggested by other studies, but can be modified if desired).

Simulation Evidence

Simulated Processes. We study Gaussian time series generated from

$$(1 - \phi\, B)\,(1 - 2\rho\, \cos(\pi/6)\, B + \rho\, B^2) X_t \sim \mathsf{WN}(0, \sigma^2). \qquad (3)$$

The autocovariance function and spectrum are plotted in Figure 1, where we have set $\phi = .8$ and $\rho = .9$, and $\sigma = 1$. From the plots, it is apparent that the moderate seasonality ($\rho = .9$) is somewhat attenuated by the transient effect ($\phi = .8$), so the impact of the atomic seasonality is weaker than it would be if $\phi = 0$.

Second Example. Lower the seasonal persistency to $\rho = .8$, and dampen the transient component by setting $\phi = .3$, displayed in Figure 2.

Simulation Evidence

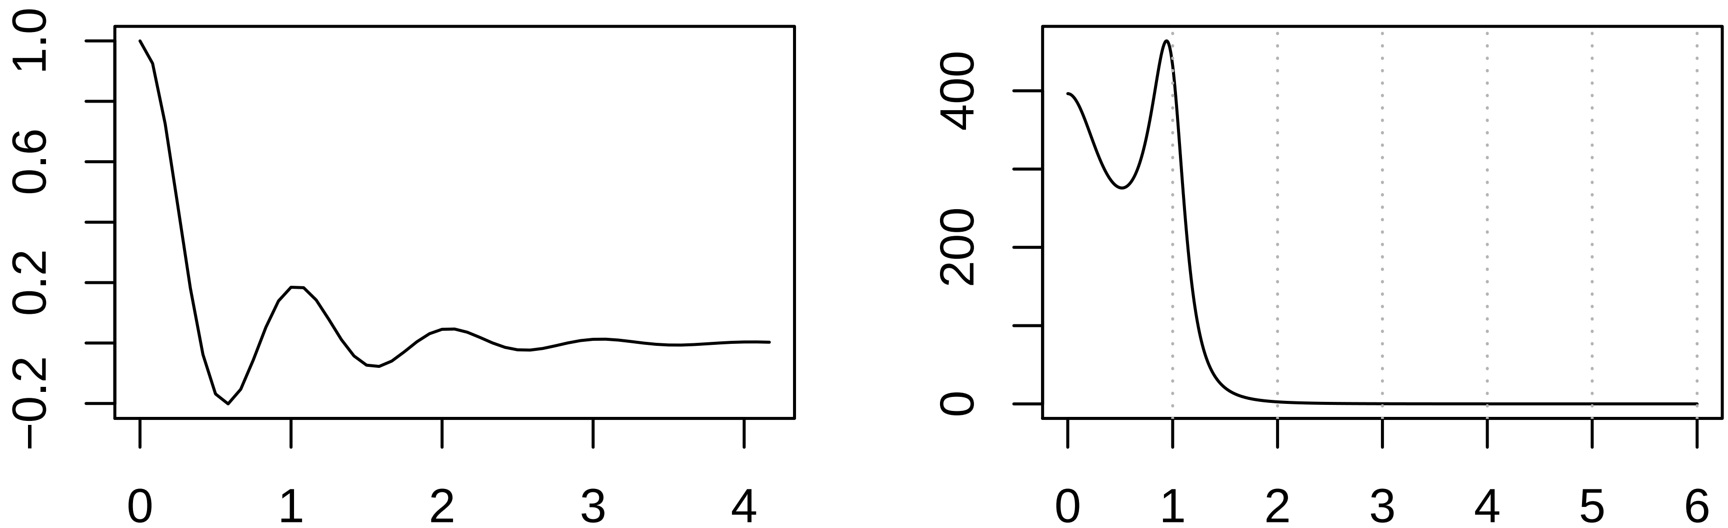

Figure 1: Autocorrelation function (left panel) and spectral density (right panel) for AR(3) process ($\phi = .8$, $\rho = .9$).

