ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

WHAT DOES THE VIX ACTUALLY MEASURE? AN ANALYSIS OF THE CAUSATION OF SPX AND VIX

Merav Ozair†

† Department of Finance and Risk Engineering, Polytechnic Institute at NYU, USA

Abstract. *We examine the causality relationship between the S&P500 (SPX) and the VIX. Our contention that a circular mechanism which "feeds" itself that can be explained by "cause and effect", is supported by the empirical findings on the intraday, minute bar, time series of the SPX and the VIX. The findings are supported across different samples and estimation models and show that: (1) the SPX shock to the VIX time series is not only significant but also persistent; (2) the VIX follows a serial pattern of significant reversal (in the first lag) followed by momentum in the subsequent lags (and beyond the first 10 minutes); (3) the VIX endures a "permanent" market impact, while the SPX sustains a "transitory" one; and (4) the SPX shock on the VIX system remains in the system long enough to account for 70% of the variance of the VIX, suggesting a predictive power of the SPX to the current movement of the VIX.*

Keywords: *causality, Vector Autoregressive (VAR), Volatility Index (VIX), S&P500 Index (SPX), shocks, market impact, reversal, momentum, autocorrelation, cointegration*

JEL Classifications: C58; C55; G02; G19

Introduction

Understanding market volatility has long been a quest of both researchers and practitioners. In 1993, the Chicago Board Options Exchange (CBOE) introduced the CBOE Volatility Index (VIX), originally designed to measure the market's expectation of 30-day volatility, implied by at-the-money S&P 100 Index option prices. In 2003, the VIX was updated to reflect a new way to measure volatility, based on the S&P500 Index (SPX) and estimates expected volatility by weight- averaging a wide range of strike prices of put and call options on the SPX. Principally, the VIX supposed to capture the future volatility of the SPX, and hence predict the future movement of the S&P500. However, does the VIX actually represent the future direction of the SPX in current market conditions? This is the primary question that we attempt to answer in this study.

It has well been documented that implied volatility is a reasonable forecast of future realized volatility (e.g., Granger and Poon (2003), and Anderson, Bollerslev, Christofferson and Diebobld (2005) – for historical review). Surprisingly few prior studies deal with the topic of this paper - the possible relationship between implied volatility and future stock returns. Yet, market participants, in particular traders, are well aware of it. A widespread belief among them holds that swings in implied volatility value are associated with fear in the market, whereas a decline indicates complacency. As a measure of fear and complacency, implied volatility is often used as a contrarian indicator: prolonged and/or extremely high VIX readings indicate a high degree of anxiety – or even panic – among option traders, and are regarded as a bullish indicator1. Prolonged and/or extremely low readings indicate a high degree of complacency, and are generally regarded as a bearish indicator. In 2008, the VIX had remained in the low 20's when all knew that problems were spinning out of control, and later in the year spiked, correcting its previous assumptions. It spiked, however, far beyond reality as panic drove option premiums (i.e., insurance prices) into the stratosphere. It again (quite quickly) over compensated as it fell. The VIX suffered huge whipsaws in 2009, 2010 and 2011 trying to over compensate and find some realm of equilibrium between perception and math.

The literature examining VIX primarily focus on the predictive power of the VIX, and hence assumes that the VIX actually measures the forward 30-day volatility of the S&P 500 and can predict the S&P500 movement a month ahead. Doran, Goldberg and Ronn (2008) use the VIX as a proxy for the S&P 500 Index volatility in their time-varying expected S&P 500 return model. Bekaert and Hoerova (2013) decompose the square VIX index into the Conditional Variance of stock returns (CV) and the equity variance premium (VP). Using different volatility forecasting models they examine the predictive power of the VIX and its two components: stock market returns and economic activity. They find that the VP is a significant predictor of stock returns, but the CV mostly is not. CV, however, robustly and significantly predicts economic activity with negative sign, whereas VP has no predictive power for future ecomonic output growth. Others study the VIX in comparison to the GARCH model and investigate whether different types of GARCH models fit the CBOE VIX. (Hao and Zhang (2013), Lui and Qiao (2012)). The key finding in this line of research is that the GARCH model (examined with different types of GARCH) consistently underestimates the VIX, and concludes that the discrepancy is due to the variance risk premium not captured by the GARCH.

These studies not only assume the predictive power of the VIX, but also assume it is "priced" correctly. The VIX is a measure of implied volatility and is based on options prices (i.e., the option premium – how much market participants are willing to pay for protection from market movements). In that sense, we assume that market participants "know" the "correct" price of the option. If we believe this is indeed true then we can also believe that the implied volatility represented by the VIX, calculated from option prices, is also a correct measure of the future volatility of the market and can thus predict the movement of the S&P 500.

1 On the day of the "flash crash" when the Dow plummeted 900 point just to recover in minutes, the VIX spiked more than 60% , supporting the "fear indicator" hypothesis.

ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

There is (to some extent) a circular mechanism that "feeds" itself. The VIX, as a measure of implied volatility, is a function of the options and the underlying (i.e., the SPX). When market participants view the value of the quoted VIX they may react to it either with adjustments to their option prices or with actions in the market (i.e., the S&P 500). This mechanism may explain why during the years between 2008 and 2011, the VIX over/under reacted and then corrected itself. In order to decipher this mechanism and better understand the "cause and effect" relationship we postulate two main (and one secondary) hypotheses: (1) If VIX is a forward looking measure of the S&P500 future volatility, we would expect a leading relationship, meaning the VIX movement leads the S&P 500 and hence, we would expect that VIX "granger causes" the S&P 500 Index; (2) the VIX measure is a function of the S&P 500, and hence implicitly determined by the values of the S&P 500 Index. Therefore, this type of relationship implies that the S&P 500 "granger causes" the VIX; (3) Secondary to the main hypotheses, a third hypothesis states a bi-directional causality relationship between the VIX and the SPX(but also postulating that the impact of the S&P 500 Index (SPX) is the stronger and the more significant of the two.)

This paper is the first to examine causality between the implied volatility, the VIX, and its underlying, the SPX, and sheds light on the implication of the VIX calculations. Previous studies have examined the correlation between VIX and SPX (Zheng (2012), Brenner, Shu and Zhang (2010), Carr and Wu (2006) and Whaley (2008)). Whaley (2008), in his simple regression analysis of daily changes of the VIX to daily changes of the SPX and a conditional rate of change in the S&P 500 on the market going down or up, finds that: (1) negative relationship of change in VIX to change in S&P 500; (2) the relation in the VIX and the SPX is asymptotic; (3) VIX is more a barometer of investors' fear of the downside than it is a barometer of investors' excitement (or greed) in a market rally. The evidence in Whaley's study merely documents correlation and is not intended to express causality.

We examine the intraday interaction of the two minute bar time-series of the VIX and the SPX[2]. Our four main findings offer a new insight as to the interactions of the VIX quotes with the market SPX movement and to the implications of the VIX calculations. First, SPX significantly and robustly "granger causes" the VIX. This causality test is supported and is evident not only in any sample examined but also by using different types of analysis. The VIX causality, however, even though indicated by the Granger Causality test that we have a bi-directional causality relationship, is not supported by any other tests, such as the estimated VAR model coefficients but in particular the Impulse Response and the Variance Decomposition analyses. These two tests evidently show that the SPX is certainly not affected by the VIX, whereas the VIX, on the contrary, is significantly affected by the SPX and about 40% of its variance can be explained by the SPX movement.

Second, we observe a pattern in the minute returns/level time series and especially in the VIX time series. The SPX seems to strongly and positively relate to its first lag. This impact dies, on average after the fourth lag, and the significance of the coefficients on the second to fourth lags is reduced significantly. The VIX, however, is significantly related to all lags

2 Ozair (2011) examine the interaction of the daily time series of the SPX and VIX and find a consistent and significant ganger causality of the SPX on the VIX, and no opposite causality relationship.

estimated in the model and follows a pattern which can be interpreted as a correction followed by a momentum, whereas the magnitude of the correction (observed in the first lag) is, on average, about six fold compared to following lags. Hence, one can explain the VIX minute changes or level values, that on average there is an immediate correction (or reversal) and then some element of momentum (which will depend on the magnitude of the change in the VIX for that previous minute). This can explain the over/under reaction of the VIX and the correction that follows, which implies that VIX's first reaction tends to over/under estimate the shock/news.

Third, in every sample and estimation method (VAR or VECM; returns versus levels) we observe that the current level of VIX is related/affected by all lags (with the exception of the second lag) included in the estimation model. This implies that the VIX time series is much more autocorelated than the SPX. Any shock to the SPX will die relatively quickly, while the VIX will carry on the impact of a shock for a relatively long period of time. In Market Microstructure literature, we can refer to that as a "permanent" market impact – VIX has a "permanent" market impact whereas the SPX market impact seems to be more transitory. This observation is of particular interest when developing best executions strategies and optimizing transaction costs3.

Forth, there is a cointegration relationship of first order between the VIX and the SPX time series. The main finding when analyzing the VECM model lies in the Variance Decompositions. This variance analysis shows that in the first period following the shock the decomposition of variance for the VIX is 30% - 70%, explained by the SPX and VIX respectively. This decomposition, however, flips with time – i.e., by the twelfth period (and even sooner) the decomposition become 70%-30% explained by the SPX and VIX, respectively. This suggests that the SPX shock stays in the VIX system longer, which will also suggest that one can "predict" the direction (and to some degree the magnitude) of the VIX in the next ten minutes simply by observing the current movement of the SPX. This is a paramount observation, if one wishes to make trading decisions (whether as a trader or a hedger or an investor). It also suggests that what the current level of the VIX captures is merely the current changes in the SPX and has no predictive power or assessment on the future movement of the SPX.

Collectively, these findings contribute to the growing literature on the Variance Risk Premium and the VIX. Our paper may offer an explanation (or alternative explanation) to the aforementioned research findings – as to what it measures, what it may predict (if anything) and its validity and correctness. For example, the inability to reconcile the discrepancies of the GARCH model with the implied volatility of the VIX can be explained by the pattern of the VIX time series (as it reacts to shocks to the system) that is observed in this study.

The reminder of the paper is organized as follows. Section 2 describes the data and provids descriptive statistics and explanation on special irregularities within the data sample (section 2.1); and descriptions of the research method and some hypotheses (section 2.2). Section 3 reports the empirical analysis and the results for daily sub-samples and for the whole sample

3 This should be examined on the tradable vehicles of the VIX, such as the ETN's on the VIX and the VIX futures. It is very likely that one should, observe the same pattern with the VIX tradable vehicles as they are its derivative.

ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

period. Section 4 provides further analysis and examination of the results discussed in section 3. Section 5 illustrates the applicability of the findings in the study and section 6 concludes.

Data and Estimation Procedures

Data: Description, Cleaning and Statistics

To examine the relationship between SPX and VIX, we use the Bloomberg data4 (level I) intraday tick and minute bar. The data have been retrieved from the Bloomberg terminal and have been collected incrementally over time5. We have collected over a year worth of data from August 9th 2012 to October 3th 2013 for the tick data and for the minute bar data from October 5th 2012 to August 9th 2013 (see information in table 1). For the SPX and VIX we have 1,415,935 and 464,215 observation for the tick data, respectively; and minute bar data 100,323 and 99,950 observations, respectively. The data consists of 421 calendar days but only 289 trading days due to weekend and national holidays where the market is closed and two special days (which are unique to this period of time) due to hurricane Sandy (in total 12 days where the market closed other than on weekends). There were also three half trading days, when the market closed at 1pm a day before a major national holiday (see table 2 for details).

