Times Series Analysis (XVI) – Seasonal Models

Jianyong Sun
School of Mathematics and Statistics
Xi'an Jiaotong University

2nd Nov., 2017

Prefix

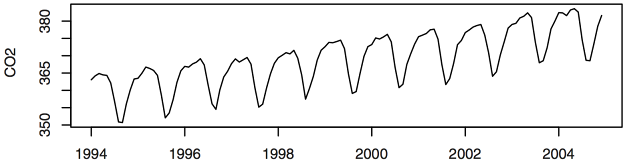

- we saw how seasonal deterministic trends might be modeled. However, in many areas in which time series are used, particularly business and economics, the assumption of any deterministic trend is quite suspect even though cyclical tendencies are very common in such series.
- This figure displays the monthly CO2 levels from January 1994 through December 2004.

Prefix

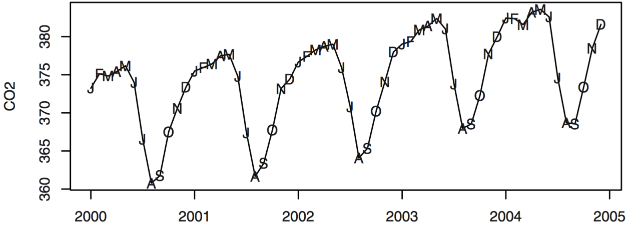

- carbon dioxide levels are higher during the winter months and much lower in the summer.
- Deterministic seasonal models such as seasonal means plus linear time trend or sums of cosine curves at various frequencies plus linear time trend could certainly be considered here.

Prefix

- we discover that such models do not explain the behavior of this time series. For this series and many others, it can be shown that the residuals from a seasonal means plus linear time trend model are highly autocorrelated at many lags.
- In contrast, we will see that the **stochastic seasonal models** developed in this chapter do work well for this series.

We begin by studying stationary models and then consider nonstationary generalizations. Let s denote the known seasonal period; for monthly series $s = 12$ and for quarterly series $s = 4$.

Consider the time series generated according to

$$Y_t = e_t - \Theta e_{t-12}$$

Notice that

$$Cov(Y_t, Y_{t-1}) = Cov(e_t - \Theta e_{t-12}, e_{t-1} - \Theta e_{t-13}) = 0$$

but that

$$Cov(Y_t, Y_{t-12}) = Cov(e_t - \Theta e_{t-12}, e_{t-12} - \Theta e_{t-24}) = -\Theta\sigma_e^2$$

It is easy to see that such a series is stationary and has nonzero autocorrelations only at lag 12.

Generalizing these ideas, we define a **seasonal MA(Q) model of order Q with seasonal period s** by

$$Y_t = e_t - \Theta_1 e_{t-s} - \Theta_2 e_{t-2s} - \cdots - \Theta_Q e_{t-Qs}$$

with **seasonal MA characteristic polynomial**:

$$\Theta(x) = 1 - \Theta_1 x^s - \Theta_2 x^{2s} - \cdots - \Theta_Q x^{Qs}$$

It is evident that such a series is always stationary and that the autocorrelation function will be nonzero only at the seasonal lags of $s, 2s, 3s, \cdots, Qs$.

In particular,

$$\rho_{ks} = \frac{-\Theta_k + \Theta_1\Theta_{k+1} + \Theta_2\Theta_{k+2} + \cdots + \Theta_{Q-k}\Theta_Q}{1 + \Theta_1^2 + \Theta_2^2 + \cdots + \Theta_Q^2}$$

for $k = 1, 2, \cdots, Q$.

For the model to be invertible, the roots of $\Theta(x) = 0$ must all exceed 1 in absolute value.

It is useful to note that the seasonal MA(Q) model can also be viewed as a special case of a nonseasonal MA model of order $q = Qs$ but with all θ-values zero except at the seasonal lags $s, 2s, 3s, \cdots, Qs$.

