

Modeling and Analysis of Time Series Data

Chapter 4: Linear time series models and the algebra of ARMA models

STATS 531, Winter 2026

Edward Ionides

Definition: A **stationary causal linear process** is a time series model that can be written as

$$[M7] \quad Y_n = \mu + g_0 \epsilon_n + g_1 \epsilon_{n-1} + g_2 \epsilon_{n-2} + g_3 \epsilon_{n-3} + g_4 \epsilon_{n-4} + \dots$$

where $\{\epsilon_n, n = \dots, -2, -1, 0, 1, 2, \dots\}$ is a white noise process, defined for all integer timepoints, with variance $\text{Var}(\epsilon_n) = \sigma^2$.

- We do not need to define any initial values. The doubly infinite noise process $\{\epsilon_n, n = \dots, -2, -1, 0, 1, 2, \dots\}$ is enough to define Y_n for every n as long as the infinite sum in [M7] converges.

Question. When does “stationary” here mean weak stationarity, and when does it mean strict stationarity?

-
- **causal** in [M7] refers to $\{\epsilon_n\}$ being a causal driver of $\{Y_n\}$. The value of Y_n depends only on noise process values already determined by time n .
 - This matches a requirement that causes must precede effects ([Bradford Hill criteria](#)).
 - **linear** refers to linearity of Y_n as a function of $\{\epsilon_n\}$.
-

The autocovariance function for a linear process

$$\gamma_h = \text{Cov}(Y_n, Y_{n+h}) \quad (1)$$

$$= \text{Cov} \left(\sum_{j=0}^{\infty} g_j \epsilon_{n-j}, \sum_{k=0}^{\infty} g_k \epsilon_{n+h-k} \right) \quad (2)$$

$$= \sum_{j=0}^{\infty} \sum_{k=0}^{\infty} g_j g_k \text{Cov}(\epsilon_{n-j}, \epsilon_{n+h-k}) \quad (3)$$

$$= \sum_{j=0}^{\infty} g_j g_{j+h} \sigma^2, \text{ for } h \geq 0. \quad (4)$$

- For the autocovariance function to be finite, we need

$$\sum_{j=0}^{\infty} g_j^2 < \infty. \quad (5)$$

- Note: we assumed we can move $\sum_{j=0}^{\infty} \sum_{k=0}^{\infty}$ through Cov.

- The interchange of expectation and infinite sums cannot be taken for granted. $\text{Cov} \left(\sum_{i=1}^m X_i, \sum_{j=1}^n Y_j \right) = \sum_{i=1}^m \sum_{j=1}^n \text{Cov}(X_i, Y_j)$ is true for finite m and n , but not necessarily for infinite sums.
- In this course, we do not focus on interchange issues, but we try to notice when we make assumptions.
- The interchange of \sum_0^{∞} and Cov can be justified by requiring a stronger condition,

$$\sum_{j=0}^{\infty} |g_j| < \infty. \quad (6)$$

- The MA(q) model that we defined in equation M3 is an example of a stationary, causal linear process.
- The general stationary, causal linear process model, M7, can also be called the MA(∞) model.

Causal and non-causal AR(1) models

The stochastic difference equation defining the AR(1) model,

$$[M8] \quad Y_n = \phi Y_{n-1} + \epsilon_n,$$

has a causal solution,

$$[M8.1] \quad Y_n = \sum_{j=0}^{\infty} \phi^j \epsilon_{n-j}.$$

It also has a non-causal solution,

$$[M8.2] \quad Y_n = - \sum_{j=1}^{\infty} \phi^{-j} \epsilon_{n+j}.$$

Question. Work through the algebra to check that M8.1 and M8.2 both solve equation M8.

Question. For what values of ϕ is the causal solution M8.1 a convergent infinite sum, meaning that it converges to a random variable with finite variance? For what values is the non-causal solution M8.2 a convergent infinite sum?

Using the MA(∞) representation to compute the autocovariance of an ARMA model

Question. The linear process representation can be a convenient way to calculate autocovariance functions. Use the linear process representation in M8.1, together with our expression for the autocovariance of the general linear process M7, to get an expression for the autocovariance function of the AR(1) model.

