Dan Silano
Time Series – Summer 2013 Project

1. Introduction

One of the more interesting objects for astronomers to analyze in the night sky happen to be a class of AGN called blazars. These AGN have a hyper-relativistic electron jet emitting synchrotron radiation within a few degrees of the line of sight to the observer. The relativistic beaming effects amplify the natural timescales of light variability and produce interesting variations and trends.

Data was obtained from the Kepler Space Telescope as a result of the guest observer program. The published paper on this object (among others) can be found here: http://adsabs.harvard.edu/abs/2013ApJ...773...89W

2. Data

The data presented consists of photometric flux readings of the object labeled “C” taken over the course of a month at 30 minute intervals.

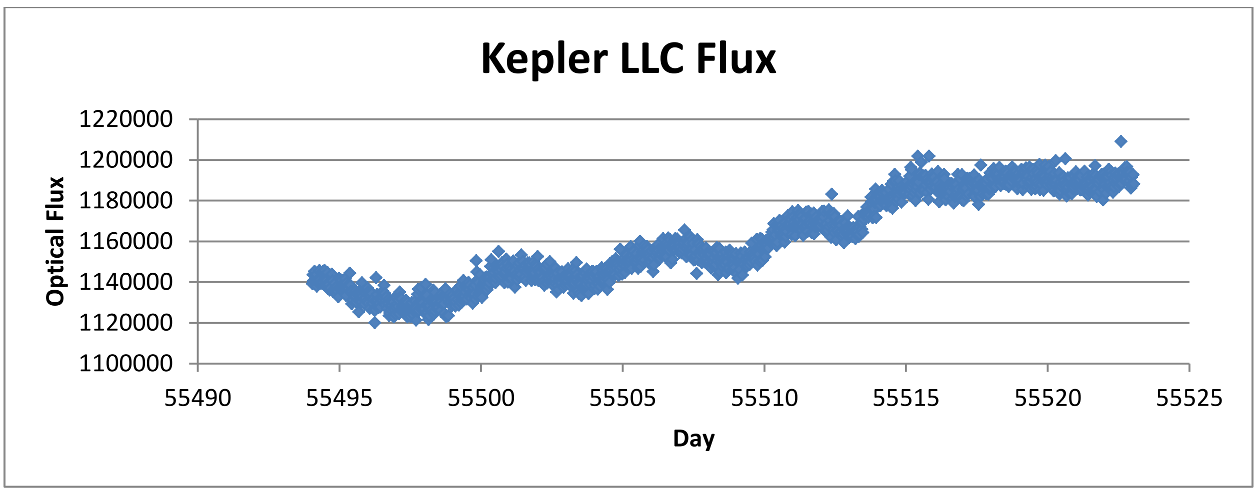

A cursory observation of the data shows that the original time series is not stationary. In addition, there appears to be a roughly linear trend being shown, and hints of ‘seasonality’. Upon closer observation, the periodic variations are not at equal time intervals, even though they appear to be close. Given that the data does not show seasonality, the next step taken was to remove the approximately linear trend.

3. First Differences & Model Choice

The graph of the first differences is plotted below:

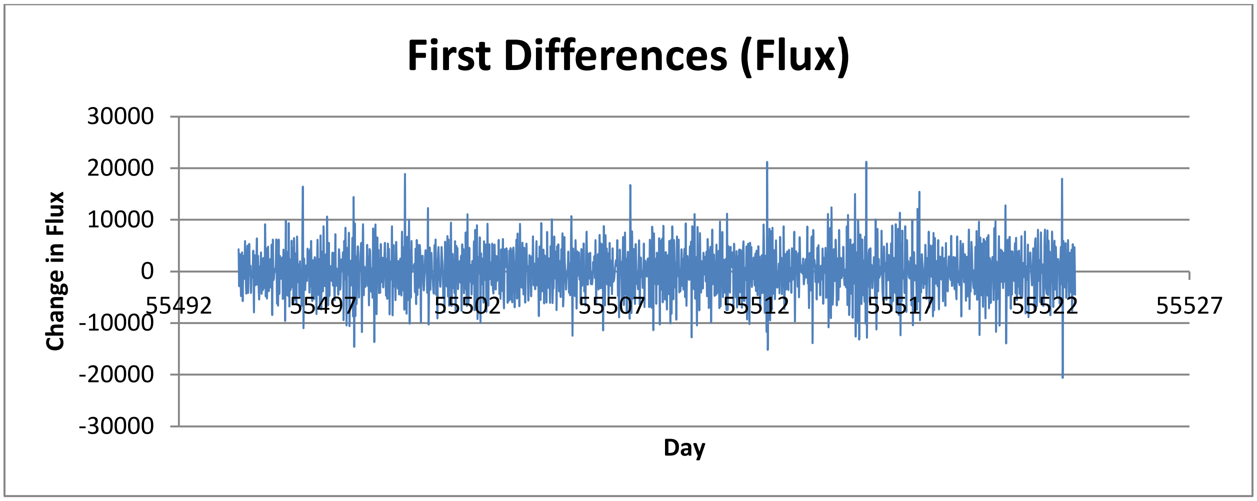

Upon visual inspection, this data series appears to be stationary. The next step was to graph the sample autocorrelation function of the set of first differences. This is pasted below:

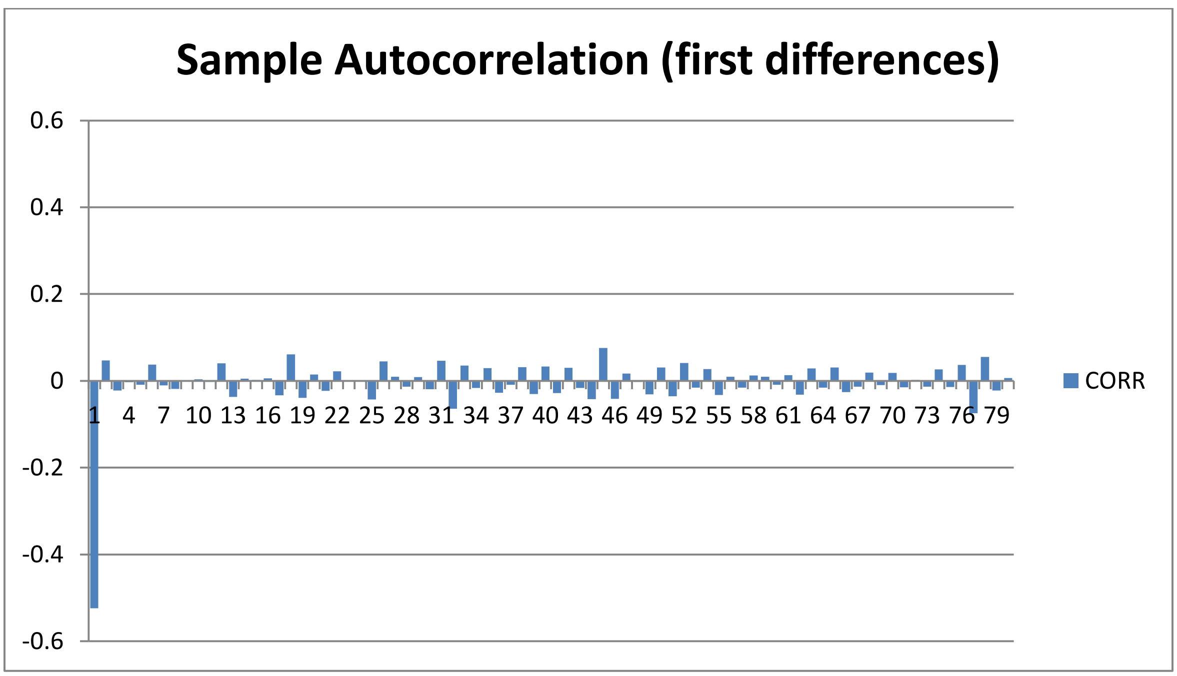

This sample autocorrelation function shows a sharp peak at the first time lag which decays immediately to negligible peaks. This sharp decline is indicative of a stationary process, and the peak at the first time lag is evidence that a MA(1) process will be the best fit.

