



JMP045: Modeling Gold Prices

ARIMA/ARMA Models, Model Comparison

Produced by

M Ajoy Kumar, Associate Professor
Siddaganga Institute of Technology
ajoy@sit.ac.in

Muralidhara A, JMP Global Academic Team
muralidhara.a@jmp.com



Modeling Gold Prices

ARIMA/ARMA Models, Model Comparison

Key ideas:

The case study deals with univariate time series modeling, in which a time series is modeled using its own past values. Hence, these models use only one data series. Univariate models are specialized models, where past values (or lags) of a series are considered as independent variable. ARIMA/ARMA is a popular univariate model, used extensively for analyzing the characteristics and forecasting time series data. This study analyzes time series data with JMP.

Background

Hari, a research assistant at a leading university, has been asked by his professor to prepare a report on gold prices in the United States. The professor wants Hari to look at the price of gold over a five-year period, analyze the characteristics of gold prices and suggest a suitable univariate model that fits the data.

The Task

Hari is entrusted with the following tasks:

- Collect daily gold prices for a five-year period.
- Study the unit root property of the data.
- Identify a suitable univariate model.
- Estimate the parameters of the best fit model.
- Perform diagnostic checking of the model.

The Data [GP.jmp](#)

Hari decided to use daily gold prices for a period of five years from January 2016 to December 2020. The data was collected from Yahoo Finance. It contained 1,245 observations of the price of gold in USD per troy ounce. The data set has two series: date and daily gold prices.

Date	Day on which the gold price is considered
GP	Gold price on a specific date/day

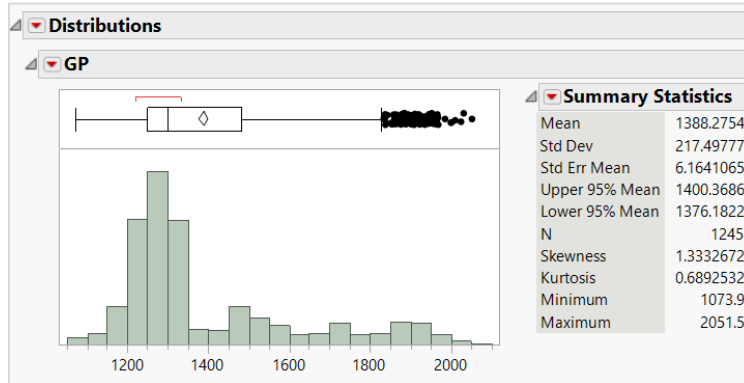
Gold price is a continuous time series variable, whereas the date is time variable.

Analysis

Descriptive statistics

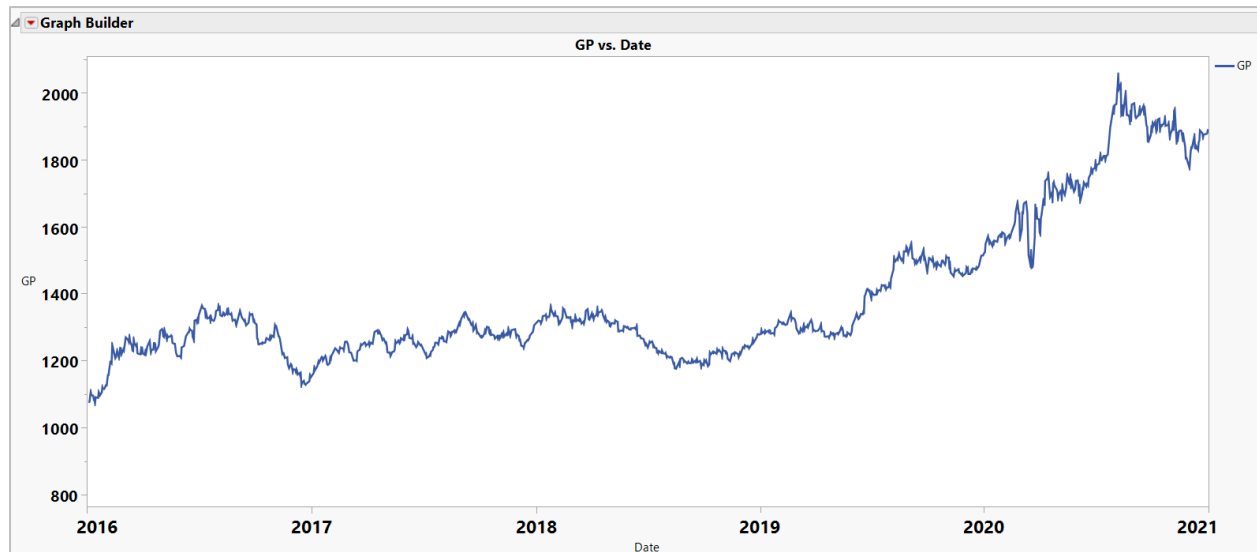
Let's explore the data using distribution and Graph Builder in JMP.