ARMA models written using the backshift operator

- The **backshift** operator B , also called the **lag** operator, is

$$BY_n = Y_{n-1}. \quad (7)$$

- The **difference** operator $\Delta = 1 - B$ is

$$\Delta Y_n = (1 - B)Y_n = Y_n - Y_{n-1}. \quad (8)$$

- Powers of the backshift operator correspond to different time shifts, e.g.,

$$B^2 Y_n = B(BY_n) = B(Y_{n-1}) = Y_{n-2}. \quad (9)$$

- We can also take a second difference,

$$\begin{aligned} \Delta^2 Y_n &= (1 - B)(1 - B)Y_n \\ &= (1 - 2B + B^2)Y_n = Y_n - 2Y_{n-1} + Y_{n-2}. \end{aligned} \quad (10)$$

- The backshift operator is linear, i.e.,

$$B(\alpha X_n + \beta Y_n) = \alpha BX_n + \beta BY_n = \alpha X_{n-1} + \beta Y_{n-1} \quad (11)$$

- Backshift operators and their powers can be added, multiplied by each other, and multiplied by a scalar.
- Mathematically, backshift operators follow the same rules as the algebra of polynomial functions.
- For example, a distributive rule for $\alpha + \beta B$ is

$$(\alpha + \beta B)Y_n = (\alpha B^0 + \beta B^1)Y_n = \alpha Y_n + \beta BY_n = \alpha Y_n + \beta Y_{n-1}.$$

- Mathematical properties we know about polynomials can be used to work with backshift operators.
 - The AR, MA and linear process model equations can all be written in terms of polynomials in the backshift operator.
-

- Write $\phi(x) = 1 - \phi_1x - \phi_2x^2 - \cdots - \phi_px^p$, an order p polynomial. The equation M1 for the AR(p) model becomes

$$Y_n - \phi_1 Y_{n-1} - \phi_2 Y_{n-2} - \cdots - \phi_p Y_{n-p} = \epsilon_n, \quad (12)$$

which can be written using the backshift operator as

$$[M1'] \quad \phi(B)Y_n = \epsilon_n.$$

- Write $\theta(x)$ for a polynomial of order q ,

$$\theta(x) = 1 + \theta_1x + \theta_2x^2 + \cdots + \theta_qx^q. \quad (13)$$

- The MA(q) equation M3 is equivalent to

$$[M3'] \quad Y_n = \theta(B)\epsilon_n.$$

-
- If $g(x)$ is a function defined by the [Taylor series](#)

$$g(x) = g_0 + g_1x + g_2x^2 + g_3x^3 + g_4x^4 + \dots, \quad (14)$$

the stationary causal linear process equation [M7] is

$$[M7'] \quad Y_n = \mu + g(B)\epsilon_n.$$

- Whatever you know or learn about working with Taylor series expansions helps you understand AR, MA and ARMA models.

The general ARMA model

Putting together M1 and M3 suggests an **autoregressive moving average** ARMA(p,q) model given by

$$[M9] \quad Y_n = \phi_1 Y_{n-1} + \phi_2 Y_{n-2} + \cdots + \phi_p Y_{n-p} + \epsilon_n + \theta_1 \epsilon_{n-1} + \cdots + \theta_q \epsilon_{n-q},$$

where $\{\epsilon_n\}$ is a white noise process. Using the backshift operator, we can write this more succinctly as

$$[M9'] \quad \phi(B)Y_n = \theta(B)\epsilon_n.$$

- Experience with data analysis suggests that models with both AR and MA components often fit data better than a pure AR or MA process.
- The general stationary ARMA(p,q) also has a mean μ ,

$$[M9''] \quad \phi(B)(Y_n - \mu) = \theta(B)\epsilon_n.$$

Obtaining the MA(∞) representation and autocovariance of the ARMA(1,1) model

Consider $Y_n = \phi Y_{n-1} + \epsilon_n + \theta \epsilon_{n-1}$.

Step 1. Put the model in the form $Y_n = g(B)\epsilon_n$.

Formally, we can write

$$(1 - \phi B)Y_n = (1 + \theta B)\epsilon_n, \quad (15)$$

which algebraically is equivalent to

$$Y_n = \left(\frac{1 + \theta B}{1 - \phi B} \right) \epsilon_n. \quad (16)$$

We can write this as

$$Y_n = g(B)\epsilon_n, \quad (17)$$

where

$$g(x) = \left(\frac{1 + \theta x}{1 - \phi x} \right). \quad (18)$$

Step 2. Work out the Taylor series expansion,

$$g(x) = g_0 + g_1 x + g_2 x^2 + g_3 x^3 + \dots \quad (19)$$

You can do this either by hand or using your favorite math software.

Step 3. Obtain the MA(∞) representation, by putting the Taylor series into the form $Y_n = g(B)\epsilon_n$.

Step 4. Obtain the autocovariance function, by using the general formula for an MA(∞) process.

Carrying out this calculation is an exercise.

Causal, invertible ARMA models

- We say that the ARMA model [M9] is **causal** if its MA(∞) representation is a convergent series.
- Recall that **causality** is about writing Y_n in terms of the driving noise process $\{\epsilon_n, \epsilon_{n-1}, \epsilon_{n-2}, \dots\}$.
- **Invertibility** is about writing ϵ_n in terms of Y_n, Y_{n-1}, \dots .
- To assess causality, we consider the convergence of the Taylor series expansion of $\theta(x)/\phi(x)$ in the ARMA representation

$$Y_n = \frac{\theta(B)}{\phi(B)} \epsilon_n.$$

- To assess invertibility, we consider the convergence of the Taylor series expansion of $\phi(x)/\theta(x)$ in the inversion of the ARMA model given by

$$\epsilon_n = \frac{\phi(B)}{\theta(B)} Y_n.$$

-
- Fortunately, there is a simple way to check causality and invertibility without calculating the Taylor series.
 - The ARMA model is causal if the AR polynomial,

$$\phi(x) = 1 - \phi_1 x - \phi_2 x^2 - \dots - \phi_p x^p$$

has all its roots (i.e., solutions to $\phi(x) = 0$) outside the unit circle in the complex plane.

- The ARMA model is invertible if the MA polynomial,

$$\theta(x) = 1 + \theta_1 x + \theta_2 x^2 + \dots + \theta_q x^q$$

has all its roots outside the unit circle.

Question. It is undesirable to use a non-invertible model for data analysis. Why?
Hint: One answer to this question involves diagnosing model misspecification.

We can check the roots using NumPy's `roots` function. For example, consider the MA(2) model, $Y_n = \epsilon_n + 2\epsilon_{n-1} + 2\epsilon_{n-2}$. The roots to $\theta(x) = 1 + 2x + 2x^2$ are

```
roots = np.roots([2, 2, 1])
print(roots)
```

`[-0.5+0.5j -0.5-0.5j]`

- Finding the absolute value shows that we have two roots inside the unit circle, so this MA(2) model is not invertible.

```
print(np.abs(roots))
```

`[0.70710678 0.70710678]`

- In this case, you should be able to find the roots algebraically. In general, numerical evaluation of roots is useful.

Reducible and irreducible ARMA models

- Write the ARMA model written as a ratio of two polynomials,

$$Y_n = \frac{\theta(B)}{\phi(B)} \epsilon_n. \quad (20)$$

If the two polynomials $\phi(x)$ and $\theta(x)$ share a common factor, it can be canceled out without changing the model.

- The **fundamental theorem of algebra** says that every polynomial $\phi(x) = 1 - \phi_1 x - \dots - \phi_p x^p$ of degree p can be written in the form

$$(1 - x/\lambda_1) \times (1 - x/\lambda_2) \times \dots \times (1 - x/\lambda_p), \quad (21)$$

where $\lambda_{1:p}$ are the p roots of the polynomial, which may be real or complex valued.

- The Taylor series expansion of $\phi(B)^{-1}$ is convergent if and only if $(1 - B/\lambda_i)^{-1}$ has a convergent expansion for each $i \in 1 : p$. This happens if $|\lambda_i| > 1$ for each i .

-
- The polynomials $\phi(x)$ and $\theta(x)$ share a common factor if, and only if, they share a common root.

- It is not clear, just from looking at the model equations, that

$$Y_n = \frac{5}{6}Y_{n-1} - \frac{1}{6}Y_{n-2} + \epsilon_n - \epsilon_{n-1} + \frac{1}{4}\epsilon_{n-2} \quad (22)$$

is **exactly the same model** as

$$Y_n = \frac{1}{3}Y_{n-1} + \epsilon_n - \frac{1}{2}\epsilon_{n-1}. \quad (23)$$

- To see this, you have to do the math! We see that the second of these equations is derived from the first by canceling out the common factor $(1 - 0.5B)$ in the ARMA model specification.

```
ar_roots = np.roots([1/6, -5/6, 1])
ma_roots = np.roots([1/4, -1, 1])
print(f"AR roots: {ar_roots}, MA roots: {ma_roots}")
```

AR roots: [3. 2.], MA roots: [2.0000001 1.9999999]

AR(2) models and oscillating behavior

- Non-random physical processes evolving through time have been modeled using differential equations ever since the ground-breaking work by Newton [1].
- The data and systems we want to study have considerable randomness (unpredictability).
- However, it is helpful to study related deterministic systems.
- The **deterministic skeleton** of a time series model is the non-random process obtained by removing randomness from a stochastic model.
- For a discrete-time model, we can define a continuous-time deterministic skeleton by replacing the discrete-time difference equation with a differential equation.
- Rather than deriving a deterministic skeleton from a stochastic time series model, we can instead add stochasticity to a deterministic model.

Oscillatory behavior modeled using an AR(2) process

- In physics, a basic model for processes that oscillate (springs, pendulums, vibrating machine parts, etc) is simple harmonic motion.
- The differential equation for a simple harmonic motion process $x(t)$ is

$$[M10] \quad \frac{d^2}{dt^2}x(t) = -\omega^2x(t).$$

- This is a second order linear differential equation with constant coefficients. Such equations have a closed form solution. You may already know that the solution to M10 is **sinusoidal**.
- Finding the solution to a linear differential equation is very similar to the task of solving difference equations which is useful elsewhere in time series analysis. It also gives a chance to review complex numbers. Let's see how it is done.

-
1. Look for solutions of the form $x(t) = e^{\lambda t}$. Substituting this into the differential equation [M10] we get

$$\lambda^2 e^{\lambda t} = -\omega^2 e^{\lambda t}. \quad (24)$$

2. Canceling the term $e^{\lambda t}$, we see that this has two solutions, with

$$\lambda = \pm \omega i, \quad \text{where } i = \sqrt{-1}. \quad (25)$$

3. The linearity of the differential equation means that if $y_1(t)$ and $y_2(t)$ are two solutions, then $ay_1(t) + by_2(t)$ is also a solution for any a and b . So, the **general solution** to M10 is

$$x(t) = ae^{i\omega t} + be^{-i\omega t}. \quad (26)$$

Here, a and b could be complex numbers.

-
4. We may suspect that $x(t) = ae^{i\omega t} + be^{-i\omega t}$ is sinusoidal by recalling the identities

$$\sin(\omega t) = \frac{1}{2i}(e^{i\omega t} - e^{-i\omega t}), \quad \cos(\omega t) = \frac{1}{2}(e^{i\omega t} + e^{-i\omega t}). \quad (27)$$

5. For physical systems, $x(t)$ is real so we know that the complex part must be zero. Thus, the two terms on the right are complex conjugates. Writing $a = (A/2)e^{i\phi}$, this implies $b = (A/2)e^{-i\phi}$ for real A and ϕ . The factor of $1/2$ is arbitrary. This gives

$$x(t) = \frac{A}{2} (e^{i(\omega t+\phi)} + e^{-i(\omega t+\phi)}). \quad (28)$$

6. Putting together these identities we get

$$x(t) = A \cos(\omega t + \phi), \quad (29)$$

which explains why the factor of $1/2$ is convenient.

Frequency, amplitude and phase for $x(t) = A \cos(\omega t + \phi)$

- ω is called the **frequency**, and ϕ is called the **phase**.
- Angle is usually measured in **radians**, so the units of ω are radians per unit time, and units of ϕ are radians.
- The **period** is $2\pi/\omega$, the time for one cycle.
- A is called the **amplitude**.
- The frequency of the oscillation is determined by ω in M10, but the amplitude and phase are unspecified constants which may be determined by initial conditions.
- It may be convenient to rescale to **cycles per unit time**,

$$x(t) = A \cos(2\pi(\omega't + \phi')) \quad (30)$$

where $\omega' = \omega/2\pi$, $\phi' = \phi/2\pi$.