The tick data include date and timestamp and price (i.e., the level quote of the index). The S&P 500 index is quoted every five seconds during the trading hours of the day (i.e., 9:30am to 4pm)6, hence the tick data has 12 entries during each minute. The VIX, on the other hand, is quoted only every 15 seconds during the trading hours of the day, and therefore its tick data has 4 entries during each minute. The S&P 500 index quotes starts exactly with the opening of the financial market (i.e., 9:30am), the VIX, however, begins its quotes, for most days7, during the second minute of the trading day. Some minutes have different number of ticks, are very few and negligible, and account for less than 0.001% and 0.01% for the SPX and VIX respectively for the entire sample8.

4 Bloomberg receives its market data through NYSE data feed.

5 Historical intraday data is available for download from Bloomberg terminal a maximum of 140 days ago.

6 For more information see http://us.spindices.com/indices/equity/sp-500

7 Out of 289 trading days in the sample, we observe the one minute delay in 279 trading days (i.e., 96.5% of trading days); for more information see table 1.

8 More information is available from the author.

Table 1: Sample Data Summary

Data Type	Sample Period	Observations (1)	Calendar Days	Trading Days (2)
Historical Intraday Tick (3)	8/9/12-10/3/13	SPX: 1,415,935 VIX: 464,215	421	289
Historical Intraday Minute Bar (4)	10/5/12-10/3/13	SPX: 100,323 VIX: 99,950	364	249
Inferred Intraday Minute Bar (5),(6)	8/9/12-10/3/13	SPX: 118,246 VIX: 116,101	421	289

Table 2: Dates Market Closed During the Sample Period

Sep 3 2012	Labor Day
Oct 29 2012	Hurricane Sandy
Oct 30 2012	Hurricane Sandy
Nov 22 2012	Thanksgiving
Dec 25 2012	Christmas
Jan 1 2013	New Year
Jan 21 2013	MLK Day
Feb 18 2013	Washington's Birthday
Mar 29 2013	Good Friday
May 27 2013	Memorial Day
July 4 2013	Independence Day
Sep 2 2013	Labor Day

Notes:

The discrepancy in number of observations between SPX and VIX is 405: 126 due to missing intraday VIX data during trading hours; 279 due to 1-minute delay of VIX at the beginning of the trading day.

Including 3 days where market closed at 1:00pm - 11/23/12, 12/24/12, 7/3/13

From Bloomberg Data Services Level 1

From Bloomberg Data Services Level 1

Inferred from Historical Intraday Tick

In all subsequent analysis, Intraday Minute Bar observations are limited to 9:30 AM to 4:01 PM of each trading day in the sample period, yielding 112748 SPX observations and 112343 VIX observations.

From the above description, it is clear that the entries of tick data for the VIX and SPX are not inline. For the purpose of conducting the analysis we had to synchronize the two time series (i.e., SPX and VIX), and chose to use the minute bar data9. To work with minute interval, we have transformed the tick data to minute intervals, after the transformation we ended-up with 118,246 and 116,101 observations for the SPX and the VIX, respectively. We were able to retrieve minute data for Bloomberg from October 5th 2012 to October 3rd 2013. For the period October 5th 2012 to October 3rd 2013 we have compared our calculated minute entries with those quoted on Bloomberg terminal (see Table 3). There are very few discrepancies which account for less than .01% for both the SPX and the VIX, and the absolute value of differences with respect to Bloomberg's data is also less than .01% for both the SPX and the VIX. Only in the case where there are discrepancies in the number of ticks per minutes, it seems that in these cases the Bloomberg tick data tend to be underestimated. The results of the comparison have supported our calculation of the minute bar and thus we will be using our calculated minute bar sample for the rest of the analysis in this research.

The "after trading hours" data (i.e., after 4pm) experience many irregularities which primarily stems from lack of liquidity (i.e., lower trading volume, larger quote spread and higher volatility) but also from computers' delays. For this reason we have restricted our sample for each trading day to "normal" trading hours, 9:30am to 4:01pm (taking in consideration some adjustments that might occur after the close of the markets). The analysis thereafter is performed on "normal" trading hours10.

Intraday data exhibits irregularities such as, duplicate observations11, large sequence of missing data and significant outliers. We have addressed each issue with accordance to the treatment documented in the market microstructure literature and verified that our results are not a consequence of special irregularities that are associated with this specific dataset. Cleansing is an important aspect of computing realized measures. The literature suggests that when there are mis-recordings of prices or hit large amounts of turbulence at the start or the end of the trading day then they may sometimes give false signals. Barndorff-Nielsen, Hansen, Lunde and Shephard (2009) have studied systematically the effect of cleaning on realised kernels, using cleaning methods which build on those documented by Falkenberry (2002) and Brownless and Gallo(2006).

9 It is quite common in the market microstructure literature to use 1 minute or 5 minute bar as the appropriate data frequency for the analysis (see Madhavan (2000))

10 This is consistent with the market microstructure literature (see for example Madhavan (2000)).

11 By duplicates, we mean that more than one tick entry at the same second (not necessarily with the same value.)

Table 3: Minute Bar Discrepancies (1)

A comparison of minute discrepancies in the Historical Intraday Minute Bar and the Inferred Intraday Minute Bar for SPX and VIX, in the sample period Oct. 5 2012 to Oct. 3 2013. In this table, the "General" row represents the number of historical minute bar observations with one or more discrepancies. Each subsequent row indicated the number of discrepancies per variable.

Sample period: Oct. 5 2012 – Oct. 3 2013

	SPX				VIX			
	# Minutes with Discrepancy	Sample size	Accuracy(3)	Average difference(4)	# Minutes with Discrepancy	Sample size	Accuracy	Average difference
General	53	100323	99.9472%		44	99950	99.9560%	
OPEN **	18	100323	99.9821%	0.0042%	12	99950	99.9880%	0.0509%
HIGH	18	100323	99.9821%	0.0057%	11	99950	99.9890%	0.0782%
LOW	15	100323	99.9850%	0.0087%	9	99950	99.9910%	0.0556%
LAST_PRICE	18	100323	99.9821%	0.0082%	8	99950	99.9920%	0.0938%
NUMBER_TICKS	50	100323	99.9502%	117.4067%	43	99950	99.9570%	107.0833%

Notes:

"minute discrepancy" is defined as when any variable of historical and inferred values of Intraday minute bar data are not equivalent. The historical minute bar data source includes the following variables: OPEN, HIGH, LOW, LAST_PRICE and NUMBER_TICKS.

For the definition of "minute bar" and "inferred minute bar" see Table 1.

Accuracy is percent of total minute observations, generally and per variable, without discrepancies in the sample period.

1 The Average Difference is the average of the absolute value of the differences between the number of Historical Intraday Minute Bar variable discrepancies and the number of Inferred Intraday Minute Bar variable discrepancies. Each difference is normalized (divided) by the number of Historical Intraday Minute Bar variable discrepancies.

Table 4: Frequency of Values for VIX and SPX.

Sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day.

Number of ticks different than 4 for the VIX					Number of ticks different than 12 for the SPX			
Value of the tick	Frequency	W/ complements (1)	Percentage (2)		Value of the tick	Frequency	W/ complements (1)	Percentage(2)
1	18	0	0.00%		1	1	1	100.00%
2	12	1	8.33%		2	1	1	100.00%
3	124	33	26.61%		3	1	0	0.00%
5	22	21	95.45%		5	2	1	50.00%
6	1	1	100.00%		6	3	0	0.00%
8	2	2	100.00%		7	1	0	0.00%
Total number of minutes for the VIX			112,343		8	1	0	0.00%
Number minutes with 4 ticks			112,164	99.84%	9	3	0	0.00%
Number minutes that don't have 4 ticks			179	0.16%	10	3	0	0.00%
With complements			58	0.05%	11	660	586	88.79%
Without complements			121	0.11%	13	843	586	69.51%
					19	1	1	100.00%
					22	1	1	100.00%
					23	1	1	100.00%

Number of ticks different than 4 for the VIX				Number of ticks different than 12 for the SPX			
Value of the tick	Frequency	W/ complements (1)	Percentage (2)	Value of the tick	Frequency	W/ complements (1)	Percentage(2)
				Total number of minutes for the SPX			112,748
				Number minutes with 12 ticks			111,226 98.65%
				Number minutes that don't have 12 ticks			1,522 1.35%
				With complements			1,178 1.04%
				Without complements			344 0.31%

Notes:

Defined as the minute together with its immediate preceding or succeeding minute averaged at 4 or 12 ticks

Percentage of minutes with complements relative to that particular frequency group

Number of total ticks per day is 4,704 (12 ticks per minute and 392 minutes) and 1,616 (4 ticks per minute and 391 minutes) for the SPX and VIX, respectively. This total number of ticks per day is stable 99.97% and 99.99% of the sample for the SPX and the VIX, respectively. Careful look at the data reveals that a "duplicate" may not be a "real duplicate" but simply an adjustment of missing consecutive entries in the preceding minute or the following minute. The SPX and the VIX by construction of the index12 should be quoted every 5 and 15 seconds which correspond to 12 and 4 entries per minute, respectively. If a minute records less than the amount of quotes required by design of the index, that implies that the "delay" has happened due to technological problems with the reporting system either on Bloomberg or the Exchanges, and therefore at the beginning of the following minute an adjustment to these missing ticks appears to complement and reflect the total 12 tick for the SPX or the 4 ticks for the VIX. (For example, a duplicate for the SPX on 1/25/2013 at 11:11:05am of 7 reflects the complement to the merely 5 ticks recorded in the previous minute 11:10am of that day13.) We can then conclude with 99.9% confidence that the data does not consist of any "real duplicate" but simply corrections to technology mishaps.

Table 4 documents the minutes in the sample that has number of entries which are different than 12 for the SPX and 4 for the VIX. It shows that the SPX has 1,522 minutes (i.e., 1.3% of total minutes sample) that have number of entries different than 12; and the VIX has 179 minutes (i.e., 0.16% of total minute sample) that have entries different than 4. Most of these differences are complemented and adjusted with respect to the preceding or the following minute to form that on average the number of ticks per minutes is 12 or 4 for the SPX and VIX, respectively. The percentage of minutes with number of ticks different than 12 for the SPX and different than 4 for the VIX, which are not complemented with its previous or following minutes is 0.3% and 0.1% respectively.

The number of irregular ticks in the SPX is more than 8 time of the number of irregular minutes within the VIX, and it seems that it is concentrated in two particular month – August 2012 and October 2012, and within these month it centres only in five days in August and 6 days in October. These phenomena might be explained by the very low volume (lowest in the past five years) the market has experienced in August 2012 and by the weak corporate results during the month of October 201214. It should be noted that all days with irregular ticks in August and October of 2012 had been complemented with tick in the following or previous minute to account for 12 or 4 ticks for the SPX and VIX, respectively.

Another problem that these data might encounter is significant outliers. Consistent with the data cleansing procedure performed by Oxford-Man Institute of Quantitative Finance (Realized Library), we consider an outlier with respect to the median absolute minute change for the day

---

12 See http://www.cboe.com/micro/vix/vixwhite.pdf for the methodology of the VIX index and https://www.sp-indexdata.com/idpfiles/indexalert/prc/active/whitepapers/Methodology_SP_US_Indices_Web.pdf for the S&P500 Index methodology

13 More information can b provided form the author

14 The market performance might be an explanation to the unusual tick irregularities, but that not to dismiss the possibility that it has also could have happened by chance with no specific explanation behind it.

(as our preliminary examination) and then with respect to the median absolute minute returns for the day.

Starting with analysing absolute minute changes, we first look at absolute changes that are at or higher than 50 times the median of the absolute change of that day. In this case, the SPX had no outliers and the VIX only had four. That seems to not be the correct measure for outliers for our sample and after further investigation 25 times the median of the absolute change on each day appeared to be as a better measurement. In this case we found 7 such events for the SPX and 91 events for the VIX.

The analysis for outliers using absolute returns (those which are 25 times the median absolute return of each day) reveals some interesting features of the data sample. The SPX has seven such outliers concentrated within 4 trading days and the VIX has 91 outliers which are concentrated within 50 trading days. For both the SPX and the VIX, the minimum outliers is about 25 times the median of the absolute return of the day and the maximum outlier is about 85 times and 73 times the median absolute daily return for the SPX and the VIX, respectively, whereas the median outlier is about 27 and 30 times the daily median absolute return for the SPX and the VIX, respectively. As is expected in most intraday data samples we observe that for the VIX 34% of the outliers appear during the first half hour of the trading day and about 24% of the outliers appear during the last half hour of the trading day, which sum-up to about 58% of all outliers in the data sample. The rest are spread out somewhat evenly during the trading day, except for the second half hour of the trading day (i.e., 10am to 10:30am) which correspond to about 10% of all outliers in the sample (see table 5).

Table 5: Outliers per 30 Minute Interval

Sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day.

Interval		VIX		SPX	
From	To	Number of outliers	Percentage	Number of outliers	Percentage
9:30a	10:00a	31	34.07%	1	14.29%
10:01a	10:30a	9	9.89%		
10:31a	11:00a	1	1.10%		
11:01a	11:30a	2	2.20%		
11:31a	12:00p	1	1.10%		
12:01p	12:30p	1	1.10%		
12:31p	13:00p	2	2.20%	1	14.29%

Interval		VIX		SPX	
From	To	Number of outliers	Percentage	Number of outliers	Percentage
13:01p	13:30p	5	5.49%	4	57.14%
13:31p	14:00p	4	4.40%	1	14.29%
14:01p	14:30p	3	3.30%		
14:31p	15:00p	5	5.49%		
15:01p	15:30p	5	5.49%		
15:31pp	16:01p	22	24.18%		
Total		91	100%	7	100%

Looking at the correlations of the number and size of the VIX outliers with the SPX absolute returns and their sign suggests interesting characteristics of the way VIX behaves in relation to movements in the SPX (see table 6). In this type of analysis we are more concerned with the direction of the relationship rather than its magnitude (as it will be difficult to obtain a significant magnitude considering the small sample of outliers). The number of VIX outliers per day has a negative correlation with both the SPX total return per day and the sign of its return. Both of these observations imply that we should expect more outliers (i.e., irregularities) in the VIX when the SPX moves down (i.e., negative returns). The analysis of the size of the VIX outlier with the SPX returns per day is less telling. With respect to the absolute SPX return it seems that the relationship is positive which would indicate that the higher the return (positive or negative) the larger is the VIX outlier. If instead of looking at the total return per day we consider the contemporaneous minute return or the preceding minute return, we then find that the size is negatively related to the sign of the previous minute return, which implies that if the SPX moved down in the previous minute it is likely that we would observe a large adjustment (i.e., change) in the VIX value. These observations insinuate that our Hypothesis 2 may have merit15.

One of the problems one encounters with financial data is “missing data”. In the case on financial markets that could often happen due to technological system glitches16. The analysis shows that the SPX sample has no missing data, whereas the VIX data encountered a few days with missing data and one day in particular with significant gap in the data.

15 Further discussion on the analysis of the examination of the hypothesis is section 2.2

16 For the past year we have experienced several technological systems mishaps – knight capital August 2012, Batched Facebook IPO April 2012, CBOE, April 2013, Nasdaq-NYSE AUGUST 2013, NYSE September 2013. These are only a few that were documented and reported to the news. In reality, however, one can observe glitches and technology mishaps almost every day, on a security basis. (the ones that are usually reported are the ones that have an impact on significant part of the financial market (and not simply an individual stock) and which last for more than a few minutes.

Table 7 summaries the statistics with regards to the VIX missing data. There are nine trading days with missing data (account for 3.1% out of total trading days in the sample) and a total of 126 missing minutes (account for 0.1% for total trading minutes in the sample considering only trading hours 9:30am to 4:01pm, and 286 whole trading days and three half trading days – 112,748 minutes). Four days have only one missing minute and one day with two missing minutes, for these minutes we have interpolated the data to fill in for the missing minutes. Two days have three missing minutes and one day with ten missing minutes. For these days we have deleted these missing observations17.

April 25, 2013, however, was an exceptional day. The CBOE experienced an outage that day since the opening of the trading day and resumed trading only at 1pm18. It is not unusual to see a trading delay in one of the 11 exchanges on which options are traded. This type of delay happens once a month. It is generally not too disruptive since banks can just reroute orders from one exchange to another. The S&P 500 options and the options on the CBOE Volatility Index (VIX), exclusively trade on the CBOE so there was no trading in those contracts while the CBOE was shut19. On April 25th the CBOE had an internal system issue caused by software problem and "not the result of any outside influence" or cyber-attack. Trading resumed in the S&P500 options contracts at 12:50 pm and in all other equity and ETF options opened by 1pm. Most of the trading functions were operating normally once they reopened, but some electronic methodology of confirming open outcry trades were being entered manually.

Table 7 shows that the VIX on April 25, 2013 had missing data from 11:06am to 12:49pm (which account for 104 minutes). It is understandable that the VIX resumes activity around 12:50pm, as the S&P and VIX options resume trading at that time. It is less clear why would Bloomberg show activity on the VIX form 9:30am to 11:05am, while the CBOE was down. The VIX quotes are derived (among other parameters) from the quotes of the underlying, S&P500 and from the prices of the options on the S&P500, which only trade on the CBOE and on April 25th, 2013 did not start trading before 12:50pm. Hence, it is unclear how those VIX quotes were calculated and whether they are reliable. For this reason we have decided to treat this outage on CBOE in the following way: (1) Include April 25th, 2013 in the sample but only for the trading

17 There was no particular information relating for these minutes in terms of particular glitches in the system. Since we do not know the reason for its absence and a simple interpolation would have distorted the sequence of the data, we have decided to delete these observations. Even though there was no stated reason for the absence of these minutes, it does not mean that it could not have been a technical mishap of the system. Nonetheless, their deletion will have a miniscule effect on the total data and its analysis as a whole, on a daily basis (2% out of daily minutes) and most definitely on the whole sample period (0.005% of total trading minutes).

18 The outage was the latest in a series of disruptions at exchanges, including Nasdaq's high-profile flub with Facebook IPO in April 2012. It also comes at a time when the financial services industry has good reason for concern about network security due to hacks. Major banks have suffered numerous denial of service attacks on their website in recent months, and Charles Schwab was attacked just earlier that week. The CBOE stressed that it had not been hacked.

19 It should be noted that the CBOE is the only exchange that rarely has any problems.

period 12:50pm to 4:01pm; (2) delete April 25th, 2013 from the sample data; (3) do not perform any analysis on a daily basis20.

To estimate the VAR model (as described below) we needed to calculate the minute returns21.We observe that the returns for the VIX and for the SPX have some minutes with zero returns: 1,956 minutes for the SPX and 49,439 minutes for the VIX, out of total 112,243 for the entire sample22, which account for 1.7% of total observations for the SPX and 44% of total observations for the VIX. This phenomenon is prevalent every trading day for the VIX and in 286 days (out of the total 289 days in the sample) for the SPX. The effect of zero returns in the SPX sample data is quite negligible (see table 8). Considering that its number of zero returns within a trading day varies as minimum as 1 minute and as maximum of 16 minutes, 0.26% and 4% of total minutes within the day and with a median of 7 minutes (i.e., 1.79%); and their appearance during the day is quite sporadic. Unlike the SPX the VIX experience quite a significant amount of zero returns during a trading day, this begs for a more thorough analysis as for the patterns of these zero returns and their relationship to movement of the SPX.

---

20 We have decided to exclude this day from any daily analysis,as this is an unusual day and any results would be biased and specific to this particular event. Options (1) and (2) were performed for robustness test.

21 The minute returns are calculated as $r_t = \ln\left({}^{P_t}/_{P_{t-1}}\right)$ where t is minute.

22 After accounting for missing data and in the sample (see discussion on table 7)

Table 6: Correlation of outliers per day and the SPX total returns per day

Sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day.

Correlation	SPX total return per day	SPX absolute total return per day	Sign of SPX total return per day	Sign of SPX return for the minute	SPX return for the previous minute	Sign of SPX return for the previous minute
Number of outliers per day – SPX	0.2647	0.1871	0.1956			
Number of outliers per day – VIX	-0.104	0.1726	-0.1271			
Size of VIX outliers (1)	0.2005	0.2324	0.1952	0.0711	0.118	-0.0078
Size of SPX outliers (2)	0.2337	0.2337	**			

Notes:

Outliers are defined as absolute returns minute by minute that are above 25 times of the median of absolute returns minute by minute for each day.

Average size of outliers in terms of the median of absolute returns for each day

All 4 days with outlier in SPX has positive total return in SPX.

ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

Table 7: VIX Missing Data Statistics (1),(2)

Date	Minutes Missing	Gap 1	Gap 2	Obs	Mean	Standard Dev.	Min	Median	Max	Skewness	Kurtosis
8/20/12	1	3:14 PM		390	14.32	0.1971	14.02	14.27	14.78	0.74	2.55
8/21/12	10	11:27 AM - 11:35 AM	11:38 AM	381	14.60	0.4130	14.04	14.58	15.44	0.35	1.94
8/29/12	2	9:58 AM	10:01 AM	390	16.68	0.1312	16.50	16.64	17.00	0.87	2.60
9/14/12	3	12:24 PM - 12:26 PM		388	14.36	0.3164	13.53	14.52	14.70	-1.32	3.57
1/11/13	1	4:00 PM		390	13.55	0.0882	13.37	13.53	13.78	1.00	4.06
1/25/13	1	2:49 PM		390	12.81	0.1003	12.50	12.82	12.99	-1.06	4.20
4/25/13	104	11:06 AM - 12:49 PM		287	13.42	0.2225	13.13	13.33	13.87	0.50	1.86
7/12/13	1	2:37 PM		390	13.93	0.0536	13.74	13.93	14.03	-0.82	4.24
9/16/13	3	1:40 PM - 1:42 PM		388	14.16	0.1577	13.87	14.15	14.47	0.10	2.03

Notes:

Calculated from Historical Intraday Minute Bar data – see Table 1 Row 2 for details

Sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day.

Table 8: Zero Return Descriptive Statistics

This table lists descriptive statistics of the number of zero returns from minute to minute during normal trading hours. In the sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day. (289 days with 0-returns in VIX, 286 days with 0-returns in SPX.)

	Mean	Std. Dev.	Smallest	1% percentile	5% percentile	10% percentile	50% percentile	90% percentile	95% percentile	99% percentile	Largest
# 0-return VIX per day	171.21	51.06	26	34	83	106	174	239	248	267	277
# 0-return SPX per day	6.86	3.13	1	1	2	3	7	11	13	15	16
% 0-return VIX per day	44.17%	12.97%	6.67%	8.72%	22.05%	28.13%	44.87%	61.28%	63.59%	68.46%	71.03%
% 0-return SPX per day	1.77%	0.82%	0.26%	0.26%	0.51%	0.77%	1.79%	2.81%	3.32%	4.09%	4.27%

Table 9: Correlations of Reruns

This table provides information on the correlation of minute zero returns with in the day and the total return for the day (sign of the total return and the absolute value). In the sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day. (289 days with 0-returns in VIX, 286 days with 0-returns in SPX.)

	# 0-return VIX per day	# 0-return SPX per day	% 0-return VIX per day	% 0-return SPX per day	Total return SPX per day	Total return VIX per day	Absolute total return SPX per day	Absolute total return VIX per day	Sign of total return SPX per day	Sign of total return VIX per day
# 0-return VIX per day	1									
# 0-return SPX per day	0.4083	1								
% 0-return VIX per day	0.9842	0.4118	1							
% 0-return SPX per day	0.377	0.9833	0.4091	1						
Total return SPX per day	0.2825	0.1429	0.2979	0.1491	1					
Total return VIX per day	-0.1919	-0.0931	-0.1948	-0.0907	-0.7688	1				
Absolute total return SPX per day	-0.294	-0.259	-0.2974	-0.2557	-0.1326	0.196	1			
Absolute total return VIX	-0.2123	-0.1721	-0.2258	-0.18	-0.2065	0.2215	0.6633	1		

	# 0-return VIX per day	# 0-return SPX per day	% 0-return VIX per day	% 0-return SPX per day	Total return SPX per day	Total return VIX per day	Absolute total return SPX per day	Absolute total return VIX per day	Sign of total return SPX per day	Sign of total return VIX per day
per day										
Sign of total return SPX per day	0.3042	0.1682	0.3135	0.1619	0.7455	-0.552	-0.0876	-0.1111	1	
Sign of total return VIX per day	-0.1865	-0.1105	-0.172	-0.0923	-0.5286	0.6995	0.1145	0.0882	-0.5659	1

ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

Table 8 shows that the VIX experiences a much more frequent appearance of zero returns during the trading day than the SPX does. Its number of zero returns within a trading day varies between a minimum of 26 minutes and a maximum of 277 minutes, 6.67% and 71.03% of total daily trading minutes and with a median of 174 minutes (i.e., 44.87%). We then investigate how this vast number of zero returns for the VIX relates to the SPX movement. Using correlation (see table 9) we again are interested more in the directions rather than the magnitude of the measure23. The analysis reveals that when correlating either the daily number of zeros returns for the VIX or the percentage of minutes (with zero returns out of daily 391 minutes) with the absolute SPX total returns per day, the direction is negative, implying that the lower is the change in SPX the more zero minute returns we'll observe in the VIX – therefore, one should expect more activity (adjustments) in the VIX when the SPX has significant movement and vice-versa when the SPX barely changes then there are less adjustment in the VIX and thus we'll observe more zero minute returns. When correlating either the number of zero minute returns or the percentage of the daily zero minute returns (out of 391 daily minutes) with the sign of the total SPX return per day, we observe positive correlation, which indicates that if the sign is positive (i.e., the SPX moves upwards) one should expect more zeros in the VIX (i.e., less activity), and the opposite when the SPX moves downwards – when the SPX is trending down the VIX will adjust more frequently. These two observations are consistent with the documented asymmetry in the equity markets, sometimes ascribed to as "leverage effect" or the "risk premium" effect – in the equity market it is unlikely that positive and negative shocks have the same impact on the volatility. In the story of news, negative news reduces the demand for the stocks because of risk aversion. The consequence decline in stock value is followed by the increased volatility as forecasted by the news.

Research Method and Hypothesis

Research Method

We are trying to build an empirical model that best explains and predicts the changes in the S&P 500 (in other words – a model which explains (and predicts) the returns of the S&P500).

We can think about the S&P 500 returns as following a Stochastic Differential Equation (SDE) of the general form:

$$\frac{dS}{S} = \text{Drift\_term} + \text{Diffusion\_term} \qquad \text{(SDE)}$$

23 The existence of zero returns in an interaday sample is a consequence of many different variants (such as volume, time of the day, news, etc.) which will not be captured by a single measurement such as correlation and hence we are only concerned on the direction and not on the magnitude.

Instead of assuming a specific SDE (such as the Geometric Brownian Motion (GBM)), we can say that we do not know the exact theoretical model of the SDE that explains returns. Therefore, we are trying to build an empirical model that may serve as a good proxy for the SDE. If we do not have any specific theory in mind then the best econometric tool for that purpose would be estimation via Vector Autoregression (VAR) model (which is a purely atheoretical estimation method).

The form of the VAR model is:

$$\mathrm{Y_t = C + A_1 Y_{t-1} + \cdots + A_p Y_{t-p} + \epsilon_t} \tag{1}$$

Where C is $\mathrm{k \times 1}$vector of constants (i.e., intercepts), $\mathrm{A_i}$ is $\mathrm{k \times k}$ matrix (for every $\mathrm{i = 1, \ldots . p}$) and ϵ_t is a $\mathrm{k \times 1}$ vector of error terms. We estimate a Bi-Variate VAR model of two time series: (i) returns of the S&P 500; and (ii) changes in the VIX.

This estimated Bi-Variate VAR model will result with two estimation equations; (i) estimated returns of the S&P 500, and (ii) estimated changes of the VIX. Each estimation model will be a function of lag-variables of S&P 500 returns and lag-variables of the changes of the VIX. More formally,

$$\begin{aligned} \mathrm{RetSPX_t} &= \mathrm{C_1 + f_1\big(RetSPX_{t-1,}, RetSPX_{t-2}, \ldots , RetSPX_{t-p,}\big)} \\ &\quad \mathrm{+g_1(ChangeVIX_{t-1}, ChangeVIX_{t-2}, \ldots , ChangeVIX_{t-p})} \end{aligned} \tag{2a}$$

And

$$\begin{aligned} \mathrm{ChangeVIX_t} &= \mathrm{C_2 + f_2\big(RetSPX_{t-1,}, RetSPX_{t-2}, \ldots , RetSPX_{t-p,}\big)} \\ &\quad \mathrm{+g_2(ChangeVIX_{t-1}, ChangeVIX_{t-2}, \ldots , ChangeVIX_{t-p})} \end{aligned} \tag{2b}$$

Where $\mathrm{C_1}$ and $\mathrm{C_2}$ are constants; and $\mathrm{f_1}$ and $\mathrm{f_2}$ are a function of the lags values of the S&P 500 returns; and $\mathrm{g_1}$ and $\mathrm{g_2}$ are a function of the lags values of the changes in the VIX; and p is the optimal number of lags to be determined via information criteria.

The way we can interpret equation (2a) is as follows: remember that we are trying to find a good empirical (estimation) model as a proxy for the SDE. Therefore, we can view the function $\mathrm{f_1}$ as a proxy for the drift and the function $\mathrm{g_1}$ as a proxy for the diffusion, which in this case both are in-line with the concept of drift and diffusion, respectively. That is $\mathrm{g_1}$as a proxy for the diffusion is in fact a function of the VIX (and the VIX is a measure of the expected volatility of the S&P 500 index in the next 30-day period), and hence could be a good representation for the randomness term of the SDE. The function $\mathrm{f_1}$, on the other hand, is a function of past returns , hence can be perceived as a good representation of the drift term of the SDE.

ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

Hypotheses

After describing the empirical models we now present a few main hypotheses.

Since we are using a VAR model we will check two main competing hypotheses (and one secondary) with respect to causality.

Hypothesis 1:

As explained above the VAR index translates, roughly, to the expected movement in the S&P 500 Index over the next 30-day. This implies that the VIX is a forward–looking measure of the S&P500 future volatility. Therefore, we would expect a leading relationship – meaning the VIX movement leads the S&P 500 – hence, we would say that VIX "granger causes" the S&P 500 Index.

Hypothesis 2:

The VIX Index is an "implied volatility" measure extracted from the options' prices. The option price is a function of the underlying, which in this case is the S&P 500 Index. This suggests that the VIX measure is a function of the S&P 500, and hence implicitly determined by the values of the S&P 500 Index. Therefore, this type of relationship implies that the S&P 500 "granger cause" the VIX.

Hypothesis 3 (secondary):

Both VIX and SPX minute return Granger Cause each other's time series, which implies that both arguments on the movement of the VIX could be correct. Although, even if we'll observe a bi-directional causality or a bi-directional feedback, we may still find that one causality is greater and more significant than the other. A bi-directional causality does not necessarily indicate that the impact is the same, it only indicates that both series affect each other (but not necessarily in the same magnitude).

We then can have a sub-hypothesis in hypothesis 3 – ***Hypothesis 3(a):***

SPX granger cause VIX and VIX granger cause SPX, however, the relation described in hypothesis 2 is stronger than the relationship described in hypothesis 1 and hence the magnitude and significance of the impact of the SPX minute returns on the VIX movement is greater than the impact of the VIX minute returns on the SPX movement.

Empirical Analysis

Analysis of Daily Data

We start with looking at the behaviour of the daily data and investigate whether the intraday minute returns are stationary on a per day analysis – checking for special patterns or irregularities that might shed light on the interpretation of our results when we then run the analysis on the full sample period August 9^{th} 2012 to October 3^{rd} 2013.

Augmented Dickey Fuller (ADF) test for staionarity was performed per day on the minute returns for the SPX and VIX, for each of the 289 trading days in the sample. The SPX return minutes are stationary (at 0 and 1 lag) for each day under both the Schwartz Information Criterion (SIC) also known as the Bayes Information Criterion and the Akaike Information Criterion (AIC)24. This is to be expected, considering that the analysis in the previous section has alluded to minor (to no) significant irregularities of the SPX minute return time series25.

The VIX minute return time series, however, has experienced irregularities, such as missing data and high frequency of outliers and zero returns. Hence, one would expect some level of non-stationarity in the VIX data. When we use the SIC test it appears that the VIX minute-return time series is stationary for each of the 289 trading days in our sample period, but with different lags26. The minimum lag is 1 and the maximum lag is 12. The lag with the highest frequency (i.e., the most trading days that are stationary with this lag) is one, 104 out of 289 (i.e., 36%); the least frequent is 12 lag with only one day (i.e., 0.35%); we have 44 days (i.e., 15%) with two lags, 42 days (i.e., 14.5%) with three lags, 37 days (i.e., 12.8%) with four lags, 19 days (i.e., 6.6%) with five lags, 18 days (i.e., 6.23%) with six lags, and 10 days (i.e., 3.5%) with seven lags – one to seven lags covers about 95% of all trading days in the sample period).

Using the AIC test we observe different results with relation to the stationarity of the VIX. Even when we use a max-lag of 30 we still do not achieve stationarity in all trading days, under the AIC test – 118 days (40.83%), 182 (63%), 226 (78.2%) and 252 (87.2%) have stationary minute-return time series at a p-value of 1%, 5%, 10% and 20% respectively.

When we make choices about the order p of an autoregression, we have to balance the marginal benefit of including more lags against the marginal cost of increased uncertainty of estimation. If we do not include enough lags, we run the risk of omitting potentially significant information contained in more distant lags. On the other hand, with too many lags we estimate more coefficients than needed. This practice may lead to estimation error. In most cases we prefer the model that has fewer parameters to estimate, provided that each one of the candidate models is correctly specified. This is called the most "parsimonious" model. The AIC does not always suggest the most parsimonious model because the AIC function is largely based on the log likelihood function. In the BIC model as $T \to \infty$, the addition of another lag would increase the BIC value by a larger ln(T) margin. Hence, asymptotically, BIC would pick the more parsimonious model than AIC might suggest. Stock and Watson (2007) recommends that we choose the model suggested by AIC rather than BIC. They argue that including more parameters is better than omitting significant parameters.

The AIC procedure, however, has been criticized because it is inconsistent and tends to over fit models. A criterion is said to be order consistent if, as sample size increases, the criterion is minimized at the true order with a probability which approaches unity. Geweke and Meese showed that for regression models, Shaibata (1976) for autoregression models and Hannan

24 At a p-value of 10%, 5% and 1% for both SIC and AIC test.

25 Test results are available from the author

26 All staionarity results are at the 1%, where the p-value for each day is very close to zero.

(1982) for ARMA model. Romero (2007) affirmed that SIC deals with the problem of inconsistency noticed in AIC. Asymptotically, the SIC is minimized at the model order having the highest posterior probability. Akaike (1977) had shown that SIC can be more successful than the AIC in estimating the degree of a polynomial regression model and in estimating the order of an autoregressive model. Gayawan and Ipinyomi (2009) observe that AIC tend to choose more complex model than the SIC and agues in favour of using SIC when applying the principal of simplicity. This supports the results of Sheek (1984) and Anne and Murphee (1988).

Since we have observed extreme difference in the strationarity test of the VIX when using SIC test as opposed to the AIC, we need to decide which test would be the most appropriate one to apply in our case. The following analysis was done in order to understand whether AIC rejection of stationarity is valid (i.e., the test is able to capture something in the data due to the VIX irregularities).

The analysis is four-fold27. First we investigate whether the days that AIC consider as non-stationary time series is due to the number of zero returns within the day. A comparison of the sample of stationary versus non-stationary days reveals that the distributions of the percentage of zero minute return in a day is quite similar and hence suggests that the non-stationarity indicated by the AIC test is that the results of the frequency of zero minute returns. Second, we examine whether it's the result of the magnitude of the absolute total daily return of the SPX. Again, the distributions of days which are stationary versus the distributions of the days that are non-stationary are quite similar and do not suggest that the AIC results is driven by differences in SPX total return for the day (which is a proxy for increased momentum in the day). Third, we look at the number of ticks in the minute and the analysis shows that the stationary days and the non-stationary days, behave rather similar. Lastly, we examine the number of outliers in a day for the two sub-samples stationary versus non-stationary and ascertain that the two samples have a similar distribution of the number of outliers per day. From the above analyses we conclude that the AIC is most likely not capturing some significant irregularities in the data but simply over fits the intraday minute data per day. Therefore, we decided to apply the SIC test to our sample data as it is most parsimonious, order consistent and does not run into the problem of over fitting28. For the rest of the analysis in this paper we will use SIC test (rather than AIC).

Analysis of time-series minute returns for the sample period29:

Vector Autoregressive Analysis:

Four samples of minute returns have been created for the time period August 9th 2012 to October 3rd 2013 – (1) excluding outliers, but including half trading day of April 25th 201330 (sample A);

27 Test results and analysis are available from the author.

28 Especially, if we would like to make predictions (as in the case of the VAR model) we should refrain from over fitting the model.

29 Sample period is referred to August 9th 2012 to October 3rd 2013

(2) excluding both outliers and April 25th 2013 (sample B); (2) including outliers and half day of April 25th 2013 (sample C); (4) including outliers but excluding April 25th 2013 (sample D).

We start our analysis with sample A and use samples B to D for robustness tests31. The first essential step in estimating with VAR model is to check for the order of integration (i.e., strationarity test). The Augmented Dickey-Fuller (ADF) test32 confirms our assumption that the two time series are indeed stationary when using either SIC or the AIC test33. For the SPX returns sample the ADF test indicates statonarity with zero lag (according to SIC) and with 8 lags (according to AIC). For the VIX returns sample the ADF indicates stationarity with 8 lags (according to SIC) and with 52 lags (according to AIC). Consistent with the discussion above as to which criterion to use, we can clearly see that the AIC incorporates a significant amount of lags (especially in the case of the VIX) and hence, is very likely to be over fitting, and thus the main reason we have chosen to follow the results of the SIC.

Since all variable are $I(0)$ (stationary), we can use the standard case of VAR model. Before we continue with the estimation we will determine the optimal order of lags. In order to determine the optimal order of lags we will use Information Criteria, which are a measure of the relative goodness of fit of a statistical model. In order to perform the Information Criteria test we need also to determine the starting number of maximum lags. Usually the maximum lag number is set to $T^{1/3}$, where T is the number of observations in the time series. Using this measure we find that the starting number of maximum lag, for sample A is about 48. Applying this max-lag value we obtain that the optimal lags are nine lags and 41 lags according to SIC and AIC, respectively. Using the nine optimal lags indicated by SIC we then estimate the VAR model (see Table 10)34.

---

30 This is the day when the CBOE suffered a major shutdown and started operating regularly only around 1pm. The half trading day is the hours that the CBOE was operating regularly (12:51pm to 4:01pm).

31 As the following analysis shows we end up choosing sample C as our main sample for further analysis and applications

32 The results are available from the author.

33 We use the most general form of ADF test with both intercept and time trend.

34 We also performed stability check on the estimated VAR model, and concluded that the model is indeed stable.

Table 10: Estimated VAR Model (Sample A)

Sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day.

	RETURN_SPX	RETURN_VIX
RETURN_SPX(-1)	0.085650	-2.158301
	(0.00361)	(0.02013)
	[23.7520]	[-107.244]
RETURN_SPX(-2)	0.007281	-0.612187
	(0.00383)	(0.02136)
	[1.90234]	[-28.6601]
RETURN_SPX(-3)	0.008181	-0.344325
	(0.00386)	(0.02157)
	[2.11678]	[-15.9631]
RETURN_SPX(-4)	-0.008593	-0.166745
	(0.00387)	(0.02157)
	[-2.22296]	[-7.72942]
RETURN_SPX(-5)	0.002857	-0.067771
	(0.00385)	(0.02150)
	[0.74170]	[-3.15213]
RETURN_SPX(-6)	-0.004951	-0.034357
	(0.00383)	(0.02136)
	[-1.29335]	[-1.60820]
RETURN_SPX(-7)	-0.003765	0.005819
	(0.00381)	(0.02125)
	[-0.98881]	[0.27386]
RETURN_SPX(-8)	-0.001299	0.024810
	(0.00377)	(0.02104)
	[-0.34451]	[1.17905]

	RETURN_SPX	RETURN_VIX
RETURN_SPX(-9)	-0.010901	0.032579
	(0.00370)	(0.02067)
	[-2.94340]	[1.57620]
	(0.00065)	(0.00361)
	[7.97131]	[-33.1332]
RETURN_VIX(-2)	0.000299	0.018537
	(0.00066)	(0.00367)
	[0.45423]	[5.05007]
RETURN_VIX(-3)	-0.00111	0.033290
	(0.00066)	(0.00370)
	[-1.67199]	[8.98597]
RETURN_VIX(-4)	0.000302	0.047311

ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

Table 10: Estimated VAR Model (Sample A) cont.

Sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day.

	RETURN_SPX	RETURN_VIX
	(0.00066)	(0.00369)
	[0.45689]	[12.8223]
RETURN_VIX(-5)	-0.001053	0.042412
	(0.00066)	(0.00366)
	[-1.60338]	[11.5732]
RETURN_VIX(-6)	-0.000527	0.036652
	(0.00065)	(0.00364)
	[-0.80851]	[10.0828]
RETURN_VIX(-7)	0.000615	0.027770
	(0.00064)	(0.00359)
	[0.95442]	[7.72781]
RETURN_VIX(-8)	0.000686	0.030144
	(0.00064)	(0.00355)
	[1.07942]	[8.49738]
RETURN_VIX(-9)	2.04E-05	0.023805
	(0.00060)	(0.00336)
	[0.03387]	[7.07594]
C	3.44E-07	-3.90E-06
	(7.5E-07)	(4.2E-06)
	[0.45815]	[-0.92828]
R-squared	0.005956	0.129522
Adj. R-squared	0.005793	0.129378
Sum sq. resids	0.006761	0.210585
S.E. equation	0.000249	0.001388

	RETURN_SPX	RETURN_VIX
F-statistic	36.40062	903.9218
Log likelihood	752524.4	564479.3
Akaike AIC	-13.76085	-10.32213
Schwarz SC	-13.75919	-10.32046
Mean dependent	3.40E-07	-4.43E-06
S.D. dependent	0.000249	0.001487
Determinant resid covariance (dof adj.)		8.34E-14
Determinant resid covariance		8.33E-14
Log likelihood		1336499.
Akaike information criterion		-24.43949
Schwarz criterion		-24.43616

The results show an interesting pattern of the relationship of each variable to its own lags. The SPX minute returns seems to strongly relate to its own first lag (at the 1% significance level), and relates to lags 2 to lag 4 (at the 5% significance level), and then to lag 9 (at the 1% significance level), with positive relation to lags one, two and three and negative relation to lags four and nine. The magnitude of lag one with respect to the second, third and fourth lag is about tenfold and with respect to the ninth lag its eight fold. The VIX returns, on the other hand, are related to all of its nine lags (at the 1% of significance level), with negative relation to its first lag and positive to all other lags – this can be interpreted as a correction followed by a momentum, whereas the magnitude of the correction is of six fold compare to the second lag and about three to four fold compared to lags three to nine. Hence, one can explain the VIX minute changes (returns) that on average there is an immediate correction (or reversal) and then some element of momentum (which will depend on the magnitude of the change in the VIX for that previous minute).

We also observe that the time series (VIX returns and SPX returns) affect each other. The SPX minute returns are only related to the VIX first lag (at a 1% significant level) with a positive but very small relationship (about 0.5%). Hence, even if a relation exists it may not be noticeable. The minute returns of the VIX, however, seem to experience a longer and a more significant relation with the SPX minute returns lags. The VIX minute returns are negatively related to the first six lags of the SPX minute returns (at a 1% significant level for the first five lags and at a 10% for the six lag). For the ninth lag the relation is positive (at a 10% significant level). From the above VAR model analysis we infer the following main points: (1) the first lag in both the VIX and SPX is significantly related to both time series; and it is always negative for the VIX and always positive for the SPX – i.e., reversal versus momentum for the VIX and SPX, respectively. (2) SPX minute returns time series does affect the VIX time series (up until 6 lags and with significant magnitude) but the opposite relationship (i.e., VIX affecting SPX), if exists, does not seem to have much of a significant impact. The following analysis shed more light on the later point.

First we perform the Granger Causality test. The test results, indicate that both time series the VIX minute returns and the SPX minute returns Granger Cause each other's time series, although the F-statistics for testing the Causality of the SPX time series on the VIX time series is significantly higher than the F-statistics of the opposite causality relationship, implying (and also consistent with the VAR estimated model results discussed above) that the SPX causality on the VIX time series is of a greater impact and magnitude. To support this observation we also conducted the Impulse Response analysis (see Figure 1), and the Variance Decomposition (see Figure 2).

Figure 1: Impulse Response (Sample A)

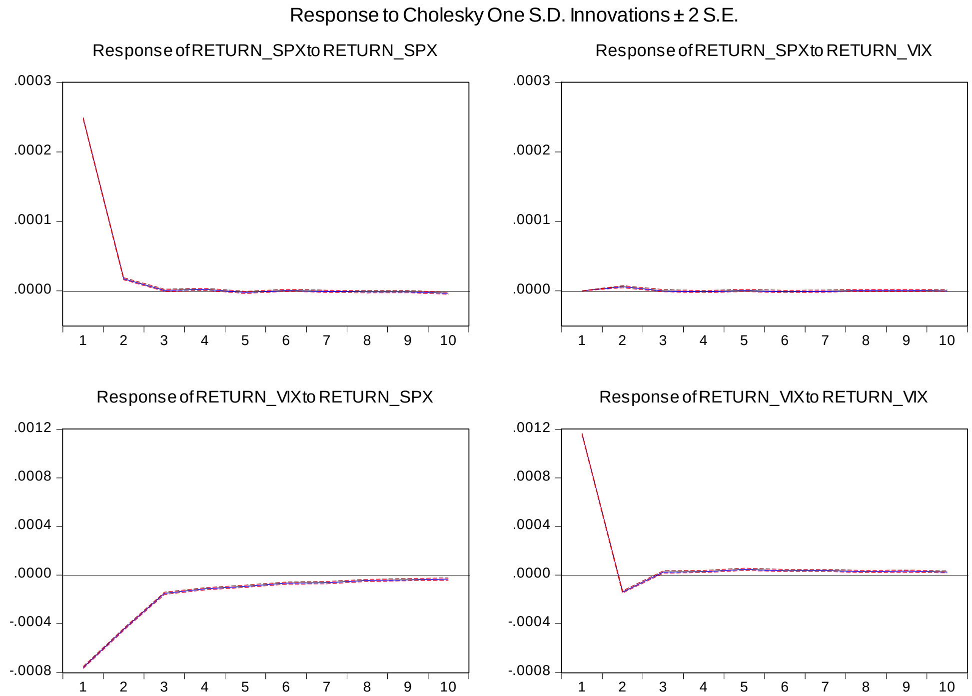

Figure 2: Variance Decomposition (Sample A)

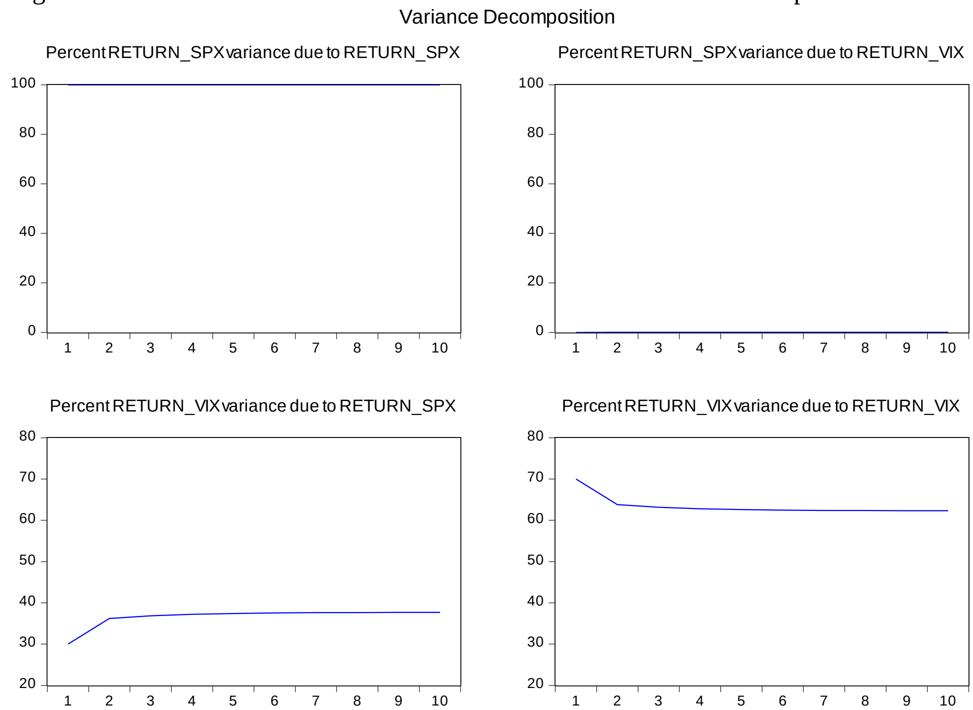

ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

It is clear from both the Impulse Response and the Variance Decomposition that the VIX time series most likely has no impact on the SPX time series but the opposite is not true. The Impulse Response analysis shows that a shock in SPX has a significant negative impact on the VIX, especially in the first three periods after the shock and then gradually dies out. A shock of the VIX, however, seems to have no effect on the SPX. When looking at the Variance Decomposition, it turns out that the SPX variance is 100% explained by its own variance, whereas the VIX variance is about 40% explained by the SPX variance, indicating that VIX movement can be explained (to some degree) by changes in the SPX, while the SPX movement cannot be explained by any such changes in the VIX35.

Robustness analysis: Analysis of sample C

For robustness check we examined the estimation of the VAR model with sample C (which is the same as sample A but does include the outliers)36. The results of sample C analysis are quite similar to those described above with respect to ample A. Starting with stationarity and ADF test 37, we again observe that both SPX and VIX minute returns time series are stationary – using SIC, SPX is stationary with zero lags and the VIX is stationary with 10 lags (which is similar to the results for sample A); using AIC we observe that both VIX and SPX are stationary with 32 lags. Although we decided to use a more parsimonious model (suggested by SIC), the result of AIC test where SPX and VIX are analysed with the same number of lags, indicates that it is very likely that sample C might be a better representation for the data and therefore a better sample for data analysis and conclusions.

Checking for optimal lag and applying 48 max-lag value we obtain that the optimal lags are eleven lags and 20 lags according to SIC and AIC38, respectively. Using the eleven optimal lags indicated by SIC we then estimate the VAR model (see Table 11)39.

35 It should be noted that when changing the ordering of the time series we still obtained the same results, which indicates that the relationship/impact is not driven by the ordering of the time series but rather represent a “real” impact of the shock of one time series on the other.

36 We wanted to establish that: (1) the results are not driven by the exclusion or inclusion of outliers, and; (2) if sample C and sample A behave quite similar then the best action would be to view sample C as the primer sample for analysis and further discussion.

37 Only shows results for SIC. The results for AIC are available from the author

38 It seems that when adding the outliers the number of lags following the AIC is reduced significantly (from 41 to 20). The SIC optimal lag does not change much it increases by two lags from 9 lags to 11 lags. This observation is another statistic that supports our inclination to use sample C as the primer data sample for analysis, it also may indicate that what we thought to be as an outlier is not really an outlier and therefore should not be excluded or ignored.

39 We also performed stability check on the estimated VAR model, and concluded that the model is indeed stable.

Table 11: Estimated VAR Model (Sample C)

Sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day

	RETURN_SPX	RETURN_VIX
RETURN_SPX(-1)	0.089364	-2.321414
	(0.00359)	(0.02104)
	[24.8643]	[-110.310]
RETURN_SPX(-2)	0.007355	-0.64267
	(0.00383)	(0.02241)
	[1.92180]	[-28.6783]
RETURN_SPX(-3)	0.008697	-0.393869
	(0.00386)	(0.02263)
	[2.25034]	[-17.4046]
RETURN_SPX(-4)	-0.007069	-0.24908
	(0.00387)	(0.02267)

ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

Table 11: Estimated VAR Model (Sample C)

Sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day

	RETURN_SPX	RETURN_VIX
RETURN_SPX(-5)	0.006065	-0.184325
	(0.00387)	(0.02264)
	[1.56889]	[-8.14256]
RETURN_SPX(-6)	-0.002358	-0.087982
	(0.00386)	(0.02259)
	[-0.61127]	[-3.89483]
RETURN_SPX(-7)	0.002156	-0.066434
	(0.00385)	(0.02253)
	[0.56040]	[-2.94925]
RETURN_SPX(-8)	0.001685	-0.006509
	(0.00382)	(0.02238)
	[0.44081]	[-0.29088]
RETURN_SPX(-9)	-0.005618	-0.008601
	(0.00380)	(0.02224)
	[-1.47882]	[-0.38667]
RETURN_SPX(-10)	0.004563	0.006028
	(0.00376)	(0.02204)
	[1.21219]	[0.27349]
RETURN_SPX(-11)	0.009480	0.040092
	(0.00370)	(0.02166)
	[2.56217]	[1.85065]
RETURN_VIX(-1)	0.005206	-0.130355
	(0.00061)	(0.00360)
	[8.47448]	[-36.2410]

	RETURN_SPX	RETURN_VIX
RETURN_VIX(-2)	0.001469	8.05E-05
	(0.00063)	(0.00367)
	[2.34353]	[0.02193]
RETURN_VIX(-3)	-0.000159	0.012646
	(0.00064)	(0.00373)
	[-0.24877]	[3.38937]
RETURN_VIX(-4)	0.000596	0.025327
	(0.00064)	(0.00372)
	[0.93631]	[6.80082]
RETURN_VIX(-5)	-0.001006	0.032128
	(0.00063)	(0.00371)
	[-1.58611]	[8.65114]

ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

Table 11: Estimated VAR Model (Sample C)

Sample period: Aug 9 2012 – Oct 3 2013, 9:30 AM – 4:01 PM for each trading day

	RETURN_SPX	RETURN_VIX
RETURN_VIX(-6)	-0.00045	0.029560
	(0.00063)	(0.00370)
	[-0.71244]	[7.99420]
RETURN_VIX(-7)	0.001117	0.022619
	(0.00063)	(0.00367)
	[1.78305]	[6.16653]
RETURN_VIX(-8)	0.001323	0.028401
	(0.00062)	(0.00364)
	[2.13033]	[7.81227]
RETURN_VIX(-9)	0.000461	0.021039
	(0.00061)	(0.00359)
	[0.75066]	[5.85244]
RETURN_VIX(-10)	0.000392	0.019292
	(0.00061)	(0.00354)
	[0.64752]	[5.44323]
RETURN_VIX(-11)	-0.000777	0.023585
	(0.00057)	(0.00334)
	[-1.36350]	[7.06500]
C	6.41E-07	-5.42E-06
	(7.6E-07)	(4.5E-06)
	[0.83983]	[-1.21235]
R-squared	0.006557	0.127710
Adj. R-squared	0.006356	0.127534

	RETURN_SPX	RETURN_VIX
Sum sq. resids	0.006912	0.236970
S.E. equation	0.000252	0.001475
F-statistic	32.66047	724.5267
Log likelihood	747817.8	555363.4
Akaike AIC	-13.73436	-10.19965
Schwarz SC	-13.73233	-10.19762
Mean dependent	6.62E-07	-6.53E-06
S.D. dependent	0.000253	0.001579
Determinant resid covariance (dof adj.)		9.80E-14
Determinant resid covariance		9.79E-14
Log likelihood		1321901.
Akaike information criterion		-24.27783
Schwarz criterion		-24.27377

The results show a similar pattern as we observed in sample A analysis, but with a few differences which, in fact, further support hypothesis 3(a) (i.e., SPX impact on VIX is more prominent than the impact of the VIX on the SPX) . Observing the relationship of each variable to its own lags: the SPX minute returns seems to strongly relate to its own first lag (at the 1% significance level), and relates to second and third lag (at the 5% significance level) and to the fourth lag (at the 7% significance level), and then to lag 11 (at the 1% significance level), with positive relation to lags one, two and three and negative to lags four and eleven. The magnitude of lag one with respect to the second, third and fourth lag is about tenfold and with respect to the eleventh lag its eight fold. The VIX returns on the other hand are related to nine of its lags (at the 1% of significance level) excluding the second and third lag (their values are also very insignificant). The first lag have a negative relation and lags forth to eleventh a positive relationship – again as we stated above when analysing sample A, one can interpret this pattern as a correction followed by a momentum, whereas the magnitude of the correction is, on average, about six fold compared to lag fourth to lag eleventh. Hence, one can explain the VIX minute changes (returns) that on average there is an immediate correction (or reversal) and then some element of momentum (which will depend on the magnitude of the change in the VIX for that previous minute).

The time series VIX minute returns and SPX minute returns influence each other. The SPX minute returns are related to the first and second lag (at a 1% significant level) with a positive but very small magnitude of impact (about 0.5%). Hence, even if a relation exists it may not be noticeable. The minute returns of the VIX, however, seem to experience a longer and a more significant relation with the SPX minute returns lags. The VIX minute returns are negatively related to the first to seven lag of the SPX minute returns (at a 1% significant level for all lags). For the eleventh lag the relation is positive (at a 7% significant level). From the above VAR model analysis we infer the following main points: (1) the first lag of the VIX and SPX is significantly related to both time series; and it is always negative for the VIX and always positive for the SPX – i.e., reversal versus momentum for the VIX and SPX, respectively. (2) the second lag of the VIX affects the SPX (but with a very insignificant magnitude) but does not have any impact on the VIX; also the third lag of the VIX has no impact on the current change of the VIX. These observations are different in what we have observed in sample A, which raises the question whether it has to do with the inclusion/exclusion of outliers. We will discuss this point in the following section (3) SPX minute returns time series does affect the VIX time series (up until 7 lags and with significant magnitude) but the opposite relationship (i.e., VIX affecting SPX), if exists, does not seem to have much of a significant impact. The following analysis shed more light on the later point.

First we perform the Granger Causality test. The test results, indicate that both time series the VIX minute returns and the SPX minute returns Granger Cause each other's time series, although the F-statistics for the testing the Causality of the SPX time series on the VIX time series is significantly higher than the F-statistics of the opposite causality relationship, implying (and also consistent with the VAR estimated model results discussed above) that the SPX causality on the VIX time series in of a greater impact and magnitude. To support this observation we also conducted the Impulse Response analysis (see Figure 3), and the Variance Decomposition (see Figure 4).

Figure 3: Impulse Response (Sample C)

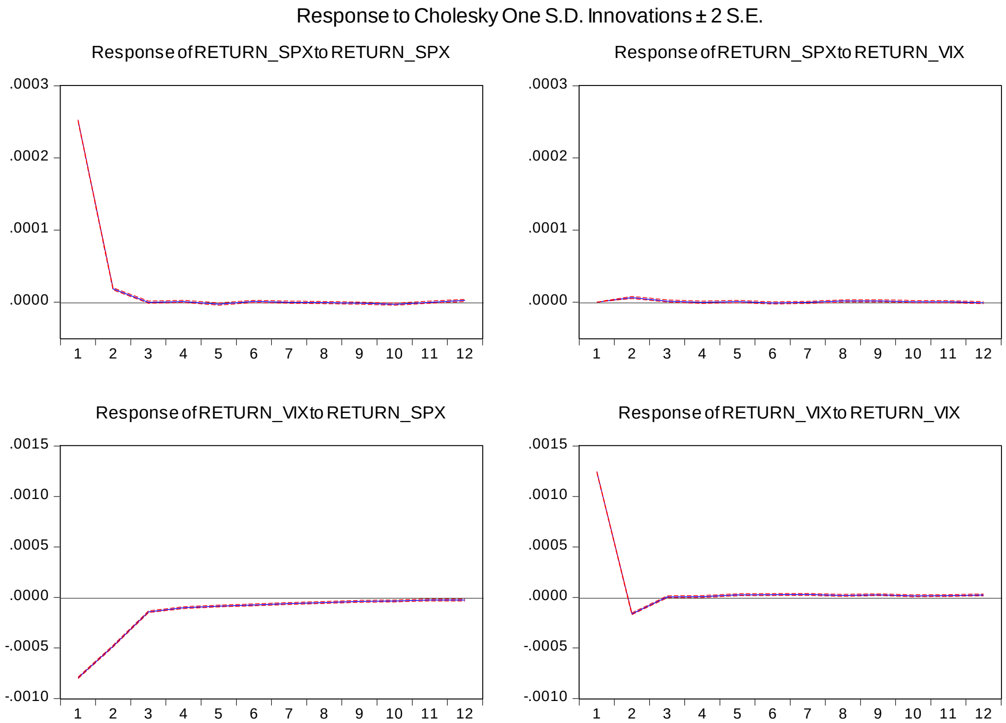

Figure 4: Variance Decomposition (Sample C)

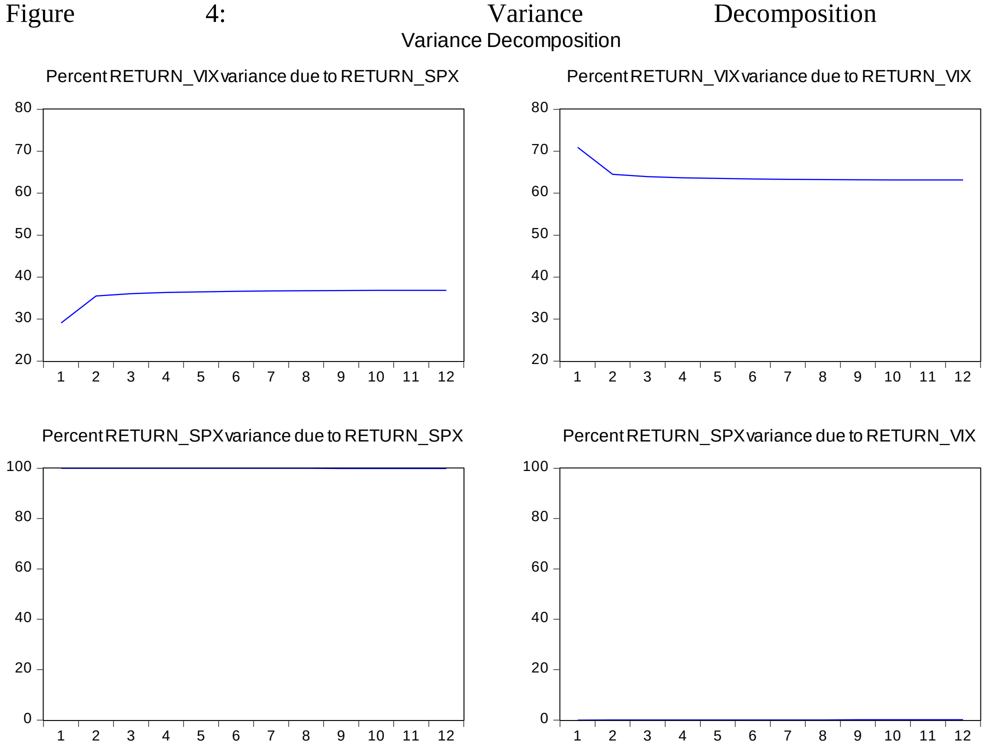

ACRN Journal of Finance and Risk Perspectives
Vol. 3, Issue 2, June 2014, p. 83 – 132
ISSN 2305-7394

It is clear from both the Impulse Response and the Variance Decomposition that the VIX time series most likely has no impact on the SPX time series but the opposite is not true. The Impulse Response analysis shows that a shock in SPX has a significant negative impact on the VIX, especially in the first three periods after the shock and then gradually dies out. A shock of the VIX, however, seems to have no effect on the SPX. When looking at the Variance Decomposition. It appears that the SPX variance is 100% explained by its own variance, whereas the VIX variance almost 40% is explained by the SPX variance, indicating that VIX movement can be explained (to some degree) by changes in the SPX, while the SPX movement cannot be explained by any such changes in the VIX40.

Further Robustness Analysis

We also ran the same tests described above on sample B and sample D41. The results for sample B are almost identical to the results of sample A; and the results of for sample D are almost identical to the results of sample C. Since the difference between samples A and B is the exclusion of April 25th 2013 in the data sample (sample B excludes this day), and the same is the difference between samples C and D (sample D excludes this day), it implies that this day has no impact on the results of the analysis, and therefore should be included.
When comparing the results of sample A to the results of sample C, it is evident that the fundamental results and conclusions are the same, which indicates that the primer result is not driven by any outliers, missing data or other changes and irregularities in the sample data – this is an important result of the analysis. We do observe, however, some minor difference between the two samples – (1) number of optimal lags is different: considering SIC the optimal lag is greater for sample C (eleven lags versus nine lags); considering AIC the optimal lags is significantly smaller for sample C (20 lags versus 41 lags). It also seems that for sample C the difference between SIC choice and AIC choice is not as wide as in the case of sample A. (2) for stationarity analysis under AIC the optimal number of lags for SPX and VIX is the same (32 lags) when using sample C. This is not evident in the case of sample A, as we obtain eight and 52 lags for SPX and VIX, respectively. (3) the VAR estimates when using sample C indicates that the VIX minute returns do not relate to their own second and third lag, while they do relate when using sample A. The first two points described above are in favour of using sample C instead of sample A, as it is indicated by the AIC measure. The third point indicates that the outliers may have an effect on the results, but it is also likely that it, in fact, captures the “true” structure/pattern of the time series. We believe that it would be best to draw conclusions from the analysis of sample C rather than sample A, and therefore use it for our further discussion and analysis.

40 It should be noted that when changing the ordering of the time series we still obtained the same results, which indicates that the relationship/impact is not driven by the ordering of the time series but rather represent a “real” impact of the shock of one time series on the other.

41 Tables are not included in the paper but are available from the author.

Further Analysis – Cointegration:

It is quite possible for random walks to be related to each other so that a regression of one random walk on the other has a stationary error term. For example let

$$\Delta x_t = \varepsilon$$

$$\Delta y_t = u$$

and let

$y_t + x_t$ be stationary. The simplest example is that $y_t = -x_t + v.$ That is, let one random walk be the negative of the other – allowing for some error. Then the sum is simply a random error with no unit root or autocorrelation.

If the combination of unit root variables is not unit root then there must be some relationship between them. This is an "if and only if" statement (Green, p. 856). If you find cointegration then a relationship exists, if not it does not. Therefore if you are interested in establishing that a relationship exists between unit root variables, this is equivalent to establishing cointegration. That relationship is called the cointegrating vector. There is a way to write a system that captures all the relationships and avoids unit roots. Consider

$$\Delta x_t = \alpha_1(\beta_1 y_{t-1} + \beta_2 x_{t-1}) + \varepsilon_t + v_t$$

$$\Delta y_t = \alpha_2(\beta_1 y_{t-1} + \beta_2 x_{t-1}) + u_t + v_t$$

This is called a Vector Error Correction Model (VECM). The error correction comes from the cointegrating relationship. The betas contain the cointegrating equation and the alphas the speeds of adjustment. If y and x are far from their equilibrium relationship, either y or x or both must change, the alphas let the data choose. The vector part of the name does not apply to the model above, but it will if the error terms are autocorrelated.

If a set of financial variable is non-stationary, but they tend to move together overtime, then we say that they are bounded by some relationship in the long-run. Long-run is a state of equilibrium where there is no tendency to change since variables are in balance. Long-run models are often termed "static" model, but there is no necessity to actually achieve equilibrium at any point in time. All that required is that the system moves toward the equilibrium defined by the long-run relationship. Since the equilibrium (steady-state) is rarely observed, it is interesting to consider short run evolution especially for predictions. The VECM is a representation of this short-run dynamic. When the system deviate from its equilibrium it is called "error", but it will be corrected in the following way: (1) if it is positive, there is a downward correction in the current period; (2) conversely, a negative error induced upward correction; (3) in a steady-state, there is no error correction, neither changes in y or x. In addition to the error correction mechanism, VECM is able to deal with non-stationarity $I(1)$

variables provided they are cointergrated, since changes in y and x plus the error term are all stationary. Therefore, estimating VECM can help not only obtain short term dynamic but also obtain the long-term equilibrium.