Simulation Evidence

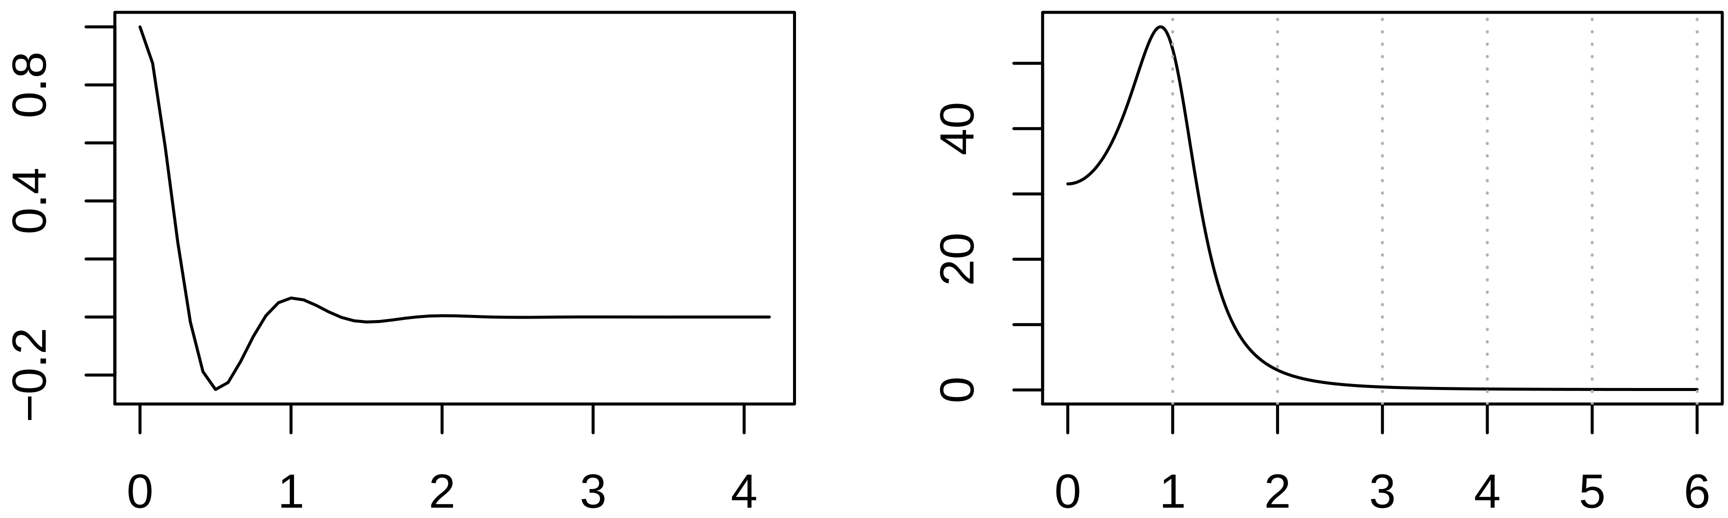

Figure 2: Autocorrelation function (left panel) and spectral density (right panel) for AR(3) process ($\phi = .3$, $\rho = .8$).

Simulation Evidence

Simulations. Both processes generated with sample size $T = 12n$ and $n = 5, 10, 15, 20$. We take null with $\omega = \pi/6$ and either $\rho = .9, .8$, with either $p = 3$ is known or selected via BIC.

- Size: $\rho = .9$ for first process (Table 1) or $\rho = .8$ for second process (Table 2)

- Power: $\rho = .8$ for first process (Table 3) or $\rho = .9$ for second process (Table 4)

Simulation Evidence

α	5 years	10 years	15 years	20 years
.10	.149	.116	.113	.101
.05	.092	.067	.061	.054
.01	.039	.026	.021	.018
.10	.139	.116	.108	.109
.05	.080	.061	.057	.056
.01	.027	.014	.013	.012

Table 1: Size simulations from an AR(3) DGP (corresponding to Figure 1) based on a null hypothesis of .9-persistent seasonality at frequency $\pi/6$. Results are for known AR order (first three rows) and unknown AR order (last three rows).