Exhibit 1 Summary Statistics of Gold Price



To create, *Analyze>Distribution>Y = GP>OK*. Under the red triangle next to *Distributions*, select *Stack* to align the output horizontally. Under the red triangle next to *Summary Statistics*, select *Customize Summary Statistics* and then *N*, *Skewness*, *Kurtosis*, *Minimum* and *Maximum*. Click *OK*.

Exhibit 2 Movement of Daily Gold Price (2016-2020)



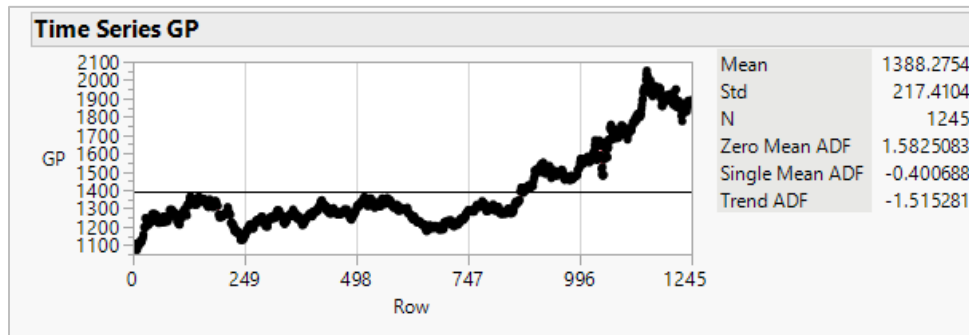
To create, *Graph>Graph Builder>Y = GP, X = Date*. Select *line graph* from the chart options. Click *Done*.

The basic descriptive characteristics of the data are presented in Exhibit 1 and the graph showing movement of gold prices during the five years is given in Exhibit 2. The gold prices fluctuated between \$1,073 and \$2,051 during the five years, with a mean of \$1,388 and standard deviation of \$217. Exhibit 2 shows an upward trend in prices during this period, especially after 2019.

Stationarity of the data

Stationarity of the data series is a prerequisite for most of the econometric models. From Exhibit 2, it is evident that the gold price is not a stationary series. In order to confirm the same, the augmented Dickey-Fuller (ADF) test is used.

Exhibit 3 ADF Test for Gold Price



To create, Analyze>Specialized Modeling>Time Series>Y = GP, X, Time ID = Date. Click OK.

The null hypothesis for the ADF test is that the series has unit root, or the series is non-stationary. The ADF test returns three test statistics:

- Zero Mean ADF: A test against a random walk with zero mean.
- Single Mean ADF: A test against a random walk with a non-zero mean.
- Trend ADF: A test against a random walk with a non-zero mean and a linear trend.

The test statistic is expected to be negative; therefore, it must be more negative (less) than the critical value for the hypothesis to be rejected. The values shown for the Zero Mean, Single Mean and Trend ADF in JMP are the Tau statistics associated with the Dickey-Fuller test. Because Dickey and Fuller produced tables for the critical values associated with the distribution of the Tau statistic, and because the associated p-values would only be approximations, the JMP developer decided not to display approximate p-values for these statistics. The critical values for the ADF test at a 5% level are -2.86 without trend and -3.41 with trend for large samples.

