

Modeling Seasonal Influenza-Like Illness with SARMA under a COVID-19 Structural Break

1 Abstract

This report analyzes weekly Influenza-Like Illness (ILI) data from the CDC ILINet system. The series shows strong annual seasonality ($s = 52$), a long-term trend, and a clear structural disruption associated with the COVID-19 pandemic. A log transformation is applied to stabilize variance, and a piecewise quadratic trend with a pandemic indicator is used to capture changes in level and slope. The remaining dependence structure is modeled using a Seasonal Autoregressive Moving-Average (SARMA) model. Based on the Akaike Information Criterion (AIC) comparison, a SARMA(2, 1) \times (0, 1)₅₂ specification is selected. Residual diagnostics suggest the fitted model adequately captures the temporal dependence, with residuals behaving close to white noise, aside from one notable outlier during the COVID-19 period.

2 Introduction

Seasonal influenza is a persistent public health concern in the United States, with annual epidemics leading to substantial morbidity. Weekly Influenza-Like Illness (ILI) activity is tracked through the CDC's ILINet surveillance system, which provides a long historical record of respiratory disease patterns. These patterns were sharply disrupted during the COVID-19 pandemic: beginning in early 2020, masking, social distancing, and changes in healthcare-seeking behavior led to a marked drop in reported ILI cases, making standard time series modeling more challenging. Using ILINet data from 2010 to 2026, this study examines the trend and seasonal structure of national ILI activity, explicitly accounts for the COVID-19-related structural break, and fits a SARMA model to the detrended series. It is worth noting that parallel trends observed in other respiratory disease studies—such as the time series analysis of COVID-19 in Ireland (Anonymous 2025)—showed a primary peak of confirmed cases between 2020 and 2022. This external context aligns with the substantial fluctuations observed in the ILITOTAL data during the same period, suggesting a broader pandemic-driven impact on public health metrics.

3 Methods

3.1 Data Preprocessing

The ILINet dataset provides weekly national ILI counts from the CDC surveillance system. Week 53 observations, which arise as an artifact of the ISO week calendar rather than reflecting a complete influenza season, are removed to maintain a consistent 52-week annual cycle. The response variable, ILITOTAL, shows increasing variability as the level of the series rises. A log transformation is therefore applied to stabilize the variance and simplify the seasonal structure.

3.2 Trend Removal

The log-transformed series exhibits a nonlinear long-term trend and a clear structural change beginning in early 2020, coinciding with the onset of the COVID-19 pandemic. To account for this behavior, we fit a piecewise quadratic regression of the form:

$$\log(\text{ILITOTAL}_t) = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_3 C_t + \beta_4 (t \cdot C_t) + \beta_5 (t^2 \cdot C_t) + \varepsilon_t,$$

where t denotes elapsed time from the start of the series and $C_t = \mathbf{1}(t \geq 2020.4)$ is an indicator equal to 1 for observations around May 2020 and 0 otherwise. The interaction terms allow both the slope and curvature of the trend to shift after the pandemic onset. The fitted values from this regression are treated as the estimated trend, and the residuals

$$\hat{\varepsilon}_t = \log(\text{ILITOTAL}_t) - \log(\widehat{\text{ILITOTAL}}_t)$$

are used as the detrended series in subsequent analysis. (Ionides 2026)

3.3 SARMA Modeling

The detrended residuals retain strong annual seasonality with period $s = 52$ weeks, as evidenced by a sinusoidal ACF pattern with peaks at multiples of lag 52. To achieve stationarity, a lag-52 seasonal difference is applied to the detrended series prior to model fitting. The seasonally differenced series is then modeled as SARMA(p, q) \times (P, Q)₅₂, defined by

$$\phi(B)\Phi(B^{52})\tilde{\varepsilon}_t = \psi(B)\Psi(B^{52})\epsilon_t,$$

where B is the backshift operator, ϕ and ψ are the non-seasonal AR and MA polynomials of orders p and q , and Φ and Ψ are the seasonal AR and MA polynomials of orders P and Q . Model orders are selected by minimizing the Akaike Information Criterion (Shumway and Stoffer 2017),

$$\text{AIC} = -2\ell(\hat{\theta}) + 2D,$$

where $\ell(\hat{\theta})$ is the maximized log-likelihood and D is the number of estimated parameters. The grid search covers $p, q \in \{0, 1, 2\}$ and $P, Q \in \{0, 1, 2\}$.

