Demetrios Sofronis, Prof. Alexandra Chronopoulou 12/16/16

Detecting the Change in Memory of a Financial Time Series

Abstract

Our analysis focuses on deciding when there is a significant change in the value of the Hurst exponent, H. The problem arises when we want to conduct any assessment based on the long range dependence in the volatility of the financial time series. Thus, we set out to analyze how the memory changes with time and with the replacement of old data points with new ones. Finally, we are able to statistically pinpoint when each significant change occurs in the variance of H and observe the change point's historical context.

Introduction

After concluding on the presence of long-range dependence in certain financial time series in our project over the summer, this semester, our main focus was to investigate how far in the past the dependence goes or in other words how much data should we use to appropriately estimate the memory parameter. Primarily, we wanted to pinpoint how and when the memory structure changes through time and then understand the reason why. In addition, we wanted to link the historical events that trigger changes in the memory parameter, H.

Preliminary Analysis

Using the whittle estimator, we started by calculating the values for the Hurst exponent using a daily rolling window for the logarithmic squared returns of the S&P 500 and the difference in returns (for stationarity purposes) for the Volatility Index VIX for the same period of time. In order to avoid skewed data from a crash in the stock market due to the 2008 financial crisis, we started by calculating the value for the Hurst exponent right after the crash, reflecting a period with less immediately visible observations. We continued to do so for a yearly interval, shifting the window by one day in the future until the end of 2015. The significance of the rolling window is that it helps us understand which events trigger a change in the value of the Hurst index, H.

As it is clear in Figure 1, the S&P 500 and the VIX share a similar pattern in their dependence structure. Moreover, we proceeded with our assumption that the logarithmic squared returns of the S&P 500 were an accurate proxy to estimate the dependence in the volatility. The graph shows that the two indices have the same change points

Figure 1

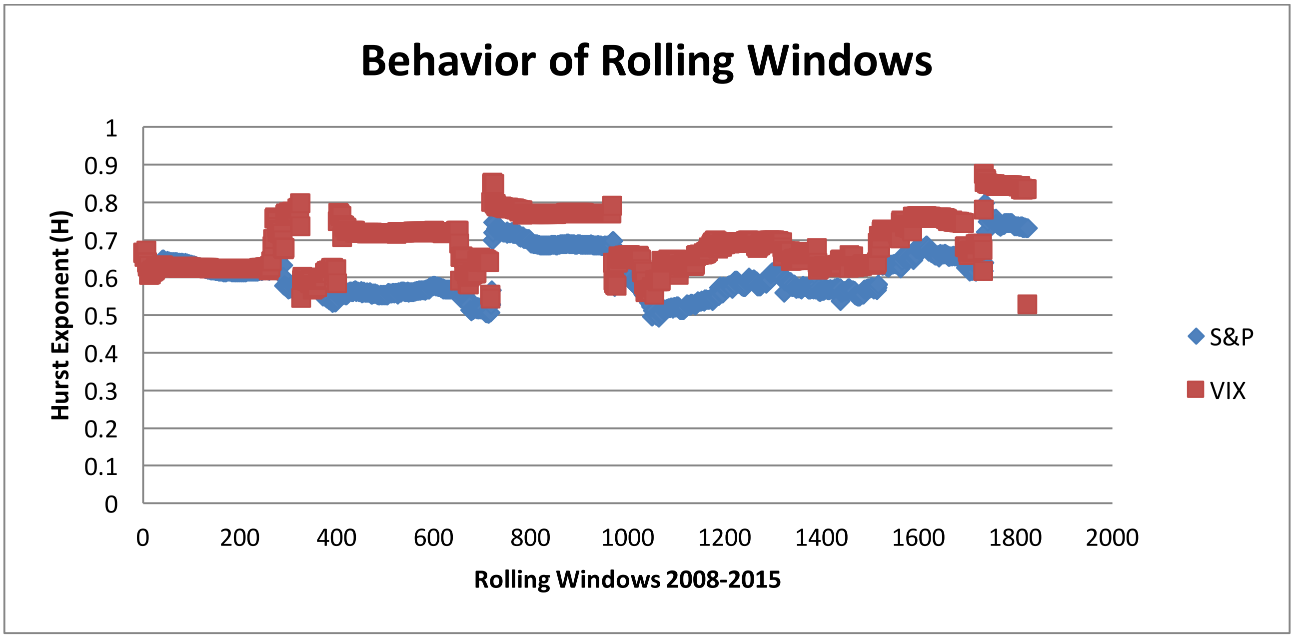

Change Detection

To accurately conclude on where each change point occurs, we used various off-line detection methods. Using the "changepoint" package in R, we were able to analyze our data accordingly. The final result was marking the exact day where a change occurred. Primarily, we conducted our testing using binary segmentation. Essentially, the variances would be cut into sections of relative stability, and we would be alerted when significant changes occurred. Being less hypersensitive than other methods, our testing provided us with what are statistically the most relevant change points.

In Figure 2, we can see a plot of the change points of the Hurst exponent, and their relevant events in figure 3.

Figure 2

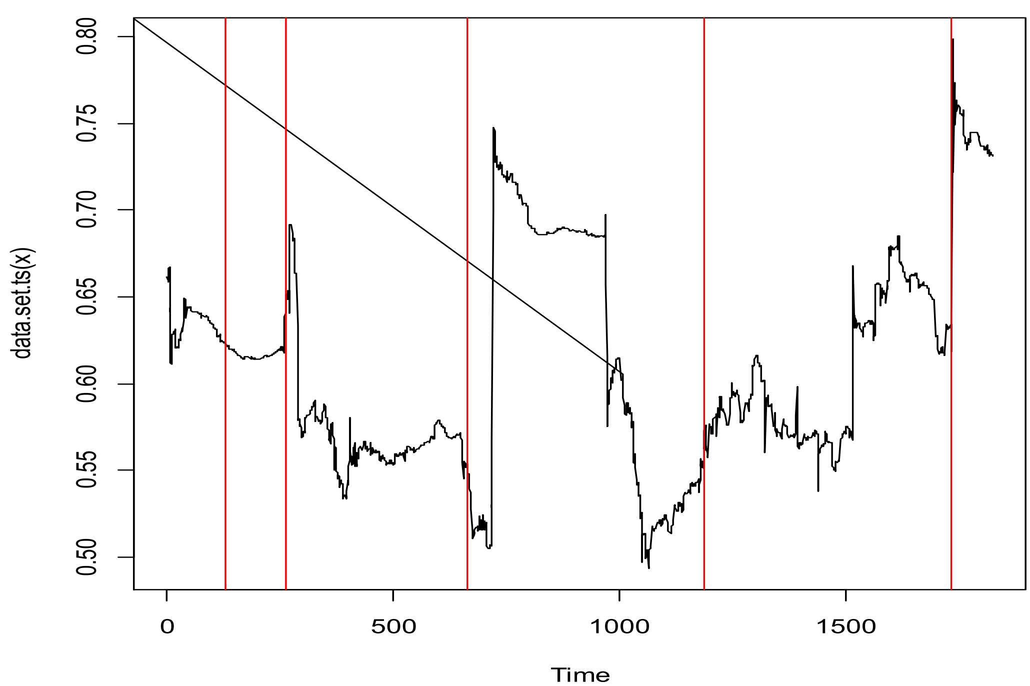

Figure 3

L'Aquila Earthquake in Italy	4/6/2009
News of Stimulus Creating Jobs	10/14/2009
Australian Bank's Credit Cut	5/18/2011
Japanese Economy Officially goes into Recession	5/19/2011
Dow and S&P record worst days of 2013	6/20/2013
Dow in particular Plunges, S&P drops as well	8/21/2015

We saw prolonged intervals of consistency with abrupt changes in between. Certainly, when utilizing the market's memory for prolonged intervals, it is reasonable to infer the persistence of a particular value of the Hurst Exponent. However, the question that arises going forward is how can one accurately assess when to reevaluate for its value?

Significance of Change Points

What grasps our attention is that each change point does in fact have a historical significance as indicated in figure 3. In fact, this change in the variance is when we see times of great uncertainty. Various natural disasters such as earthquakes as well as significant financial movements correspond to the points where each change occurs. In fact, aside from the offsetting of one day of the official Japanese recession, the red markers in the change point graph directly correspond to the dates above. For example, a major change occurred on August 21st 2015 when the Dow plummeted more than 500 points. Such observations allow us to conclude that the markets and their memory are truly correlated to the current events of a time period. With that being said, this allows one to have a better assessment of the Hurst Exponent and when it is applicable to reevaluate. As we would intuitively expect, periods of global uncertainty increase the market uncertainty, reflected in the change points of the Hurst exponent's variance. Such behavior would reasonably call for a recalculation of the memory parameter.

Conclusion

We can deduce a multitude of valuable information from our analysis. First we validated the assumption that the S&P 500 could be a robust means of estimating memory. However, the main question that needed to be addressed was when does one absolutely need to reevaluate their current value for the Hurst Exponent? We answered this question by illustrating how exactly the Hurst Exponent changes through time with the replacement of data points. Interestingly enough, when we add a new point that reflects the data of a day of historical uncertainty, the memory reacts accordingly. This allows one to make the valid assumption that a cyclical reassessment between major current events will provide a more accurate estimate of memory. What our results suggest is an assessment on the efficiency of the market and its ability to immediately reflect the most recent information in the estimation of the memory parameter.