- A discrete time version of M10 is a deterministic linear difference equation, replacing $\frac{d^2}{dt^2}$ by the second difference operator, $\Delta^2 = (1 - B)^2$. This corresponds to a deterministic model equation,

$$\Delta^2 y_n = -\omega^2 y_n.$$

- Adding white noise, and expanding out $\Delta^2 = (1 - B)^2$, we get a stochastic model,

$$[M11] \quad Y_n = \frac{2Y_{n-1}}{1 + \omega^2} - \frac{Y_{n-2}}{1 + \omega^2} + \epsilon_n.$$

- Model M11 may be appropriate to describe systems that have semi-regular but somewhat erratic fluctuations, called **quasi-periodic** behavior. Such behavior is evident in business cycles or wild animal populations.

We look at a simulation from M11 with $\omega = 0.1$ and $\epsilon_n \sim \text{iid } N[0, 1]$. From our exact solution to the deterministic skeleton, we expect that a typical period of the oscillations should be $2\pi/\omega \approx 60$.

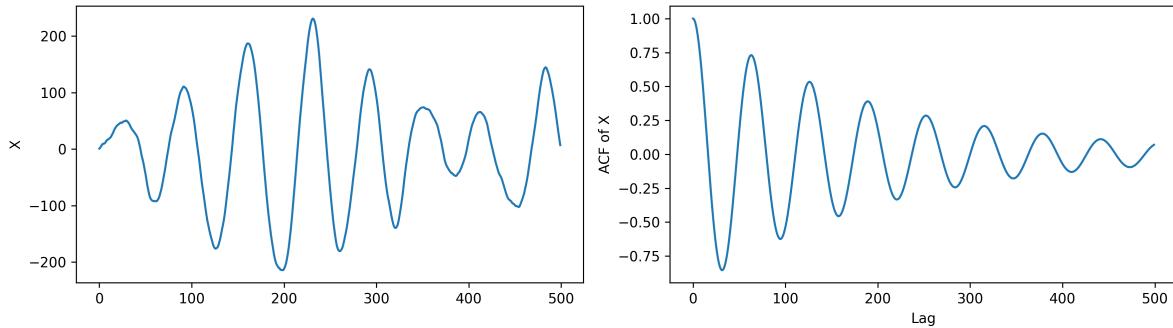
```

omega = 0.1
ar_coefs = [2/(1+omega**2), -1/(1+omega**2)]
# Using ArmaProcess for simulation
ar_poly = np.r_[1, -np.array(ar_coefs)]
ma_poly = np.array([1])
ar_process = ArmaProcess(ar_poly, ma_poly)
X = ar_process.generate_sample(nsample=500,scale=1.0)

fig,(ax1,ax2) = plt.subplots(1,2,figsize=(12,3.5))
ax1.plot(X); ax1.set_ylabel('X')
# Compute theoretical ACF
acf_vals = ar_process.acf(lags=500)
ax2.plot(acf_vals); ax2.set_ylabel('ACF of X')
ax2.set_xlabel('Lag')
plt.tight_layout(); plt.show()

```

Quasi-periodic AR(2) simulation



-
- Quasi-periodic fluctuations are **phase locked** when the random perturbations are not able to knock the oscillations away from being close to their initial phase.
 - Eventually, the randomness should mean that the process is equally likely to have any phase, regardless of the initial phase.

Question. What is the timescale on which the simulated model shows phase locked behavior? Equivalently, on what timescale does the phase of the fluctuations lose memory of its initial phase?

Further reading

- Section 2.2 of Shumway and Stoffer [2] introduces the backshift operator.
- Section 3.1 develops the theory of ARMA models in a similar way to this chapter.
- Section 3.2 gives a difference equation approach to calculating ARMA autocovariance functions which gives an opportunity to practice algebra similar to our study of the AR(2) model.

Acknowledgments

- Compiled on January 18, 2026 using Python.
- Licensed under the [Creative Commons Attribution-NonCommercial license](#). Please share and remix non-commercially, mentioning its origin.
- We acknowledge [previous versions of this course](#).

References

- [1] Isaac Newton. *Philosophiae Naturalis Principia Mathematica*. 1687.
- [2] Robert H Shumway and David S Stoffer. *Time Series Analysis and its Applications: With R Examples*. 4th. Springer, 2017.