3.4 Model Diagnostics

Model adequacy is assessed through residual diagnostics: the residual ACF for remaining autocorrelation, a normal Q-Q plot to assess approximate normality of innovations, and the Ljung-Box test as a formal test of the white noise hypothesis. The autocorrelation function is defined as

$$\rho(h) = \frac{\gamma(h)}{\gamma(0)},$$

where $\gamma(h) = \text{Cov}(Y_t, Y_{t+h})$ is the autocovariance at lag h .

4 Results

4.1 Data Exploration

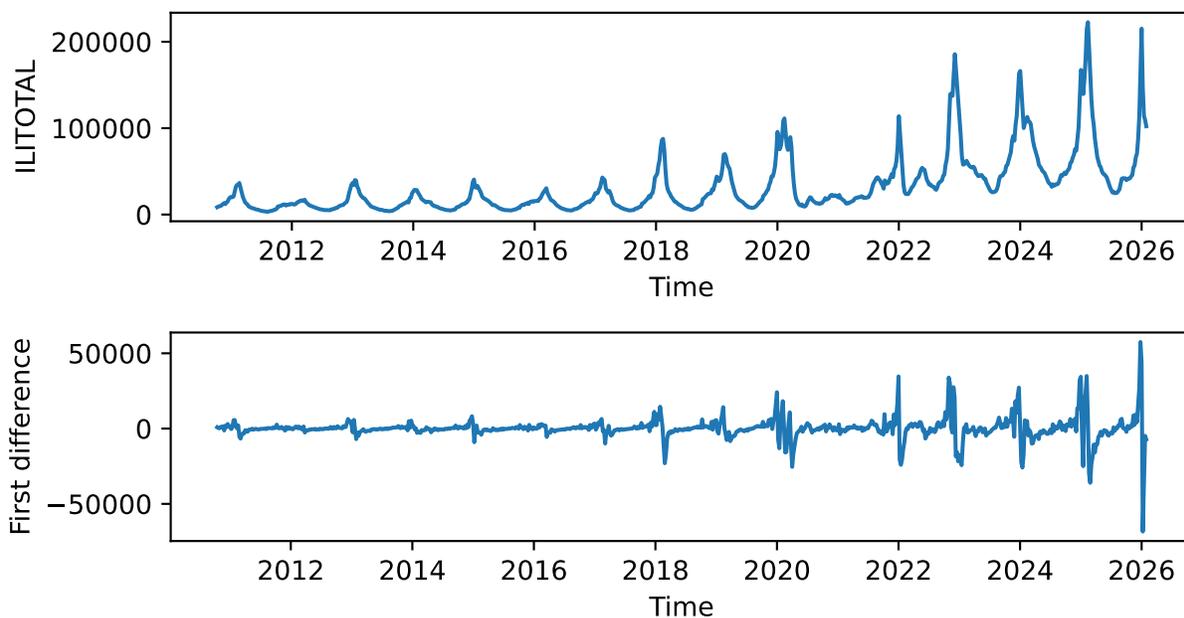


Figure 1: Raw ILITOTAL counts (top) and first differences (bottom) over time.

The raw time series exhibits strong annual seasonality and growing variability over time, particularly after 2020 (Figure 1). The first-difference plot confirms increasing variance, indicating the series is non-stationary and that a variance-stabilizing transformation is warranted. Since all values are positive, a log transformation is applied.

After log transformation, the variance is substantially stabilized (Figure 2). However, a nonlinear upward trend remains, and a prominent dip is visible from 2020 to 2022, consistent with COVID-19 suppression of ILI reporting.

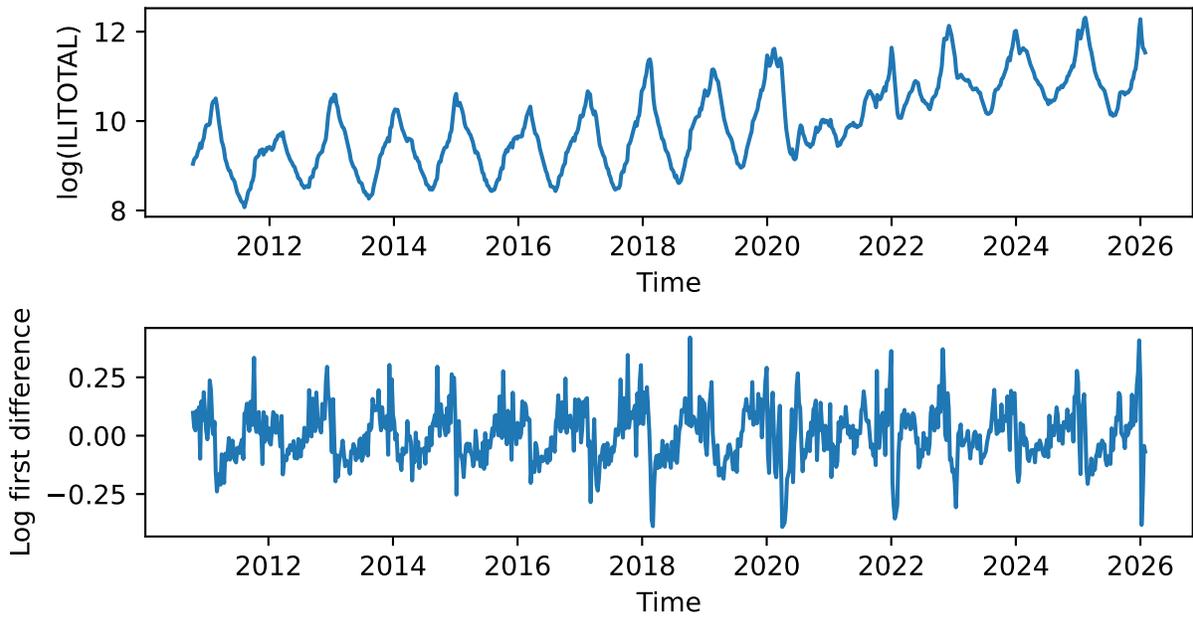


Figure 2: Log-transformed ILITOTAL (top) and log first differences (bottom).

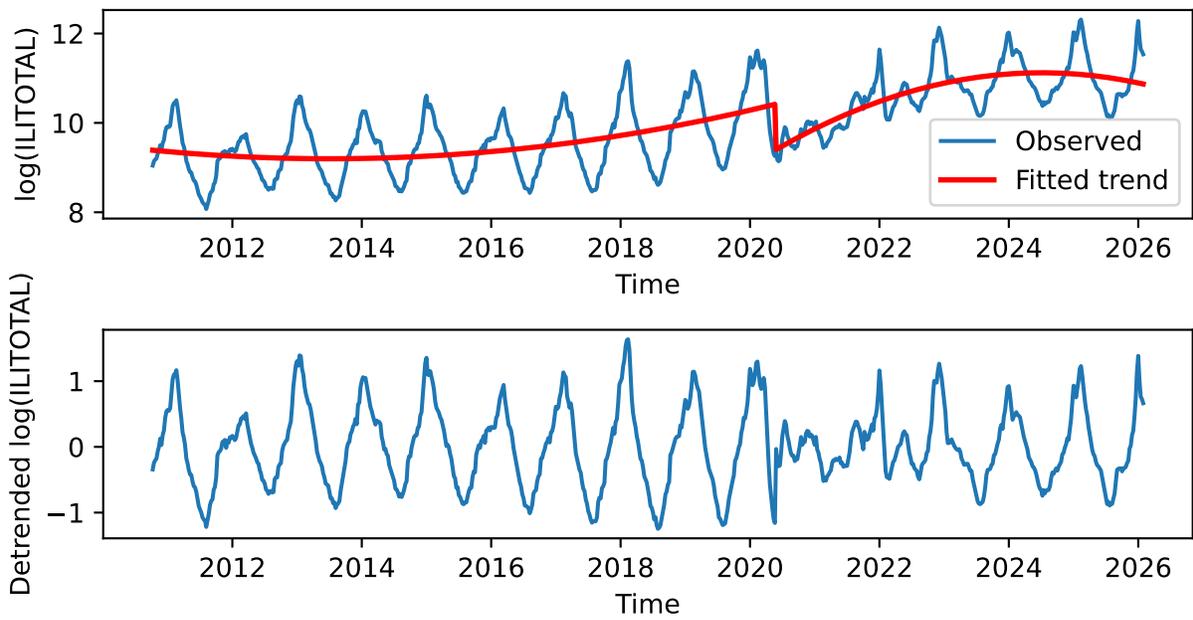


Figure 3: Log ILITOTAL with fitted piecewise quadratic trend (top) and detrended residuals (bottom).

4.2 Trend Estimation and Detrending

The piecewise quadratic regression provides a good fit to the long-term trajectory of the series (Figure 3). All six predictors are statistically significant ($p < 0.001$), confirming that the pre- and post-COVID trend structures differ meaningfully. The detrended residuals are centered around zero, but the ACF still exhibits a sinusoidal pattern with peaks at multiples of lag 52, indicating persistent seasonal non-stationarity.

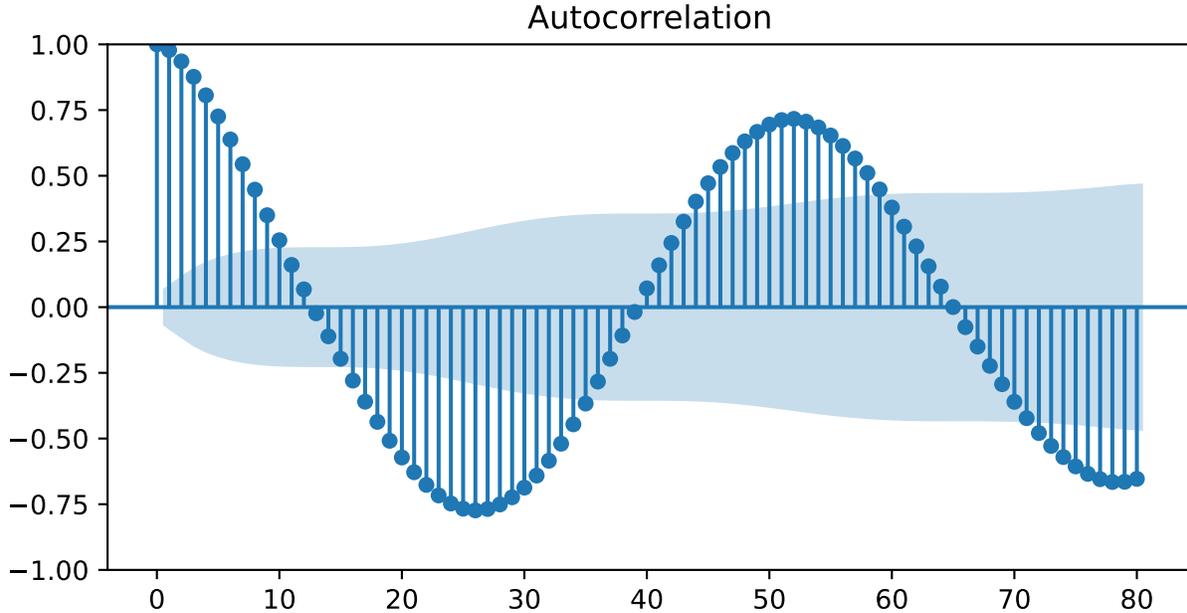


Figure 4: ACF of detrended $\log(\text{ILITOTAL})$, showing persistent seasonal autocorrelation at multiples of lag 52.

4.3 Stationarity via Seasonal Differencing

After applying a lag-52 seasonal difference, both the time plot and ACF indicate a weakly stationary series (Figure 5). The sinusoidal pattern is eliminated, and autocorrelations decay toward zero at non-seasonal lags, justifying the application of a SARMA model.

4.4 Model Selection via AIC

The AIC grid search identifies $\text{SARMA}(2, 1) \times (0, 1)_{52}$ as the model with the lowest AIC. The key AIC table for seasonal $(P, Q) = (0, 1)$ is shown above; full results across all seasonal orders are provided in the Supplementary Material.

4.5 Model Fitting and Diagnostics

	lb_stat	lb_pvalue
1	0.0754	0.7837
52	68.3559	0.0636

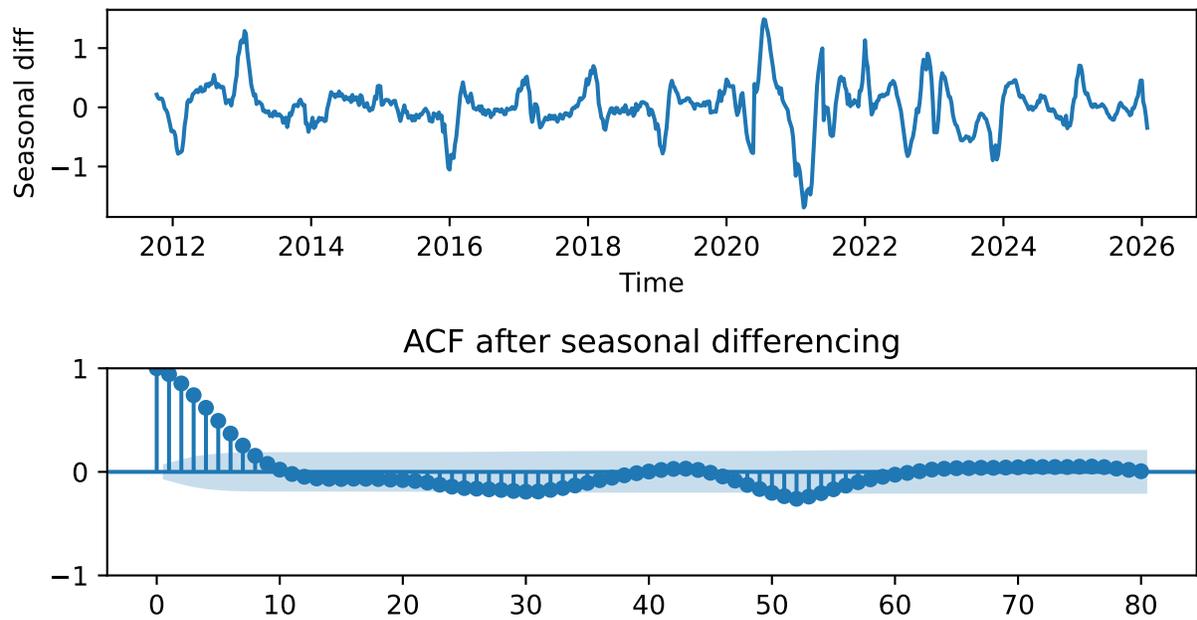


Figure 5: Seasonally differenced detrended series (top) and its ACF (bottom), indicating approximate stationarity.

Table 1: AIC values for SARMA(p,q)x(P,Q)_52 with seasonal (P,Q)=(0,1). The minimum AIC corresponds to SARMA(2,1)x(0,1)_52.