4. Fitting the MA(1) & Residual Analysis

The model to be fit is a MA(1) process following: $Y_t = \mu + e_t - \theta e_{t-1}$.

The model is fit in excel using a macro looping from θ = -1 to 1 in steps of 0.001. At each step, the sum of the squares of the residuals is analyzed. The value of θ for which the squares of the residuals are minimized was found to be 0.848. In addition, μ was found to be 34.527.

To determine whether the residuals are a white-noise process, three tests were performed: the Durbin-Watson test, a Q-Q plot of the sample autocorrelation function to determine residual normality, and the Ljung-Box test. The computed value of the Durbin-Watson test statistic was 2.095, which proves there is no serial correlation between the 1-period lag of the residuals. The Q-Q plot of the residuals is shown below:

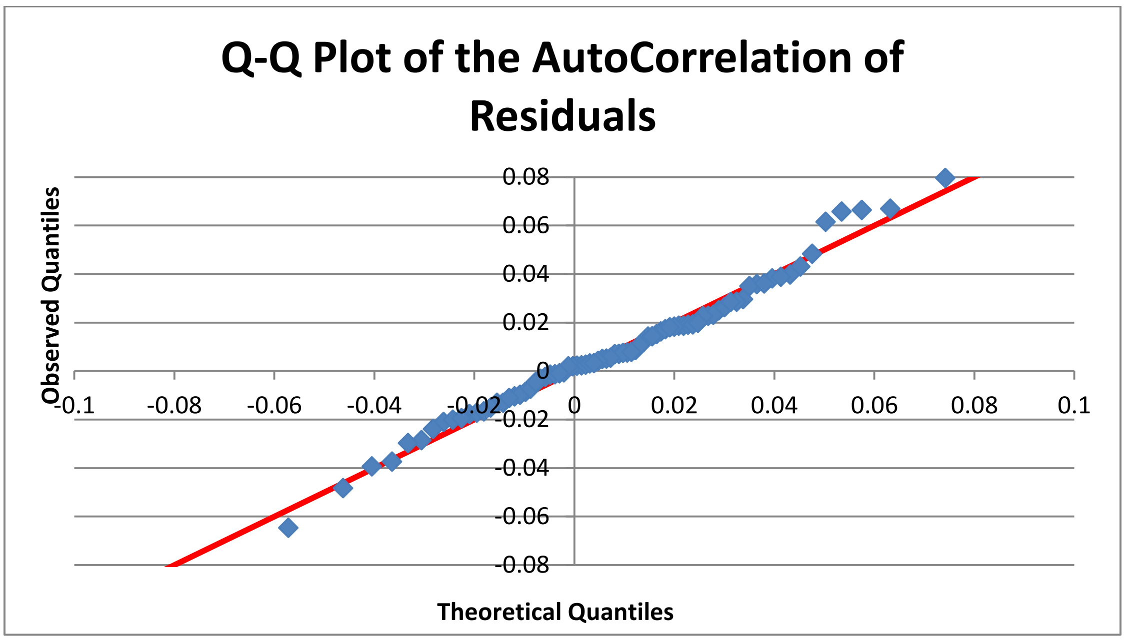

The Q-Q plot shows that the residuals approximately follow a normal distribution. In addition, the standard deviation of the autocorrelation is 0.263, where $1/\sqrt{N}$ = 0.265. Lastly, the Ljung-Box test statistic was computed to be 88.194 using 80 values, whereas the χ^2 distribution given 80 degrees of freedom at the 95% significance level is 101.88, which is greater than the computed value showing that we will not reject the null hypothesis that the residuals are independently distributed. The passing of each of these three tests show that the residuals follow a white-noise process, and validate the choice of the MA(1) model and computed parameter.

5. Conclusions: Sample MA(1) Projections

The final test of the MA(1) model is to do a projection of the future flux using the fitted model parameters. A Box-Muller transform was used to generate Gaussian random numbers which were used to calculate residuals based on a normal distribution (justification from section 4). The resulting residuals were used to calculate the projected first differences, which were then used to back into the flux. The original data is plotted in blue, the projected data is in red: