2011

Applying Localized Realized Volatility Modeling to Futures Indices

Luella Fu
Claremont McKenna College

CLAREMONT McKENNA COLLEGE

APPLYING LOCALIZED REALIZED VOLATILITY MODELING TO FUTURES INDICES

SUBMITTED TO

PROFESSOR MARK HUBER

AND

DEAN GREGORY HESS

BY

LUELLA FU

FOR

SENIOR THESIS

SPRING/2011

MONDAY, APRIL 25, 2011

Table of Contents

Introduction

The inspiration for this study comes from the article "Localized Realized Volatility Modeling" by Ying Chen, Wolfgang Karl Härdle, and Uta Pigorsch, in which the authors develop a new approach to forecasting financial volatility. Currently, long-term memory models are popular for forecasting stock and other asset prices, but this study looks at an alternative approach by examining a short-term memory model, the *localized realized volatility model*, also referred to as the *LAR procedure*.

The model approaches the task of forecasting by limiting the interval from which future values are extrapolated to only the set that is closest in time to the realized volatility values it seeks to forecast. The approach, however, seeks to find the longest interval over past data for which even a simple linear regression can be extended into the future. The authors call this the *interval of homogeneity*.

Though clearly a model designed to capture future prices only up to a relatively short period ahead, the authors found that the LAR procedure performed as well as and sometimes better than long-term memory models specifically designed to capture volatility further into the future. In the original literature, the authors also found that the LAR procedure had many useful properties, such as robustness against poorly fit equations to the data, against misspecification of parameters, and against changes in time intervals. These benefits provide the incentive to investigate the localized realized volatility model further.

Description of This Study

While the original study applies the LAR procedure to the S&P 500 futures indices, we decided to apply the LAR procedure to similar indices, those of the Nikkei 225, a stock market index for the Tokyo Stock Exchange, and the CAC 40, a French stock market index. We wanted to see if the LAR process would yield similarly high-performing results.

Limited access to minute-by-minute financial data required us to use a much shorter period than the 20 years of S&P500 data over which the original study was conducted. Whereas the original study used the minute-by-minute data of S&P 500 index futures from January 2, 1985 to February 4, 2005, we restricted our attention to the 163 days of available data from the Nikkei 225 index futures and the 167 days of available data from the CAC 40 index futures. The data range from August 29, 2010 to March 3, 2011 for both stock indexes, and the difference is rooted in a difference of days on which the indices were closed. The smaller data sets meant that we also considered different interval candidates to find the interval of homogeneity. For the Nikkei 225, our potential intervals were [4 days, 7 days, 10 days, 13 days, 16 days, 19 days, 22 days]. For the CAC40, The first interval of 4 days corresponds to the smallest interval length allowed by the original article. Because the LAR procedure is robust against time interval changes though, these modifications should not have had a significantly altering effect on our analysis. It is interesting to note that similarly, the abrupt disruption to the Nikkei 225 data caused by a series of natural disasters Japan faced should also not weaken the LAR process.

It should also be mentioned that due to time and software constraints, we did not repeat many of the comparison tests which Chen, Härdle, and Pigorsch performed. Instead, since the intervals of homogeneity were stressed in the study, and since that concept appears to be the main innovation in the LAR process, we, too, focused our comparative analysis on the impact of the intervals of homogeneity. Specifically, we compared LAR forecast results with forecast predictions made by applying an AR(1) process. The two are related in that the LAR process implemented here uses an AR(1) process, but they differ in that the LAR process derives its parameters from the most relevant interval from the historical set while the AR process derives its parameters from the entire available set.

We also considered the same rolling windows that Chen, Härdle, and Pigorsch used to make their predictions: 1 day, 5 days, and 10 days. Though our set is smaller, 10 days is well within the scope of the intervals we considered.

As in the original study, we used approximately the first fourth of the days from our data as a training set and forecasted the rest of the three-fourths data. This way, we were able to compare LAR predictions against AR predictions and also against actual realized volatility over those times. We make these comparisons by calculating the root mean square error of the difference between our extrapolations and the actual realized volatility . Also note that the final expression of the data for which we make our predictions is as the natural log of realized volatility. We first convert the raw price data into the natural log of realized volatility, which we will notate log(RV). Then, we make predictions for log(RV), and the data we compare is also in terms of log(RV).

More detail on the actual process of generating log(RV) data, creating the intervals, and developing forecasts for the LAR predictions are available in the original article by Chen, Härdle, and Pigorsch (2010).

Findings

Our findings are most concisely stated in the following two tables:

Table 1

Nikkei 225 Index	RSMFE		
	Rolling Window		
	1 Day	5 Days	10 Days
AR	1.0842	1.0984	1.1625
LAR	1.0582	1.0575	1.0581
Interval of homogeneity	4 days		

Table 2

CAC 40 Index	RMSFE		
	Rolling Window		
	1	5	10
AR	1.6054	1.5932	1.7114
LAR	1.3751	1.3845	1.3784
Interval of homogeneity	4 days		

Though our RMSFE statistics agreed with the original study's in the sense that the LAR model overall seemed to more closely predict actual realized volatility, the fact that the interval of homogeneity was the shortest one for both indices is surprising, considering that 6 months was the median interval of homogeneity for the S&P 500 in the original study. Also, it is somewhat surprising that the window which is broader than the interval of homogeneity yielded a lower RMSFE for both indices.

When we graphed the LAR and AR predictions against the actual log(RV), it became even more apparent that there was an issue with the interval of homogeneity. It is displayed below:

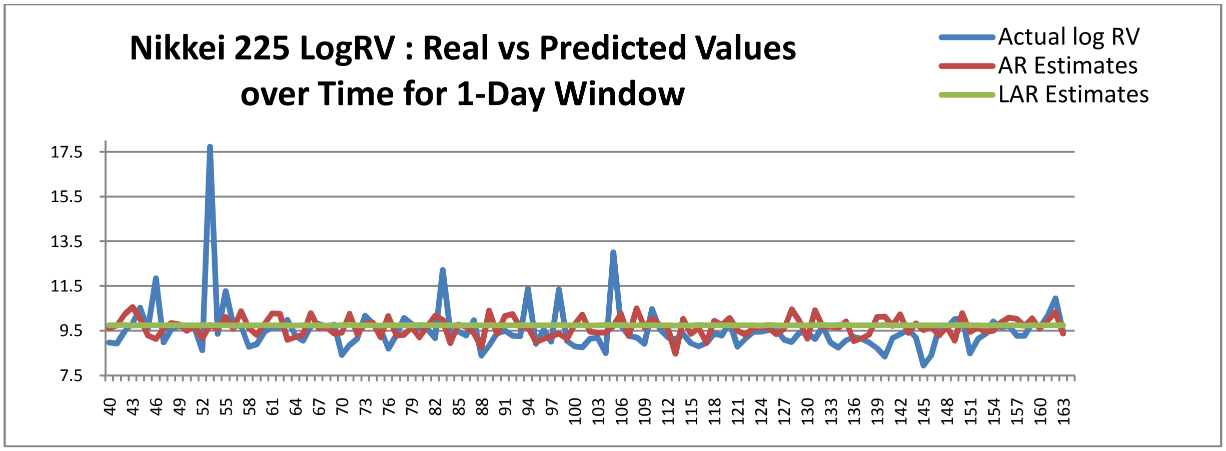

Figure 1: Log(RV) is plotted on the vertical axis, and day number (counting from the first day of data on August 29, 2010) is on the horizontal axis.

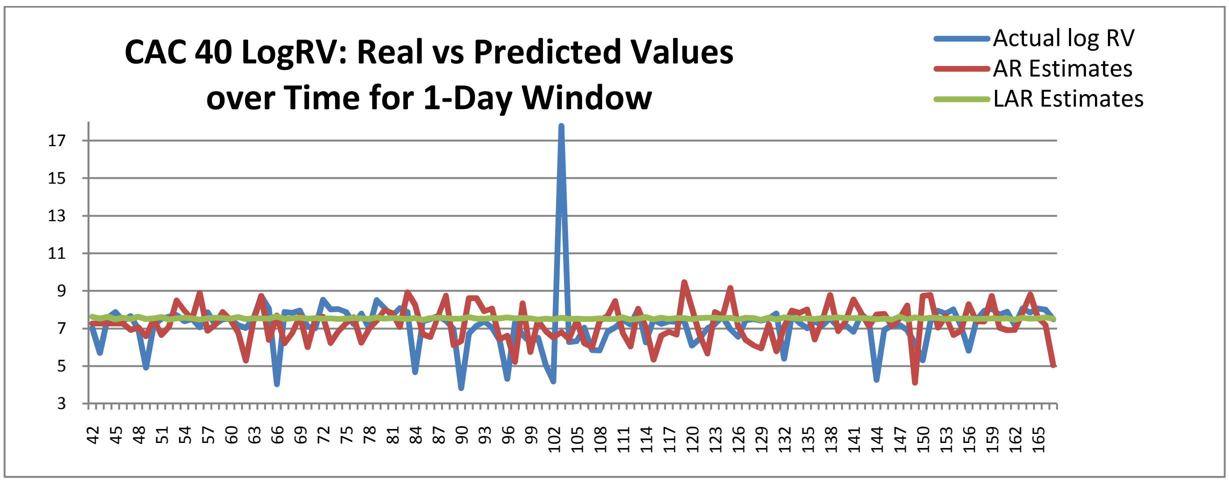

Figure 2: Log(RV) is plotted on the vertical axis, and day number (counting from the first day of data on August 29, 2010) is on the horizontal axis.

It is clear that the few number of data points in the interval of homogeneity gave rise to a small slope and a small variance, which in turn yielded an almost flat linear model for our LAR process when applied to our indices.

The graph for a rolling window of 5 days and a rolling window of 10 days look similar and are included at the end of this article in the section called, “Additional Graphs”.

Conclusion

Though it is interesting to observe how on average, a forecast that clearly performs very poorly can still appear to outperform a forecast that better matches actual data, our results were inconclusive due to the fact that the issue concerning the interval of homogeneity could not be resolved. Most likely, the error is specific to this study and lies either in the calculations for finding critical values, which determine the longest interval of homogeneity, or in the algorithm for iteratively accepting larger intervals past the first one. The problem would have to be further investigated.

Additional Figures and Information

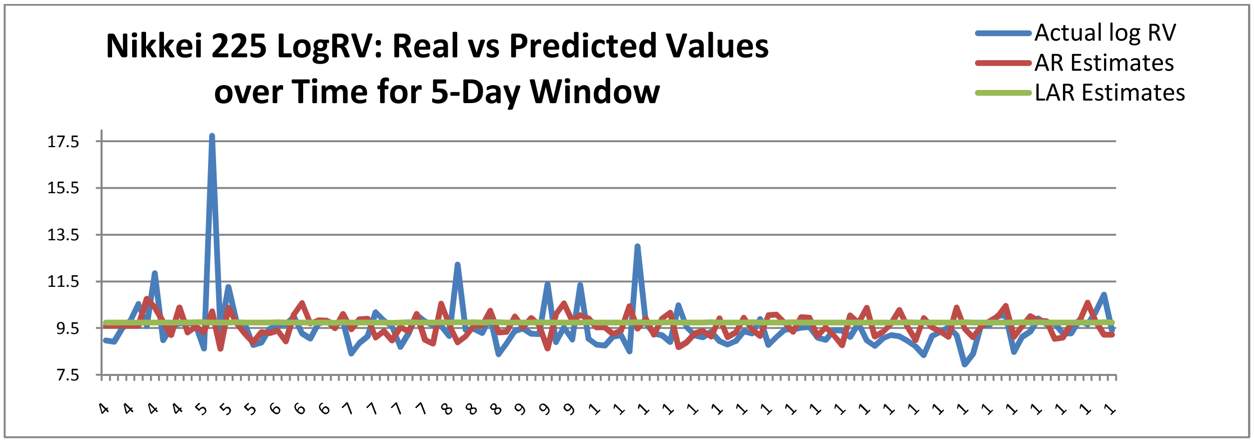

Figure 3: Log(RV) is plotted on the vertical axis, and day number (counting from the first day of data on August 29, 2010) is on the horizontal axis.

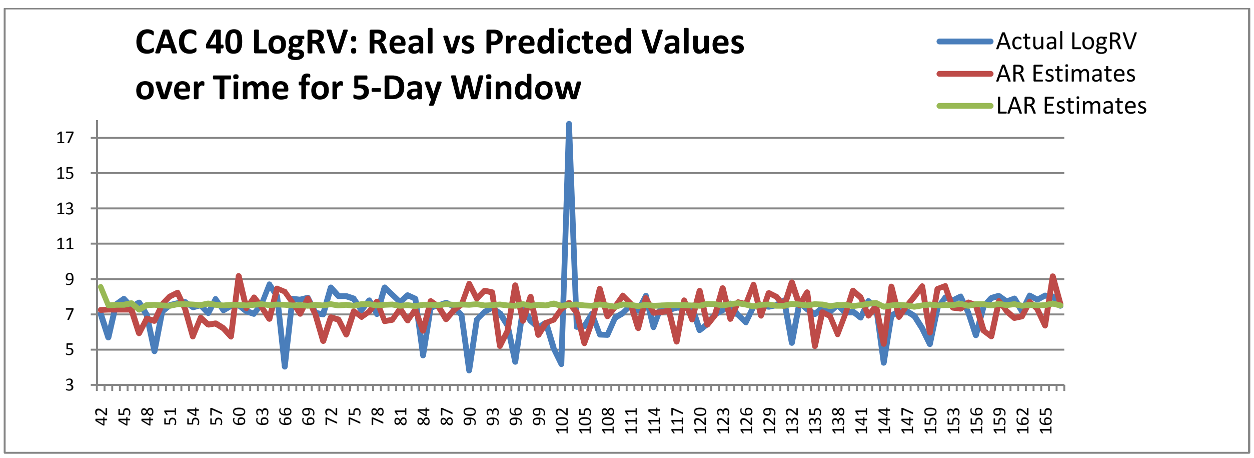

Figure 4: Log(RV) is plotted on the vertical axis, and day number (counting from the first day of data on August 29, 2010) is on the horizontal axis.

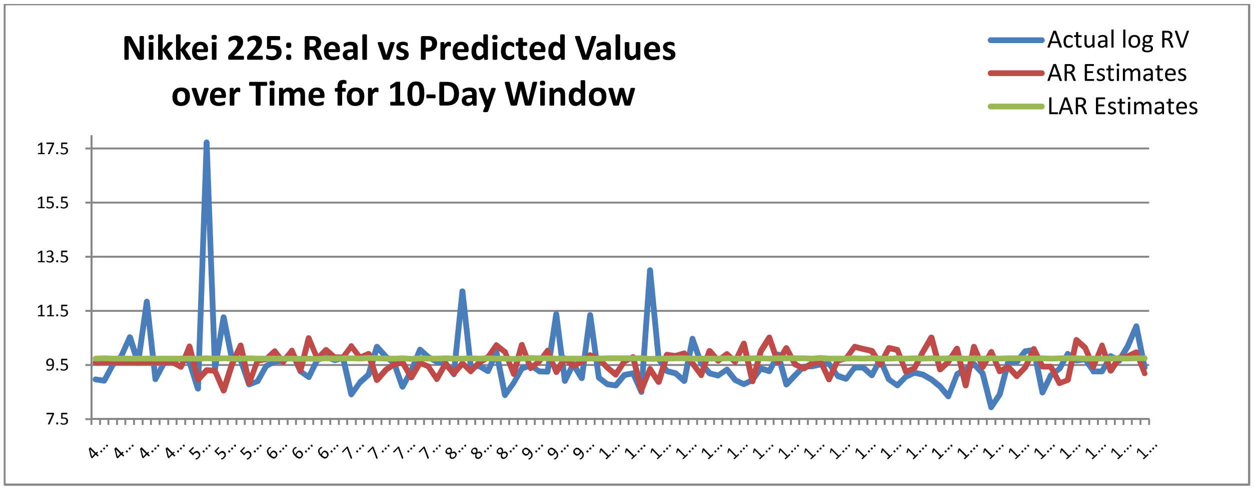

Figure 5: Log(RV) is plotted on the vertical axis, and day number (counting from the first day of data on August 29, 2010) is on the horizontal axis.

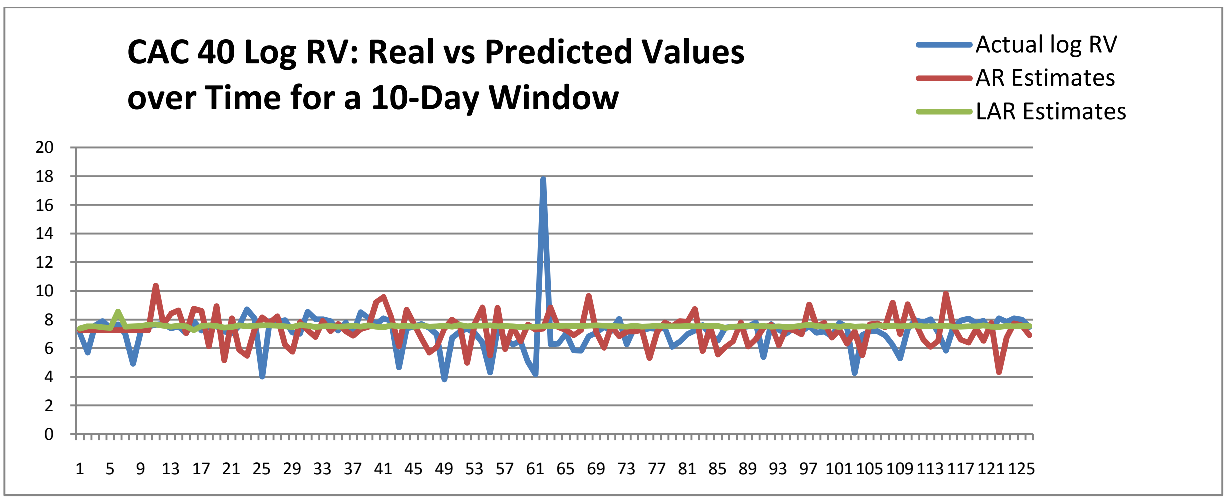

Figure 6: Log(RV) is plotted on the vertical axis, and day number (counting from the first day of data on August 29, 2010) is on the horizontal axis.

Bibliography

Chen, Ying, Wolfgang K. Härdle, and Uta Pigorsch (2010). "Localized Realized Volatility Modeling." *Journal of the American Statistical Association,* 105, 492. [1376-1393]