The result of the ADF test for gold prices is given in Exhibit 3. The test statistic of all three ADF tests are above the critical value of 5%. Hence, the null hypothesis is accepted, and it is concluded that the gold price series is non-stationary.


Differencing the data

Differencing is a technique used for transforming a non-stationary data series to stationary. Most of the financial market data becomes stationary on first differencing. We shall create a new variable called FDGP, which is the first difference of the gold prices; it can be determined through multiple ways in JMP.

The easiest way is by selecting the column GP, right-click and then choose New Formula Column>Row>Difference. This will create a new column, Difference [GP], which will have the first difference values of GP. You can rename the column as FDGP by double-clicking the column header. The differenced series is saved as FDGP.

Now let us explore the FDGP using Graph Builder.

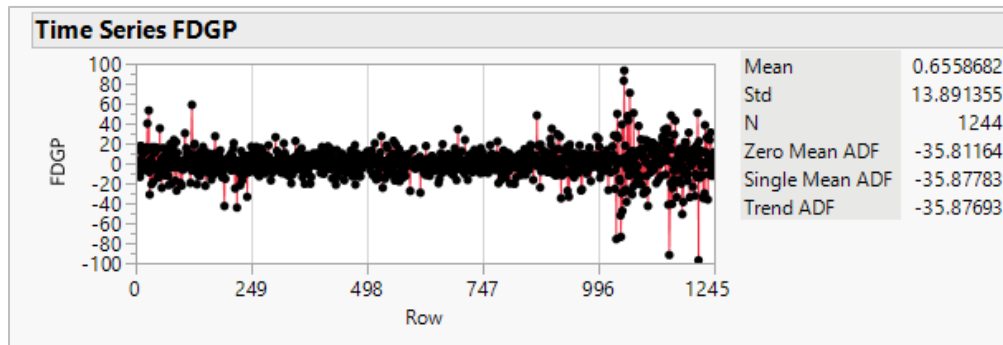
Exhibit 4 First Difference of Daily Gold Prices (2016-2020)



To create, Graph>Graph Builder>Y = FDGP, X = Date. Select line graph from the chart options and click Done.

From Exhibit 4, it can be observed that the series is stationary. However, to confirm the same, an ADF test is performed.

Exhibit 5 ADF Test for First Difference of Gold Prices



To create, Analyze>Specialized Modeling>Time Series>Y = FDGP, X, Time ID = Date. Click OK.

Exhibit 5 gives the results of the ADF test of FDGP. Since the test statistics are all less than the critical value, the null hypothesis is rejected. So, we can conclude that the first difference of gold prices is a stationary series.

ARIMA & ARMA models

An autoregressive (AR) process is one where the current value of a variable depends on its past values. The number of past values (lags) that determine the current value is known as the order of the AR model. Thus, an AR (3) model would use three past values of the data for modeling the current value. In more general terms an AR(p) model is specified as:

$$y_t = \alpha + \sum_{i=1}^p \beta_i y_{t-i} + u_t$$

A moving average (MA) process is one where the current value of a variable depends on the past and current values of the white noise disturbance terms (error terms). The number of past white noise terms included in the model is known as the order of the MA model. An MA (q) model with (q) lags is specified as:

$$y_t = \alpha + \sum_{i=1}^q \phi_i u_{t-i} + u_t$$

The autoregressive moving average (ARMA) process is a combination of the AR and MA processes. In the ARMA model, the current value of a variable depends on the past values of the variable itself and the past and current white noise disturbance terms. ARMA (p,q) model represents an ARMA process with (p) lags of AR terms and (q) lags of MA terms. The model is specified as:

$$y_t = \alpha + \sum_{i=1}^p \beta_i y_{t-i} + \sum_{i=1}^q \phi_i u_{t-i} + u_t$$

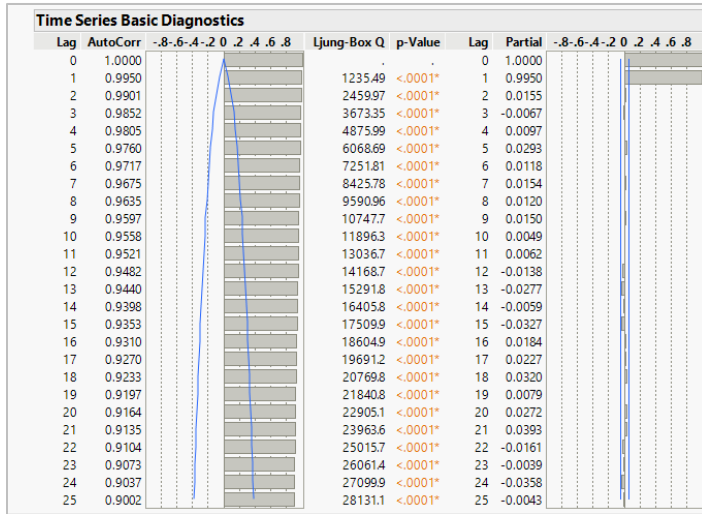
Building ARMA models involves three steps: identification, estimation and diagnostic checking. Identification deals with choosing the right order of the model that captures the dynamic features of the data. The order of the model can be decided via a graphical method by plotting the autocorrelation function (ACF) and partial autocorrelation function (PACF). Another method for deciding the order is to use information criteria. Once the order is identified, the parameters of the model are estimated. Finally, diagnostic checking is done by testing the residuals of the selected model for autocorrelation.

An ARMA model is suitable for stationary data. Using non-stationary data for modeling is called the autoregressive integrated moving average (ARIMA) model, where the order of integration is built into the model. For an ARIMA (p,d,q) model, (d) represents the order of integration, (p) represents the lags of AR term, and (q), the lags of MA term. So, if the data turns stationary on first differencing, (d) would be specified as 1. For example, ARIMA (2,1,2) would be used for modeling a time series that is integrated of first order; it will have two lags of AR and MA terms each.

ACF and PACF

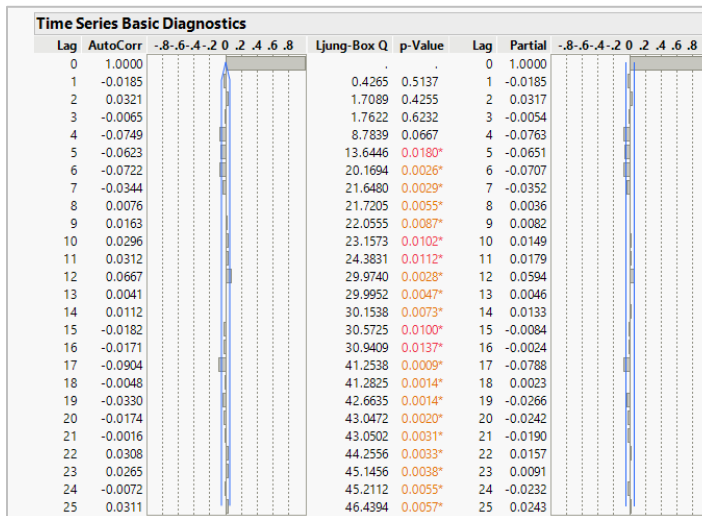
Autocorrelation function (ACF) and partial autocorrelation function (PACF) plots are used to determine the order of the model. The ACF of a non-stationary series does not decay. In case of a stationary series, if ACF geometrically decays and PACF has significant spikes up to a specific lag, an AR model is the best fit. If ACF exhibits significant spikes up to a specific lag and PACF geometrically decays, an MA model is fitted. Alternatively, if both ACF and PACF plots are geometrically decaying, an ARMA model is considered suitable.

Exhibit 6 ACF and PACF Plots for Gold Prices



To create, Analyze>Specialized Modeling>Time Series>Y = GP, X, Time ID = Date. Click OK.

Exhibit 7 ACF and PACF Plots for First Difference of Gold Prices



To create, Analyze>Specialized Modeling>Time Series>Y = FDGP, X, Time ID = Date. Click OK.

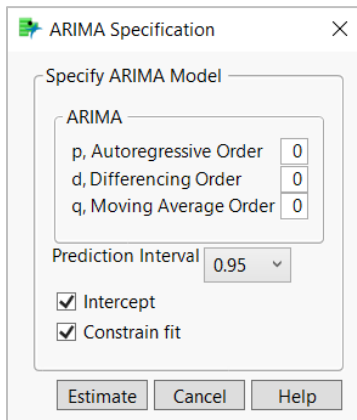
The ACF and PACF plots of gold prices (GP) and first difference (FDGP) are shown in Exhibits 6 and 7.

We can see that the ACF plot of GP in Exhibit 6 does not decay since GP is a non-stationary series. The ACF and PACF of FDGP shown in Exhibit 7 do not show any significant spikes up to the fourth lag. The pattern of decaying of ACF and PACF is also not clear. Unfortunately, while using real data, a clear pattern is rarely seen, making it difficult to interpret ACF and PACF plots. In such cases, the order of the model is determined using information criteria.

Model building

We shall build various ARMA models using the first difference of gold prices (FDGP), which is the stationary series, by specifying the parameters of autoregressive order (p) and moving average order (q). The differencing order (d) is set at 0 for all these models. The models estimated here include AR (1), MA (1), ARMA (1,1), AR (2), MA (2), ARMA (1,2), ARMA (2,1), ARMA (2,2), AR (3), MA (3), ARMA (1,3), ARMA (2,3), ARMA (3,1), ARMA (3,2) and ARMA (3,3).

Exhibit 8 ARMA/ ARIMA Model Building



To create, Analyze>Specialized Modeling>Time Series>Y = FDGP, X, Time ID = Date. Click OK. Under the red triangle next to Time Series FDGP, choose ARIMA for the option to specify the values for p and q . Once you have specified them, click Estimate. Select the ARIMA option again from the dropdown and repeat the process to build different ARMA models.

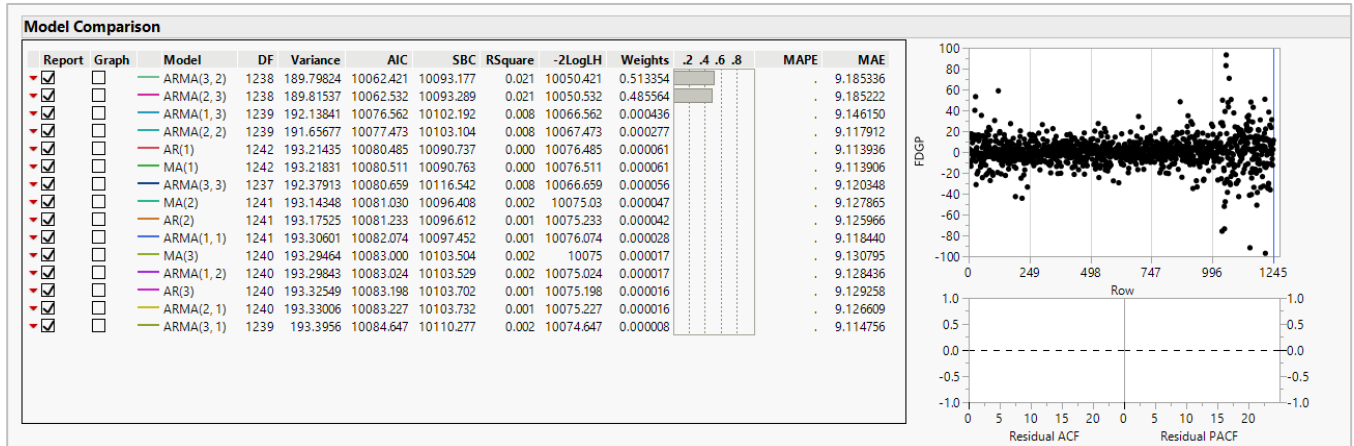
In this case, ARMA models are estimated using the first difference of gold prices (FDGP), which is the stationary series.

Model comparison and identification

Once a model is fit, the Model Comparison report is produced by JMP. This report contains the Model Comparison table and plots for the models. Each time a new model is fit, a new row is added with unique color coding. The Model Comparison table summarizes various fit statistics for each model fitted to the same time series data, which can be used for selecting the most suitable model.

Information criteria are measures of model fit, which include a penalty for adding extra parameters. Hence, the objective is always to minimize the value of information criteria. Among several criteria available, the most popular is the Akaike's information criteria (AIC). The AIC values of these model are compared to identify the most suitable one, which would be the one with the lowest AIC value. By default, JMP sorts the models by the AIC statistic in increasing order. The various ARMA models, along with the fit indices, are shown in Exhibit 9.

Exhibit 9 Model Comparison Using AIC



It can be observed from Exhibit 9 that ARMA (3,2) has the minimum AIC value of 10062.421. Hence, the most suitable ARMA model for gold prices is identified as ARMA (3,2).

Model estimation: ARMA or ARIMA?

Once the model has been identified, the next step is to estimate the parameters of the model. Whether to use an ARMA model or an ARIMA model is irrelevant since both of the models would give the same result. Let us try this by estimating ARMA (3,2) using FDGP and ARIMA (3,1,2) using GP. The order of integration (d) is specified as 1 in the ARIMA model as the gold prices are integrated of first order. The parameter estimates of ARMA (3,2) model and ARIMA (3,1,2) model are shown in Exhibit 10.

Exhibit 10 ARMA (3,2) Model for FDGP and ARIMA (3,1,2) Model for GP

Model: ARMA(3, 2)								Model: ARIMA(3, 1, 2)							
Model Summary								Model Summary							
DF				1238	Stable	Yes		DF				1238	Stable	Yes	
Sum of Squared Errors				234970.224	Invertible	Yes		Sum of Squared Errors				234970.224	Invertible	Yes	
Variance Estimate				189.798242				Variance Estimate				189.798242			
Standard Deviation				13.7767283				Standard Deviation				13.7767283			
Akaike's 'A' Information Criterion				10062.4208				Akaike's 'A' Information Criterion				10062.4208			
Schwarz's Bayesian Criterion				10093.1773				Schwarz's Bayesian Criterion				10093.1773			
RSquare				0.02111571				RSquare				0.99600022			
RSquare Adj				0.01716222				RSquare Adj				0.99598407			
MAPE				.				MAPE				0.64331064			
MAE				9.18533566				MAE				9.18533566			
-2LogLikelihood				10050.4208				-2LogLikelihood				10050.4208			
Parameter Estimates								Parameter Estimates							
Term	Lag	Estimate	Std Error	t Ratio	Prob> t	Constant Estimate	Mu	Term	Lag	Estimate	Std Error	t Ratio	Prob> t	Constant Estimate	Mu
AR1	1	1.502491	0.0913539	16.45	<.0001*	0.18732751	0.65467366	AR1	1	1.502491	0.0913539	16.45	<.0001*	0.18732751	0.65467366
AR2	2	-0.703897	0.0939864	-7.49	<.0001*			AR2	2	-0.703897	0.0939864	-7.49	<.0001*		
AR3	3	-0.084733	0.0329148	-2.57	0.0102*			AR3	3	-0.084733	0.0329148	-2.57	0.0102*		
MA1	1	1.537327	0.0884367	17.38	<.0001*			MA1	1	1.537327	0.0884367	17.38	<.0001*		
MA2	2	-0.778408	0.0744297	-10.46	<.0001*			MA2	2	-0.778408	0.0744297	-10.46	<.0001*		
Intercept	0	0.654674	0.3291510	1.99	0.0469*			Intercept	0	0.654674	0.3291510	1.99	0.0469*		

To create, Analyze>Specialized Modeling>Time Series>Y = FDGP, X, Time ID = Date. Click OK. Under the red triangle next to Time Series FDGP, choose ARIMA. Put p = 3 and q = 2 to estimate the ARMA (3,2) model. Similarly build the ARIMA (3,1,2) model using GP data.

In Exhibit 10, we can observe that the parameter estimates for both the models are the same. Hence, the choice of ARMA or ARIMA is irrelevant in terms of correctly specifying the order of integration. Exhibit 10 also shows that Prob>|t| values for all the AR and MA terms are less than 0.05. Hence, the null

hypothesis that the lags of AR and MA terms are not significant is rejected at 5%. It is evident that all the three AR terms and two MA terms are statistically significant in the model.

Diagnostic checking

Diagnostic checking of the model is done by looking at the autocorrelation of the residuals. Once you build an ARMA or ARIMA model in JMP, it will by default provide residuals information along with model summary and parameter estimates.

Exhibit 11 ACF and PACF Plots of Residuals

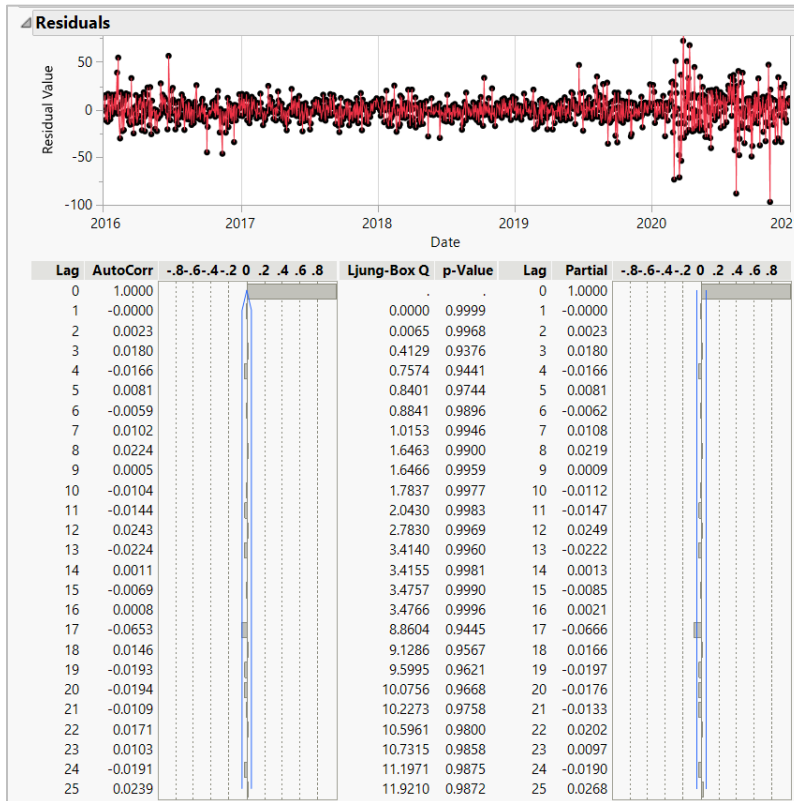


Exhibit 11 shows the ACF and PACF of residuals of the ARMA (3,2) model. It can be observed that there is no autocorrelation in the residuals (as seen in the second column, AutoCorr). Thus, the ARMA (3,2) model is correctly specified and is able to adequately capture the dynamic features of the data series.

Summary

Statistical insights

To summarize, in this case, the scheme of analysis using JMP involved the following:

- Generating summary statistics of the data series.
- Visually representing the data using Graph Builder.
- Differencing the data.
- Testing stationarity using ADF Test.
- Generating ACF and PACF plots.
- Estimating ARMA/ARIMA models.
- Creating model comparisons based on AIC.
- Conducting parameter estimations of ARMA model.
- Using residuals for diagnostic testing.

Implications

Hari can draw the following conclusions from the analysis:

- The gold prices in the United States are non-stationary in nature and exhibited an upward trend during the five-year period.
- The gold price series turned stationary on first differencing.
- An ARMA (3,2) model captures the behavior of gold prices in the United States; the same can be used for forecasting gold prices.

JMP features and hints

This study used the Distribution platform to display histograms and summary statistics; it also used Graph Builder to visualize the data in a time series manner. Time series analysis, which is under Specialized Modeling Platform, was used to conduct an augmented Dickey-Fuller test for stationarity.

Transformations are applied to create new columns, followed by checking for the stationarity of the data series. ARMA and ARIMA models were built by specifying the parameters. A Model Comparison report was used to select the model.

Exercise

The price of silver for a five-year period (2016-2020) is available in SP.jmp. Perform the scheme of analysis explained in this study. Identify the best ARMA model for silver price and perform diagnostic checking.