Seasonal (P,Q)=(0,1)			
	MA0	MA1	MA2
AR0	635.41	-76.22	-547.51
AR1	-1057.15	-1126.63	-1173.38
AR2	-1171.08	-1203.68	-1199.09

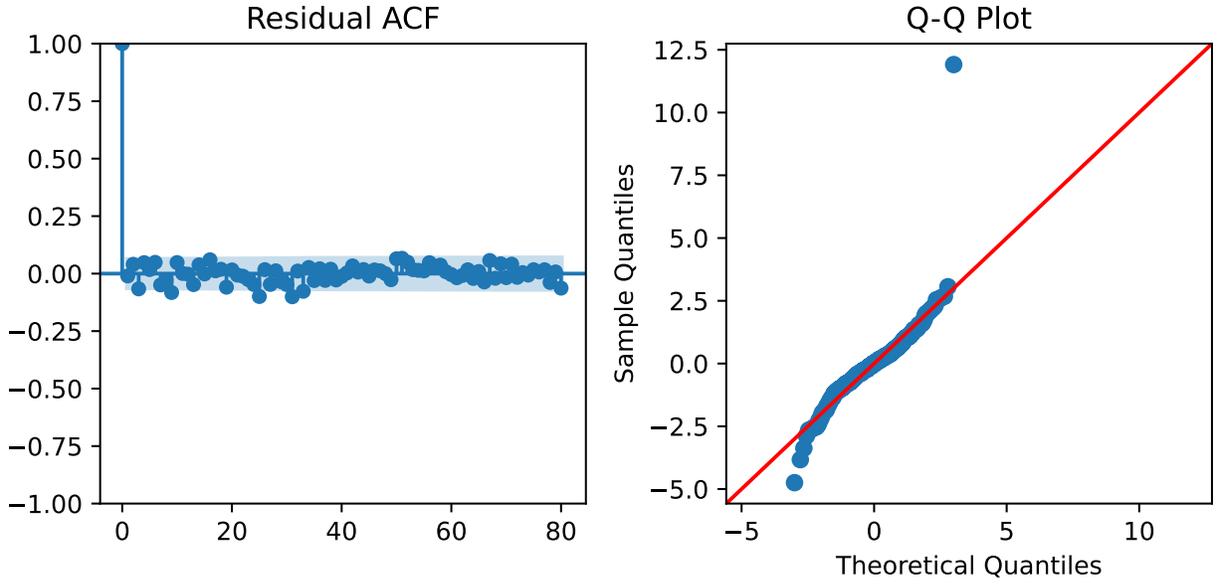


Figure 6: Residual diagnostics for the fitted SARMA(2,1)x(0,1)₅₂ model: ACF (left) and Q-Q plot (right).

The residual ACF shows all autocorrelations after lag 0 within or very close to the 95% confidence bands, indicating no significant remaining serial dependence (Figure 6). The Ljung-Box test fails to reject the null hypothesis of white noise at both lag 1 and lag 52 ($p > 0.05$), supporting model adequacy. The Q-Q plot shows residuals are approximately Gaussian, with the exception of one prominent outlier at the upper tail, likely corresponding to an anomalous week during the early COVID-19 period not fully captured by the piecewise trend.

5 Conclusion

This study analyzed weekly CDC ILINet data from 2010 to 2026 using a two-stage approach: a piecewise quadratic regression to remove the long-term trend and COVID-19-induced structural break, followed by SARMA modeling of the seasonally differenced detrended residuals. The selected SARMA(2, 1) × (0, 1)₅₂ model provides an adequate fit, with residuals consistent with white noise as confirmed by the residual ACF and Ljung-Box test. The residuals are approximately Gaussian, with one notable outlier attributable to extreme disruption in ILI reporting during the early COVID-19 period. This suggests that while the piecewise trend effectively captures the broad structural break, a single anomalous observation persists that may warrant further investigation, such as an additional indicator variable for the most acute disruption weeks.

6 Acknowledgments

Data were sourced from the CDC FluView ILINet surveillance system (Centers for Disease Control and Prevention 2026). Statistical method and the some of the code pattern follow the STATS 531

course notes (Ionides 2026). The SARMA modeling framework references Shumway and Stoffer (Shumway and Stoffer 2017). AI tools (Gemini, ChatGPT) were used to assist with code debugging and discussing alternative model specifications. Final model selection, analysis, and interpretation were conducted by the authors.

7 References

- Anonymous. 2025. “Ireland Covid Cases.” https://ionides.github.io/531w25/midterm_project/project04/blinded.html.
- Centers for Disease Control and Prevention. 2026. “FluView: ILINet Outpatient Influenza-Like Illness Surveillance Network.” <https://www.cdc.gov/fluview/>.
- Ionides, Edward. 2026. “Modeling and Analysis of Time Series Data (STATS 531) Course Notes.” <https://ionides.github.io/531w26/>.
- Shumway, Robert H., and David S. Stoffer. 2017. *Time Series Analysis and Its Applications: With R Examples*. 4th ed. Springer.

8 Supplementary

The study is also constrained by certain limitations. The existence of week 53 indicates repetitive week counts, which could lead to range misarrangement and mismatched time series coefficients at yearly connections. Furthermore, without causal inference and introducing more confounding factors, the finite explanatory power of treating the COVID pandemic as a dummy variable—even if it matches the observed trend—also restrains the power of the time series model at its assumption of stationarity. In the future, we can look into additional factors that could explain the major dip better and explore a more efficient way of time arrangement than simply removing week 53.

Full AIC tables across all seasonal (P, Q) combinations are provided below for completeness.

```
=== Seasonal (P,Q)=(0,0) ===
      MA0      MA1      MA2
AR0  748.67    15.77  -435.18
AR1 -897.58  -951.68  -986.33
AR2 -980.89 -1016.68 -1013.09
```

```
=== Seasonal (P,Q)=(0,1) ===
      MA0      MA1      MA2
AR0  635.41   -76.22  -547.51
AR1 -1057.15 -1126.63 -1173.38
AR2 -1171.08 -1203.68 -1199.09
```

```
=== Seasonal (P,Q)=(0,2) ===
      MA0      MA1      MA2
AR0  514.57  -122.04  -525.82
AR1 -954.32 -1015.73 -1054.69
```

AR2 -1054.22 -1080.62 -1075.98

=== Seasonal (P,Q)=(1,0) ===

	MA0	MA1	MA2
AR0	652.53	-35.29	-478.88
AR1	-940.96	-998.86	-1049.22
AR2	-1032.43	-1075.30	-1075.47

=== Seasonal (P,Q)=(1,1) ===

	MA0	MA1	MA2
AR0	645.03	-119.00	-570.88
AR1	-1059.65	-1127.16	-1172.29
AR2	-1167.59	-1202.14	-1197.52

=== Seasonal (P,Q)=(1,2) ===

	MA0	MA1	MA2
AR0	512.78	-125.48	-534.89
AR1	-970.35	-1032.85	-1075.65
AR2	-1073.16	-1100.06	-1095.46

=== Seasonal (P,Q)=(2,0) ===

	MA0	MA1	MA2
AR0	591.76	-40.28	-450.21
AR1	-881.47	-937.58	-986.16
AR2	-970.54	-1010.88	-1010.42

=== Seasonal (P,Q)=(2,1) ===

	MA0	MA1	MA2
AR0	554.19	-90.64	-503.53
AR1	-942.65	-1010.40	-1053.34
AR2	-1046.09	-1076.22	-1074.23

=== Seasonal (P,Q)=(2,2) ===

	MA0	MA1	MA2
AR0	498.98	-134.54	-542.87
AR1	-970.18	-1031.62	-1074.34
AR2	-1068.75	-1098.40	-1093.76