Seasonal ARIMA Models

Seasonal autoregressive models can also be defined. Consider

$$Y_t = \Phi Y_{t-12} + e_t$$

where $|\Phi| < 1$ and e_t is independent of $Y_{t-1}, Y_{t-2}, \cdots$. It can be shown that $|\Phi| < 1$ ensures stationarity. Thus it is easy to argue that $E(Y_t) = 0$; multiplying the equation by Y_{t-k} ,taking expectations, and dividing by γ_0 yields

$$\rho_k = \Phi \rho_{k-12} \text{ for } \ k \geq 1$$

Clearly,

$$\rho_{12} = \Phi \rho_0 = \Phi \qquad \text{and} \qquad \rho_{24} = \Phi \rho_{12} = \Phi^2$$

More generally,

$$\rho_{12k} = \Phi^k \text{ for } \ k = 1, 2, \cdots$$

Furthermore, setting $k = 1$ and then $k = 11$ and using $\rho_k = \rho_{-k}$ gives us

$$\rho_1 = \Phi\rho_{11} \qquad \text{and} \quad \rho_{11} = \Phi\rho_1$$

which implies that $\rho_1 = \rho_{11} = 0$. Similarly, one can show that $\rho_k = 0$ except at the seasonal lags 12, 24, 36, $\cdots$. At those lags, the autocorrelation function decays exponentially like an AR(1) model.

With this example in mind, we define a **seasonal AR(P) model of order P and seasonal period s** by

$$Y_t = \Phi_1 Y_{t-s} + \Phi_2 Y_{t-2s} + \cdots + \Phi_P Y_{t-Ps} + e_t$$

with **seasonal characteristic polynomial**:

$$\Phi(x) = 1 - \Phi_1 x^s - \Phi_2 x^{2s} - \cdots - \Phi_P x^{Ps}$$

As always, we require e_t to be independent of $Y_{t-1}, Y_{t-2}, \cdots$, and, for stationarity, that the roots of $\Phi(x) = 0$ be greater than 1 in absolute value.

Again, this equation can be seen as a special AR(p) model of order $p = Ps$ with nonzero Φ-coefficients only at the seasonal lags $s, 2s, 3s, ..., Ps$.

It can be shown that the autocorrelation function is nonzero only at lags $s, 2s, 3s, \cdots$, where it behaves like a combination of decaying exponentials and damped sine functions. In particular, previous equations easily generalize to the general seasonal AR(1) model to give

$$\rho_{ks} = \Phi^k \text{ for } \ k = 1, 2, \cdots$$

with zero correlation at other lags.

Multiplicative Seasonal ARMA Models

Rarely shall we need models that incorporate autocorrelation only at the seasonal lags. By combining the ideas of seasonal and nonseasonal ARMA models, we can develop parsimonious models that contain autocorrelation for the seasonal lags but also for low lags of neighboring series values.

Consider a model whose MA characteristic polynomial is given by

$$(1-\theta x)(1-\Theta x^{12}) = 1-\theta x-\Theta x^{12}+\theta\Theta x^{13}$$

Thus the corresponding time series satisfies

$$Y_t = e_t - \theta e_{t-1} - \Theta e_{t-12} + \theta\Theta e_{t-13}$$

For this model, we can check that the autocorrelation function is nonzero only at lags 1, 11, 12, and 13. We find

$$
\begin{aligned}
\gamma_0 &= (1+\theta^2)(1+\Theta^2)\sigma_e^2 \\
\rho_1 &= -\frac{\theta}{1+\theta^2} \\
\rho_{11} = \rho_{13} &= \frac{\theta\Theta}{(1+\theta^2)(1+\Theta^2)}, \text{and} \\
\rho_{12} &= -\frac{\Theta}{1+\Theta^2}
\end{aligned}
$$

Multiplicative Seasonal ARMA Models

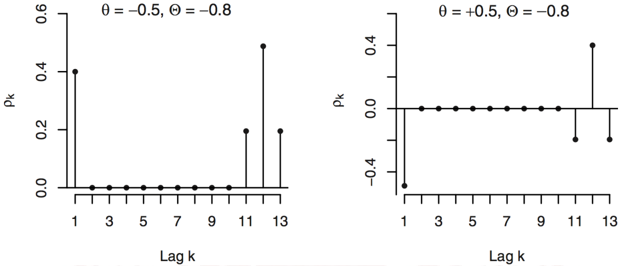

Figure: Autocorrelation from previous equations

Of course, we could also introduce both short-term and seasonal autocorrelations by defining an MA model of order 12 with only θ_1 and θ_{12} nonzero

Multiplicative Seasonal ARMA Models

In general, then, we define a **multiplicative seasonal ARMA(p,q)$\times$(P,Q)$_s$ model with seasonal period** s as a model with AR characteristic polynomial $\phi(x)\Phi(x)$ and MA characteristic polynomial $\theta(x)\Theta(x)$, where

$$\begin{aligned}
\phi(x) &= 1 - \phi_1 x - \phi_2 x^2 - \cdots - \phi_p x^p \\
\Phi(x) &= 1 - \Phi_1 x^s - \Phi_2 x^{2s} - \cdots - \Phi_P x^{Ps}
\end{aligned}$$

and

$$\begin{aligned}
\theta(x) &= 1 - \theta_1 x - \theta_2 x^2 - \cdots - \theta_q x^q \\
\Theta(x) &= 1 - \Theta_1 x^s - \Theta_2 x^{2s} - \cdots - \Theta_Q x^{Qs}
\end{aligned}$$

Multiplicative Seasonal ARMA Models

The model may also contain a constant term θ_0. Note once more that we have just a special ARMA model with AR order $p + Ps$ and MA order $q + Qs$, but the coefficients are not completely general, being determined by only $p + P + q + Q$ coefficients.

If $s = 12, p + P + q + Q$ will be considerably smaller than $p + Ps + q + Qs$ and will allow a much more parsimonious model.

As another example, suppose $P = q = 1$ and $p = Q = 0$ with $s = 12$. The model is then

$$Y_t = \Phi Y_{t-12} + e_t - \theta e_{t-1}$$

Using our standard techniques, we find that

$$\begin{aligned} \gamma_1 &= \Phi \gamma_{11} - \theta \sigma_e^2 \\ \gamma_k &= \Phi \gamma_{k-12} \qquad \text{for} \quad k \geq 2 \end{aligned}$$

After considering the equations implied by various choices for k, we arrive at

$$\begin{aligned} \gamma_0 &= \frac{1+\theta^2}{1-\Phi^2}\sigma_e^2 \\ \rho_{12k} &= \Phi^k \text{ for } k \geq 1 \\ \rho_{12k-1} &= \rho_{12k+1} = \left(-\frac{\theta}{1+\theta^2}\Phi^k\right) \text{ for } k = 0, 1, 2, \cdots, \end{aligned}$$

with autocorrelations for all other lags equal to zero.

Multiplicative Seasonal ARMA Models

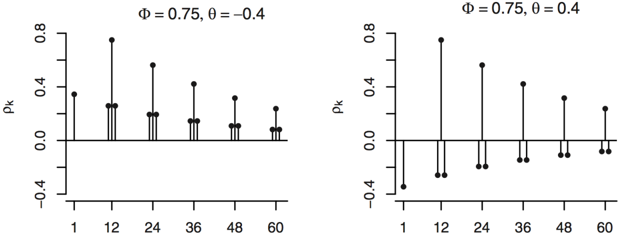

Figure: Autocorrelation from previous equations

- This figure displays the autocorrelation functions for two of these seasonal ARIMA processes with period 12: one with $\Phi = 0.75$ and $\theta = 0.4$, the other with $\Phi = 0.75$ and $\theta = -0.4$.
- The shape of these autocorrelations is somewhat typical of the sample autocorrelation functions for numerous seasonal time series. The even simpler autocorrelation function given by previous equations and displayed in previous figure also seems to occur frequently in practice (perhaps after differencing).

Nonstationary Seasonal ARIMA Models

An important tool in modeling nonstationary seasonal processes is the seasonal difference. The **seasonal difference** of period s for the series $\{Y_t\}$ is denoted $\nabla_s Y_t$ and is defined as

$$\nabla_s Y_t = Y_t - Y_{t-s}$$

For example, for monthly series we consider the changes from January to January, February to February, and so forth for successive years. Note that for a series of length n, the seasonal difference series will be of length $n - s$; that is, s data values are lost due to seasonal differencing.

Nonstationary Seasonal ARIMA Models

As an example where seasonal differencing is appropriate, consider a process generated according to

$$Y_t = S_t + e_t$$

with

$$S_t = S_{t-s} + \epsilon_t$$

where $\{e_t\}$ and $\{\epsilon_t\}$ are independent white noise series. Here $\{S_t\}$ is a "seasonal random walk", and if $\sigma_\epsilon \ll \sigma_e$, $\{S_t\}$ would model a slowly changing seasonal component.

Nonstationary Seasonal ARIMA Models

Due to the nonstationarity of $\{S_t\}$, clearly $\{Y_t\}$ is nonstationary. However, if we seasonally difference $\{Y_t\}$, we find

$$\nabla_s Y_t = S_t - S_{t-s} + e_t - e_{t-s} = \epsilon_t + e_t - e_{t-s}$$

An easy calculation shows that $\nabla_s Y_t$ is stationary and has the autocorrelation function of an MA$(1)_s$ model.

The model described by previous equations could also be generalized to account for a nonseasonal, slowly changing stochastic trend. Consider

$$\begin{aligned} Y_t &= M_t + S_t + e_t \qquad \text{with} \\ S_t &= S_{t-s} + \epsilon_t \qquad \text{and} \\ M_t &= M_{t-1} + \xi_t \end{aligned}$$

where $\{e_t\}, \{\epsilon_t\}, \{\xi_t\}$ are mutually independent white noise series.

Here we take both a seasonal difference and an ordinary nonseasonal difference to obtain (It should be noted that $\nabla_s Y_t$ will in fact be stationary and $\nabla\nabla_s Y_t$ will be noninvertible. We use previous equations merely to help motivate multiplicative seasonal ARIMA models.)

$$\begin{aligned}\nabla\nabla_s Y_t &= \nabla(M_t - M_{t-s} + \epsilon_t + e_t - e_{t-s}) \\ &= (\xi_t + \epsilon_t + e_t) - (\epsilon_{t-1} + e_{t-1}) - (\xi_{t-s} + e_{t-s}) + e_{t-s-1}\end{aligned}$$

The process defined here is stationary and has nonzero autocorrelation only at lags $1, s-1, s$, and $s+1$, which agrees with the autocorrelation structure of the multiplicative seasonal model ARMA(0,1)$\times$(0,1) with seasonal period s.

Nonstationary Seasonal ARIMA Models

These examples lead to the definition of nonstationary seasonal models. A process $\{Y_t\}$ is said to be a **multiplicative seasonal ARIMA model** with nonseasonal (regular) orders p, d, and q, seasonal orders P, D, and Q, and seasonal period s if the differenced series:

$$W_t = \nabla^d \nabla_s^D Y_t$$

satisfies an ARMA$(p, q) \times (P, Q)_s$ model with seasonal period s. We say that $\{Y_t\}$ is an ARIMA$(p, d, q) \times (P, D, Q)_s$ model with seasonal period s.

Clearly, such models represent a broad, flexible class from which to select an appropriate model for a particular time series. It has been found empirically that many series can be adequately fit by these models, usually with a small number of parameters, say three or four.

we shall simply highlight the application of these ideas specifically to seasonal models and pay special attention to the seasonal lags. As always, a careful inspection of the time series plot is the first step.

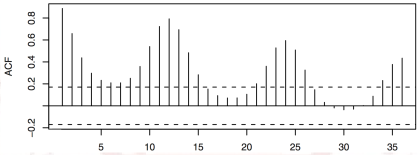

Figure: Sample ACF of CO2 Levels

This figure shows the sample autocorrelation function for that series. The seasonal autocorrelation relationships are shown quite prominently in this display. Notice the strong correlation at lags 12, 24, 36, and so on. In addition, there is substantial other correlation that needs to be modeled.

Model Specification, Fitting, and Checking

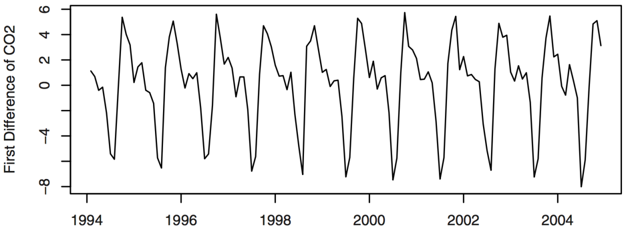

Figure: Time Series Plot of the First Differences of CO2 Levels

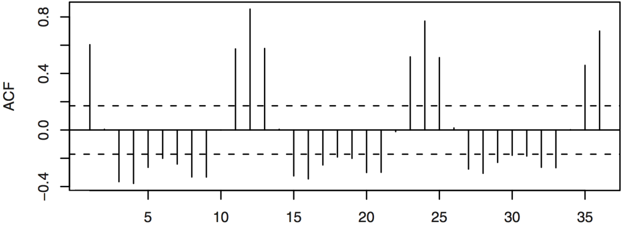

Figure: Sample ACF of First Differences of CO2 Levels

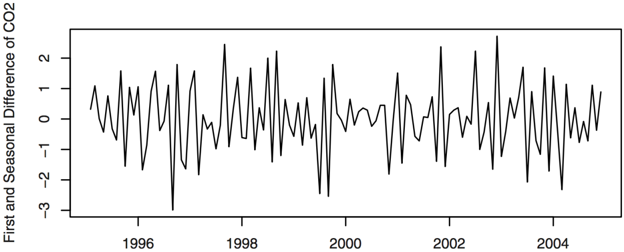

Figure: Time Series Plot of First and Seasonal Differences of CO2.

The figure displays the time series plot of the CO2 levels after taking both a first difference and a seasonal difference. It appears that most, if not all, of the seasonality is gone now.

Model Specification, Fitting, and Checking

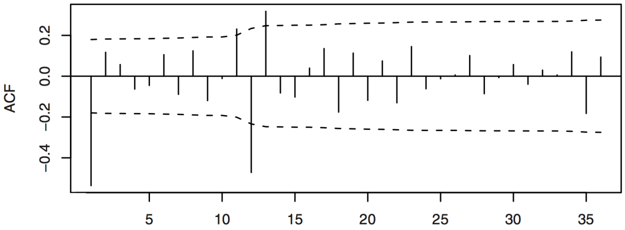

Figure: Sample ACF of First and Seasonal Differences of CO2.

This figure confirms that very little autocorrelation remains in the series after these two differences have been taken. This plot also suggests that a simple model which incorporates the lag 1 and lag 12 autocorrelations might be adequate.

We will consider specifying the multiplicative, seasonal ARIMA(0,1,1)$\times(0,1,1)_{12}$ model:

$$\nabla_{12}\nabla Y_t = e_t - \theta e_{t-1} - \Theta e_{t-12} + \theta\Theta e_{t-13}$$

which incorporates many of these requirements. As usual, all models are tentative and subject to revision at the diagnostics stage of model building.

Model Fitting

- Having specified a tentative seasonal model for a particular time series, we proceed to estimate the parameters of that model as efficiently as possible.
- As we have remarked earlier, multiplicative seasonal ARIMA models are just special cases of our general ARIMA models.
- As such, all of our work on parameter estimation carries over to the seasonal case.

Model Fitting

Coefficient	θ	Θ
Estimate	0.5792	0.8206
Standard error	0.0791	0.1137

$\hat{\sigma}_e^2 = 0.5446$: log-likelihood $= -139.54$, AIC $= 283.08$

Figure: Parameter Estimates for the CO2 Model.

This table gives the maximum likelihood estimates and their standard errors for the ARIMA(0,1,1)$\times(0,1,1)_{12}$ model for CO2 levels.

The coefficient estimates are all highly significant, and we proceed to check further on this model.

Diagnostic Checking

To check the estimated the ARIMA(0,1,1)$\times(0,1,1)_{12}$ model, we first look at the time series plot of the residuals.

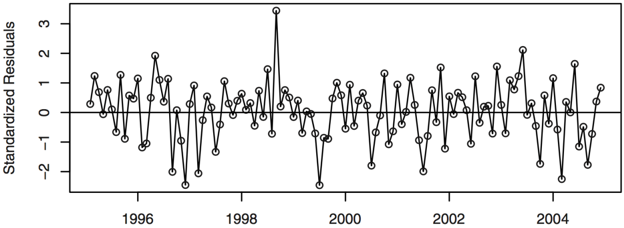

Figure: Residuals from the ARIMA(0,1,1)$\times(0,1,1)_{12}$ Model.

Th gives this plot for standardized residuals. Other than some strange behavior in the middle of the series, this plot does not suggest any major irregularities with the model, although we may need to investigate the model further for outliers, as the standardized residual at September 1998 looks suspicious.

Diagnostic Checking

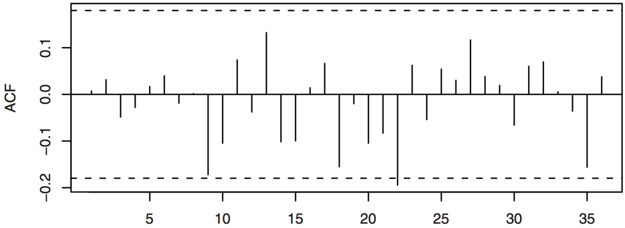

Figure: ACF of Residuals from the ARIMA(0,1,1)$\times(0, 1, 1)_{12}$ Model.

- The only "statistically significant" correlation is at lag 22, and this correlation has a value of only -0.17, a very small correlation.
- Furthermore, we can think of no reasonable interpretation for dependence at lag 22.
- Finally, we should not be surprised that one autocorrelation out of the 36 displayed is statistically significant. This could easily happen by chance alone.

The Ljung-Box test for this model gives a chi-squared value of 25.59 with 22 degrees of freedom, leading to a p-value of 0.27 — a further indication that the model has captured the dependence in the time series.

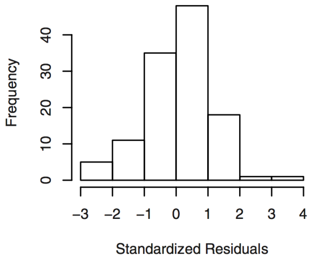

Figure: Residuals from the ARIMA(0,1,1)$\times(0,1,1)_{12}$ Model.

The figure displays the histogram of the residuals. The shape is somewhat "bell-shaped" but certainly not ideal. Perhaps a quantile-quantile plot will tell us more.

Diagnostic Checking

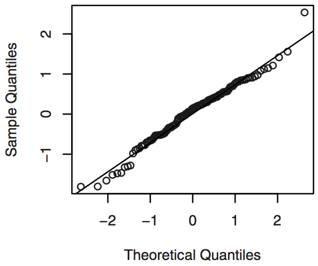

Figure: Residuals: ARIMA(0,1,1)$\times(0,1,1)_{12}$ Model.

Here we again see the one outlier in the upper tail, but the Shapiro-Wilk test of normality has a test statistic of $W = 0.982$, leading to a p-value of 0.11, and normality is not rejected at any of the usual significance levels.

Diagnostic Checking

As one further check on the model, we consider overfitting with an ARIMA(0,1,2)$\times(0,1,1)_{12}$ model.

Coefficient	θ_1	θ_2	Θ
Estimate	0.5714	0.0165	0.8274
Standard error	0.0897	0.0948	0.1224

$\hat{\sigma}_e^2 = 0.5427$: log-likelihood $= -139.52$, AIC $= 285.05$

Figure: ARIMA(0,1,1)$\times(0,1,2)_{12}$ Overfitted Model.

we see that the estimates of θ_1 and Θ have changed very little — especially when the size of the standard errors is taken into consideration. In addition, the estimate of the new parameter, θ_2, is not statistically different from zero. Note also that the estimate σ_e^2 and the log-likelihood have not changed much while the AIC has actually increased.

Diagnostic Checking

The ARIMA(0,1,1)$\times(0,1,1)_{12}$ model was popularized in the first edition of the seminal book of Box and Jenkins (1976) when it was found to characterize the logarithms of a monthly airline passenger time series. This model has come to be known as the **airline model**.

Forecasting Seasonal Models

Computing forecasts with seasonal ARIMA models is, as expected, most easily carried out recursively using the difference equation form for the model. For example, consider the model ARIMA(0,1,1)$\times(1,0,1)_{12}$.

$$Y_t - Y_{t-1} = \Phi(Y_{t-12} - Y_{t-13}) + e_t - \theta e_{t-1} - \Theta e_{t-12} + \theta\Theta e_{t-13}$$

which we rewrite as

$$Y_t = Y_{t-1} + Y_{t-12} - \Phi Y_{t-13} + e_t - \theta e_{t-1} - \Theta e_{t-12} + \theta\Theta e_{t-13}$$

The one-step-ahead forecast from origin t is then

$$\hat{Y}_t(1) = Y_t + \Phi Y_{t-11} - \Phi Y_{t-12} - \theta e_t - \Theta e_{t-11} + \theta\Theta e_{t-12}$$

and the next one is

$$\hat{Y}_t(2) = \hat{Y}_t(1) + \Phi Y_{t-10} - \Phi Y_{t-11} - \Theta e_{t-10} + \theta\Theta e_{t-11}$$

ans so forth. The noise terms $e_{t-13}, e_{t-12}, e_{t-11}, \cdots, e_t$ (as residuals) will enter into the forecasts for lead times $\ell = 1, 2, \cdots, 13$, but for $\ell > 13$ the autoregressive part of the model takes over and we have

$$\hat{Y}_t(\ell) = \hat{Y}_t(\ell - 1) + \Phi\hat{Y}_t(\ell - 12) - \Phi\hat{Y}_t(\ell - 13) \text{ for } \ell > 13$$

To understand the general nature of the forecasts, we consider several special cases.

Seasonal AR(1)$_{12}$:

The seasonal AR(1)$_{12}$ model is $Y_t = \Phi Y_{t-12} + e_t$. Clearly, we have

$$\hat{Y}_t(\ell) = \Phi \hat{Y}_t(\ell - 12)$$

However, iterating back on ℓ, we can also write

$$\hat{Y}_t(\ell) = \Phi^{k+1} Y_{t+r-11}$$

where k and r are defined by $\ell = 12k + r + 1$ with $0 \le r < 12$ and $k = 0, 1, 2, \cdots$. In other words, k is the integer part of $(\ell - 1)/12$ and $r/12$ is the fractional part of $(\ell - 1)/12$.

Forecasting Seasonal Models

If our last observation is in December, then the next January value is forecast as Φ times the last observed January value, February is forecast as Φ times the last observed February value, and so on.

Two Januarys ahead is forecast as Φ^2 times the last observed January. Looking just at January values, the forecasts into the future will decay exponentially at a rate determined by the magnitude of Φ.

All of the forecasts for each month will behave similarly but with different initial forecasts depending on the particular month under consideration.

Note the fact that the Ψ-weights are nonzero only for multiple of 12, namely,

$$\Psi_j = \begin{cases} \Phi^{j/12} & \text{for} \quad j = 0, 12, 24, \cdots, \\ 0 & \text{otherwise} \end{cases}$$

we have that the forecast error variance can be written as

$$Var(e_t(\ell)) = \left[\frac{1 - \Phi^{2k+2}}{1 - \Phi^2} \right] \sigma_e^2$$

where as before, k is the integer part of $(\ell - 1)/12$.

Forecasting Seasonal Models

Seasonal MA(1)$_{12}$: For the seasonal MA(1)$_{12}$, we have $Y_t = e_t - \Theta e_{t-12} + \theta_0$. In this case, we see that

$$\begin{aligned}
\hat{Y}_t(1) &= -\Theta e_{t-11} + \theta_0 \\
\hat{Y}_t(2) &= -\Theta e_{t-10} + \theta_0 \\
&\vdots \\
\hat{Y}_t(12) &= -\Theta e_t + \theta_0
\end{aligned}$$

and

$$\hat{Y}_t(\ell) = \theta_0 \text{ for } \ell > 12$$

Here we obtain different forecasts for the months of the first year, but from then on all forecasts are given by the process mean.

For this model, $\Psi_0 = 1, \Psi_{12} = -\Theta$, and $\Psi_j = 0$ otherwise. Thus, we have

$$Var(e_t(\ell)) = \begin{cases} \sigma_e^2 & 1 \leq \ell \leq 12 \\ (1+\Theta^2)\sigma_e^2 & 12 < \ell \end{cases}$$

ARIMA(0,0,0)$\times(0,1,1)_{12}$

The ARIMA(0,0,0)$\times(0,1,1)_{12}$ model is $Y_t - Y_{t-12} = e_t - \Theta e_{t-12}$ or $Y_{t+\ell} = Y_{t+\ell-12} + e_{t+\ell} - \Theta e_{t+\ell-12}$, so that

$$\begin{aligned}\hat{Y}_t(1) &= Y_{t-11} - \Theta e_{t-11}\\ \hat{Y}_t(2) &= Y_{t-10} - \Theta e_{t-10}\\ &\vdots\\ \hat{Y}_t(12) &= Y_t - \Theta e_t\end{aligned}$$

and then

$$\hat{Y}_t(\ell) = \hat{Y}_t(\ell - 12) \text{ for } \ell > 12$$

It follows that all Januarys will forecast identically, all Februarys identically, and so forth.

Forecasting Seasonal Models

If we invert this model, we find that

$$Y_t = (1-\Theta)(Y_{t-12} + \Theta Y_{t-24} + \Theta^2 Y_{t-36} + \cdots) + e_t$$

Consequently, we can write

$$\begin{aligned}
\hat{Y}_t(1) &= (1-\Theta)\sum_{j=0}^{\infty} \Theta^j Y_{t-11-12j} \\
\hat{Y}_t(2) &= (1-\Theta)\sum_{j=0}^{\infty} \Theta^j Y_{t-10-12j} \\
&\vdots \\
\hat{Y}_t(12) &= (1-\Theta)\sum_{j=0}^{\infty} \Theta^j Y_{t-12j}
\end{aligned}$$

Forecasting Seasonal Models

From this representation, we see that the forecast for each January is an EWMA of all observed Januarys, and similarly for each of the other months.

In this case, we have $\Psi_j = 1 - \Theta$ for $j = 12, 24, \cdots$, and zero otherwise. The forecast error variance is then

$$Var(e_t(\ell)) = [1 + k(1 - \Theta)^2]\sigma_e^2$$

where k is the integer part of $(\ell - 1)/12$.

Forecasting Seasonal Models

ARIMA(0,1,1)$\times(0,1,1)_{12}$:

$$Y_t = Y_{t-1} + Y_{t-12} - Y_{t-13} + e_t - \theta e_{t-1} - \Theta e_{t-12} + \theta\Theta e_{t-13}$$

the forecast satisfy

$$\begin{aligned}
\hat{Y}_t(1) &= Y_t + Y_{t-11} - Y_{t-12} - \theta e_t - \Theta e_{t-11} + \theta\Theta e_{t-12} \\
\hat{Y}_t(2) &= \hat{Y}_t(1) + Y_{t-10} - Y_{t-11} - \Theta e_{t-10} + \theta\Theta e_{t-11} \\
&\vdots \\
\hat{Y}_t(12) &= \hat{Y}_t(11) + Y_t - Y_{t-1} - \Theta e_t + \theta\Theta e_{t-1} \\
\hat{Y}_t(13) &= \hat{Y}_t(12) + \hat{Y}_t(1) - Y_t + \theta\Theta e_t
\end{aligned}$$

and

$$\hat{Y}_t(\ell) = \hat{Y}_t(\ell-1) + \hat{Y}_t(\ell-12) - \hat{Y}_t(\ell-13) \text{ for } \ell > 13$$

To understand the general pattern of these forecasts, we can use the representation

$$\hat{Y}_t(\ell) = A_1 + A_2\ell + \sum_{j=0}^{6}\left[B_{1j}\cos\left(\frac{2\pi j\ell}{12}\right) + B_{2j}\sin\left(\frac{2\pi j\ell}{12}\right)\right]$$

where the A's and B's are dependent on $Y_t, Y_{t-1}, \cdots$, or, alternatively, determined from the initial forecasts $\hat{Y}_t(1), \hat{Y}_t(2), \cdots, \hat{Y}_t(13)$. This result follows from the general theory of difference equations and involves the roots of $(1-x)(1-x^{12}) = 0$.

Forecasting Seasonal Models

- Notice that the equation reveals that the forecasts are composed of a linear trend in the lead time plus a sum of periodic components.
- However, the coefficients A_i and B_{ij} are more dependent on recent data than on past data and will adapt to changes in the process as our forecast origin changes and the forecasts are updated.
- This is in stark contrast to forecasting with deterministic time trend plus seasonal components, where the coefficients depend rather equally on both recent and past data and remain the same for all future forecasts.

Prediction Limits

Prediction limits are obtained precisely as in the nonseasonal case. We illustrate this with the carbon dioxide time series.

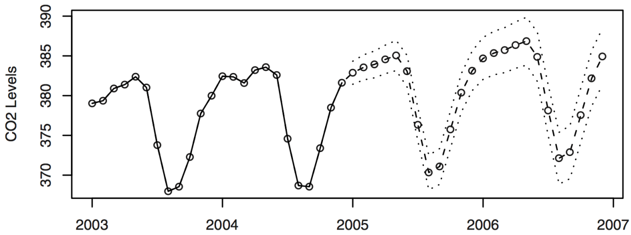

Figure: Forecasts and Forecast Limits for the CO2 Model

This figure shows the forecasts and 95% forecast limits for a lead time of two years for the ARIMA(0,1,1)$\times(0,1,1)_{12}$ model that we fit. The last two years of observed data are also shown. The forecasts mimic the stochastic periodicity in the data quite well, and the forecast limits give a good feeling for the precision of the forecasts.

Prediction Limits

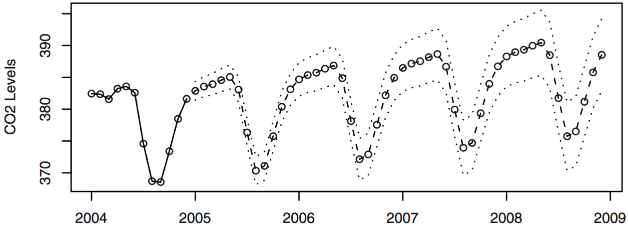

Figure: Long-Term Forecasts for the CO2 Model

This figure displays the last year of observed data and forecasts out four years. At this lead time, it is easy to see that the forecast limits are getting wider, as there is more uncertainty in the forecasts

Summary

- Multiplicative seasonal ARIMA models provide an economical way to model time series whose seasonal tendencies are not as regular as we would have with a deterministic seasonal trend model.
- Fortunately, these models are simply special ARIMA models so that no new theory is needed to investigate their properties. We illustrated the special nature of these models with a thorough modeling of an actual time series.

Questions?