Simulation Evidence

α	5 years	10 years	15 years	20 years
.10	.135	.115	.109	.104
.05	.077	.058	.057	.053
.01	.021	.012	.013	.011
.10	.132	.113	.108	.107
.05	.073	.059	.053	.054
.01	.020	.014	.009	.012

Table 2: Size simulations from an AR(3) DGP (corresponding to Figure 2) based on a null hypothesis of $.8$-persistent seasonality at frequency $\pi/6$. Results are for known AR order (first three rows) and unknown AR order (last three rows).

Simulation Evidence

α	5 years	10 years	15 years	20 years
.10	.832	.937	.983	.994
.05	.757	.907	.970	.989
.01	.564	.805	.923	.971
.10	.477	.700	.856	.930
.05	.336	.566	.755	.872
.01	.111	.277	.484	.665

Table 3: Power simulations from an AR(3) DGP (corresponding to Figure 1) with null hypothesis of $.8$-persistent seasonality at frequency $\pi/6$. Results are for known AR order (first three rows) and unknown AR order (last three rows).

Simulation Evidence

α	5 years	10 years	15 years	20 years
.10	.116	.205	.341	.473
.05	.060	.110	.203	.313
.01	.015	.028	.057	.096
.10	.281	.498	.682	.812
.05	.175	.362	.550	.702
.01	.051	.148	.294	.448

Table 4: Power simulations from an AR(3) DGP (corresponding to Figure 2) with null hypothesis of $.9$-persistent seasonality at frequency $\pi/6$. Results are for known AR order (first three rows) and unknown AR order (last three rows).

Simulation Evidence

Summary.

- Size is adequate for 10 years of data when p is known or estimated, with coverage being similar for both cases
- Power was good in the case of the first process, though greatly reduced if the model order was unknown
- Power is much lower with second process, although somewhat greater when the model order was unknown

NZ immigration

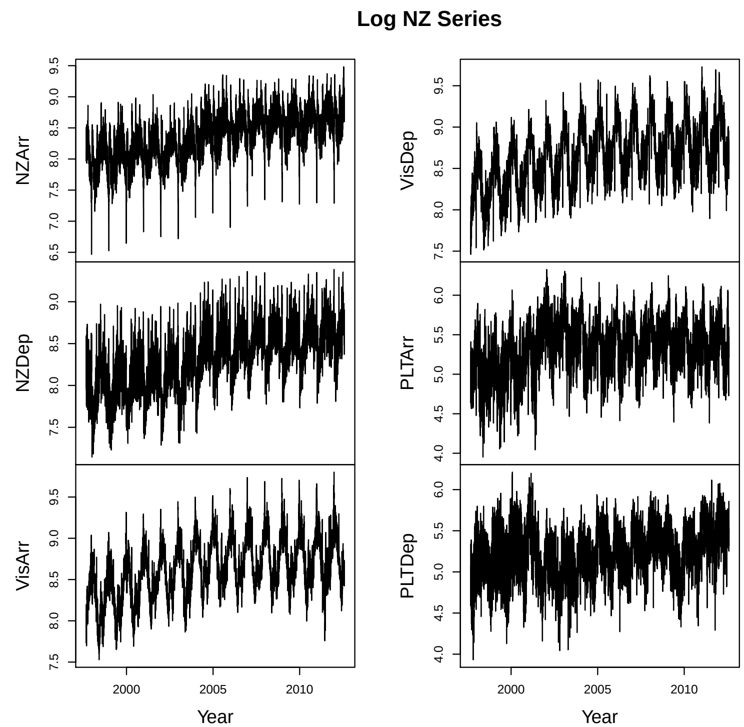

Figure 3: Log of six New Zealand immigration series (9-1-1997 through 7-31-2012), showing trend, seasonal, and weekly effects.

NZ immigration

We apply the seasonality diagnostics to the differenced logged data for frequencies $\omega = 2\pi j/7$ for $1 \leq j \leq 3$ and $\omega = 2\pi/365.25$. An AR($\widehat{p}$) model is fitted with $\widehat{p}$ selected by AIC (see Table 5), along with the identified intervals for ρ (all values between $.5$ and 1 such that the p-value exceeds $.01$).

We find that high order AR processes are needed (due to the autocorrelation present at an annual period), and that strong seasonality is present at both annual and weekly frequencies for all six series (with the exception of the first weekly seasonal for the VisDep series) .

NZ immigration

Series	$\widehat{p}$	$2\pi/365.25$	$2\pi/7$	$4\pi/7$	$6\pi/7$
NZArr	447	$\{1\}$	$\{1\}$	$[.998, 1]$	$[.998, 1]$
NZDep	405	$\{1\}$	$\{1\}$	$[.998, 1]$	$[.998, 1]$
VisArr	391	$\{1\}$	$\{1\}$	$[.997, .999]$	$[.999, 1]$
VisDep	404	$\{1\}$	$\emptyset$	$[.999, 1]$	$[.999, 1]$
PLTArr	391	$[.999, 1]$	$\{1\}$	$[.999, 1]$	$[.999, 1]$
PLTDep	398	$[.999, 1]$	$[.999, 1]$	$[.999, 1]$	$[.999, 1]$

Table 5: Intervals for ρ, such that the corresponding null hypothesis is not rejected at a 1% level. Rows correspond to each of the six component series.

Census Data

We examined a collection of 233 monthly time series published by the U.S. Census Bureau:

- 65 time series of Retail Trade and Food Services (MRTS)
- 22 time series of Wholesale Trade: Sales and Inventories (MWTS)
- 4 time series of Manufacturers' Shipments, Inventories, and Orders (M3)
- 87 time series of Manufacturing and Trade Inventories and Sales (MTIS)
- 55 time series of New Residential Construction (RES)

Census Data

All are monthly with a start date of January 1992 or later, and with end date of September 2019. A variety of features are present in these series: varying degrees of persistence and evolution in seasonal patterns; presence of outliers and calendrical effects; varying degrees of aggregation.

We extract seasonal adjustments and irregular components by application of the software X-13ARIMA-SEATS, using either the X11 option or the SEATS option. Then we apply the Q_s statistic as well as the Root diagnostic for under-adjustment and over-adjustment. For the latter two diagnostics, we set the values of ρ between $.98$ and 1. We tally for each of the five batches of series the incidences of adequacy, presented in Tables 6 and 7.

Census Data

X11 Test	MRTS	MWTS	M3	MTIS	RES
SA Q_s	63	22	86	4	53
SA Under	64	22	87	4	55
SA Over	3	2	4	1	1
Irr Q_s	63	22	86	4	55
Irr Under	64	22	87	4	55
Irr Over	4	0	4	0	1
Total	65	22	87	4	55

Table 6: Number of series seasonally adjusted by X11 that are deemed to be adequate, according to whether the Q_s test, the Root test of under-adjustment, or the Root test for over-adjustment is applied, for either the Seasonally Adjusted (SA) component or the Irregular (Irr).

Census Data

SEATS Test	MRTS	MWTS	MTIS	M3	RES
SA Q_s	64	22	85	4	52
SA Under	65	22	87	4	55
SA Over	3	1	2	0	2
Irr Q_s	65	22	86	4	55
Irr Under	65	22	87	4	55
Irr Over	6	1	2	0	2
Total	65	22	87	4	55

Table 7: Number of series seasonally adjusted by SEATS that are deemed to be adequate, according to whether the Q_s test, the Root test of under-adjustment, or the Root test for over-adjustment is applied, for either the Seasonally Adjusted (SA) component or the Irregular (Irr).

Conclusion

- A new paradigm for assessing cyclicality and seasonality is introduced, where oscillations in the Wold coefficients (and autocorrelations) are measured through AR root magnitudes
- This concept addresses several of the criteria set forth, and is demonstrated in simulation to be promising as a test of residual seasonality
- Further work: need theory for case of unit-root AR polynomials, to allow for testing raw data for seasonality

Email: tucker.s.mcelroy@census.gov