For the analysis of this section we use the sample of the level values of the indexes (not differences or returns). This sample is the level minute bar for the time period August 9th 2012 to October 3rd 2013, it includes outliers and account for missing data with the same procedure applied in sample A to D (Henceforth this level sample will be called sample E). The level values of the indexes are non-stationary (but the first difference is), that is the cointegration is of order $I(1)$). Thus we have two random walk time series, that might have a long-term relationship (i.e., they are cointergrated). Using the Johansen Cointegration test, we find that there is one cointegration equation42 (in both cointegration test measures – Trace and max-eigenvalue.) The error correction term is only significant for the VIX time series (at 1% significance level) but with a very negligible magnitude, which might be a result of the minute bar interval (i.e., there is not much of a disequilibrium "correction" within one minute. It could very well be that the disequilibrium "correction" may take longer. Moreover, this result is consistent with the long-period of autocorrelation). When estimating the VECM we observe that the model's estimated coefficients, which are present also in the VAR model (i.e, the lags of the difference one to eleventh) seems to have the same structure in all aspects – the magnitude, the direction and even the significance (i.e., at what level of significance, if any) – hence, all the discussion on the VAR model in the previous section applies here. It seems that the twelfth lag is somewhat not relevant (the twelfth lag seems to only be relevant when the twelfth lag of the VIX affect the current level of VIX)43. In every sample and estimation method (VAR or VECM; returns versus levels) we observe that the current level of the VIX is related/affected by all (with the exception of the second lag) lags included in the estimation model. This implies that the VIX time series is much more autocorelated than the SPX. Any shock to the SPX will die relatively quickly, while the VIX will carry on the impact of a shock for a relatively long period of time. In Market Microstructure literature, we can refer to that as a "permanent" market impact – VIX has a "permanent" market impact whereas the SPX market impact seems to be more transitory44.

The analyses of the Impulse response and the Variance Decomposition are much more telling. Figure 5, shows that a shock from in the SPX results has a positive impact on the SPX

42 The test is performed with twelve lags as this is the optimal lag for this dataset considering SIC (it is 23 lags with AIC). These results are similar to the optimal lag results for the return model using sample C (the one lag difference can be attributed to the fact that the returns sample by construction is always missing one observation.)

43 We have estimated the VECM model with 11 lags (instead of 12) and found very similar results (only minor differences in the value of the coefficients, but direction and significance are all the same). We can then conclude that even though, the optimal lag as indicated by SIC is 12 lags, for parsimonious reason we should use 11 lags, as there is no loss of information by using 11 lags instead of 12 lags. Tables relating to the VECM model for 12 and 11 lags are available from the author.

44 This observation will be studied further in a research project which will assess permanent and transitory market impact for both the SPX and the VIX (for permanent/transitory market impact see for example – Almgren, Thum, Hauptmann and Li (2005), Almgren and Chris (2000) and Almgren (2003)).

time series, but dies very quickly (within the second time period, i.e., second minute). On the other hand this same shock will have a negative impact on VIX, and will last longer, it will gradually decay up until the 12 period (i.e., this chock has an impact for more than ten minutes). The shock from the VIX has a positive effect on the VIX time series and with almost the same magnitude of the shock from the SPX (on the VIX), but in opposite direction, although the VIX shock seems to die faster than the SPX shock. Indicating that the VIX variations could probably be better explained by changes in the SPX than by its own time series changes (the following variance decomposition analysis supports this assumption). The SPX time series shows no response to shocks in the VIX.

Figure 5: Impulse Response (Level Values of the Indexes)

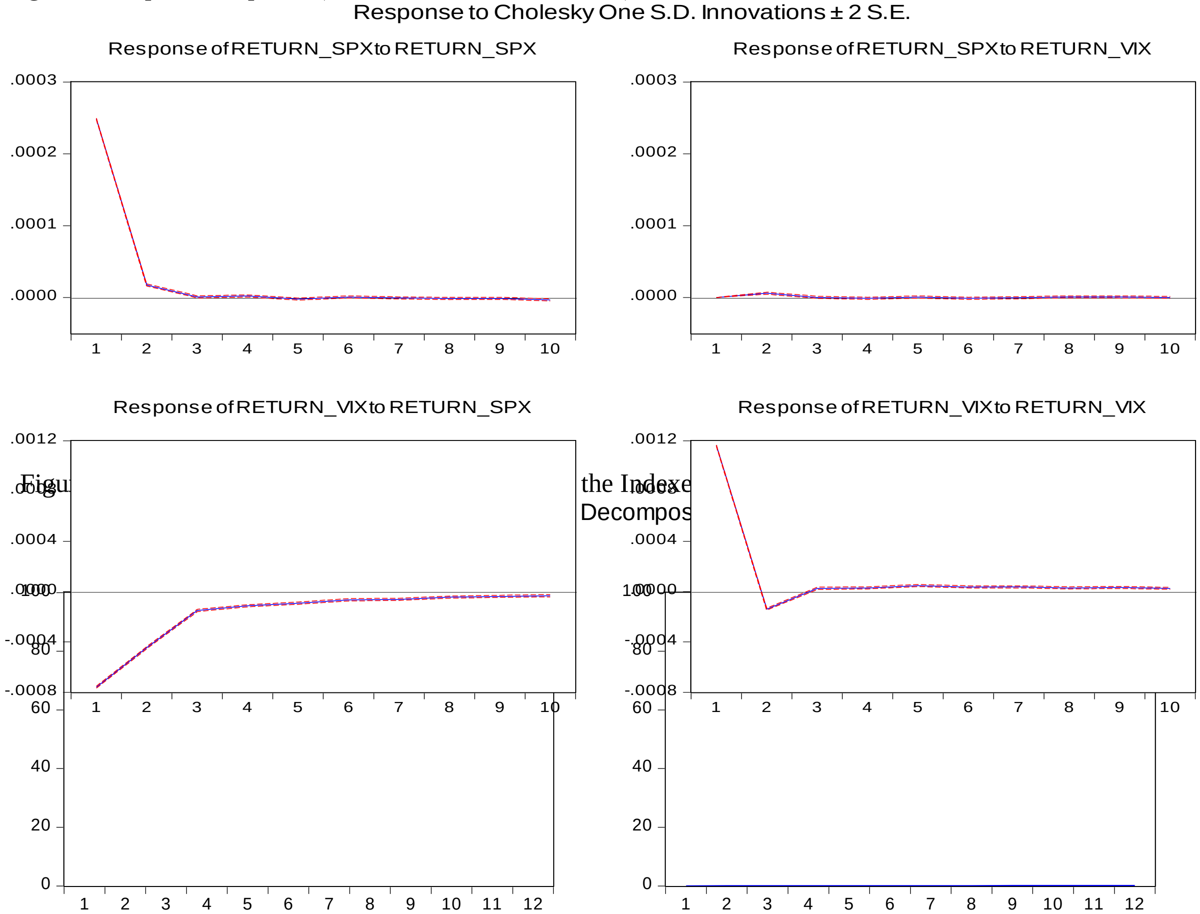

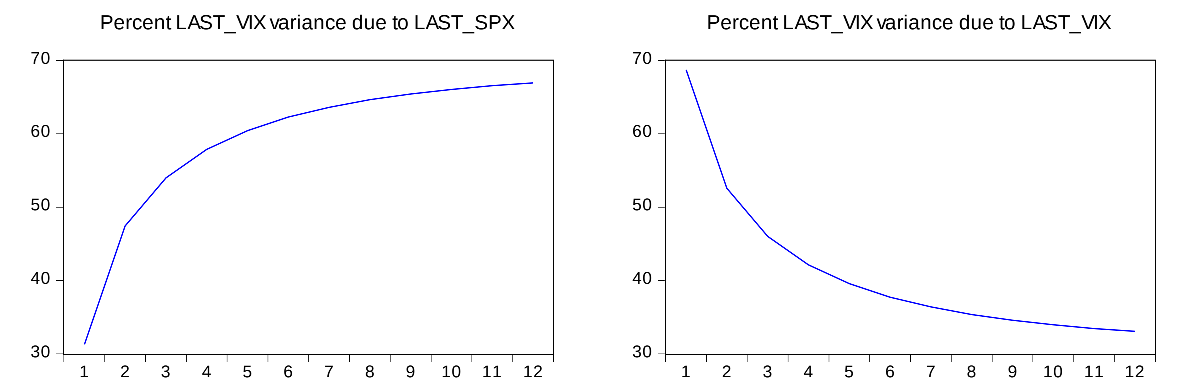

The variance decomposition analysis shown in Figure 6, clearly support our hypothesis that SPX movement cannot be explained by change in the VIX, but the opposite impact does hold. In the first period the decomposition of variance for the VIX is 30% - 70% explained by the SPX and VIX, respectively. This decomposition, however, flips with time – i.e., by the twelfth period (and even sooner) the decomposition become 70%-30% explained by the SPX and VIX, respectively. This suggests that the SPX shock stays in the VIX system longer, which will also suggest that one can “predict” the direction (and to some degree the magnitude) of the VIX in the next ten minutes simply by observing the current movement of the SPX. This is a paramount observation, if one wishes to make trading decision (whether as a trader or a hedger or an investor). It also suggests that what the current level of the VIX captures is merely the current changes in the SPX and has no predictive power or assessment on the future movement of the SPX.

Some Applications

The above analysis suggests an interesting relationship between the movement of the time series of the SPX and the VIX. The question is whether we can make use of these observations for trading/investing decision purposes. One way to show the applicability of the study’s results is by looking at the special event on April 25th, 2013, when the CBOE had a major glitch and had to shut down the exchange for the first half of the trading day (the CBOE opened for regular trading at 1pm, whereas the S&P options contracts resumed at 12:50pm). Since the S&P 500 options exclusively trade on the CBOE, there can be no quoting of the VIX while there is no trading in the options45. For the first half of the April 25th, 2013, we did not know the new levels of the VIX.

The main finding of this study is that one can make an informed assessment about the VIX movement simply by knowing the current movement of the SPX. On April 25th 2013, even though the VIX was not quoted, the SPX was, as all equity exchanges were operating regularly. Hence, by observing how the SPX behaved in the first half of the day one could at least make predictions of the direction of the VIX when the CBOE starts quoting it when it resumes trading on the exchange. The SPX closed the day earlier at 1578.16 and was moving up in early morning due to positive economic reports and then remained steady at a level of about 1589 (about 0.7% increase) until the opening of the trading at the CBOE. Since we know that the SPX negatively affects the VIX, we should expect that the VIX should open lower than its close of 13.61 the day before. Figure 7 shows that, indeed when the VIX opened at 12:50pm it was at a lower level than the Wednesday close of 13.28 (about a 2% decrease). It is evident that we could at least surmise the direction of the VIX, but could we also say something about the level of the VIX (i.e., the magnitude of the change).

45 The VIX is calculated based on the underlying (i.e., SPX) and a variety of options on the SPX (in, out, or at the money), and is derived as the 30 day implied volatility.

Figure 7: SPX versus VIX on April 25th 2013

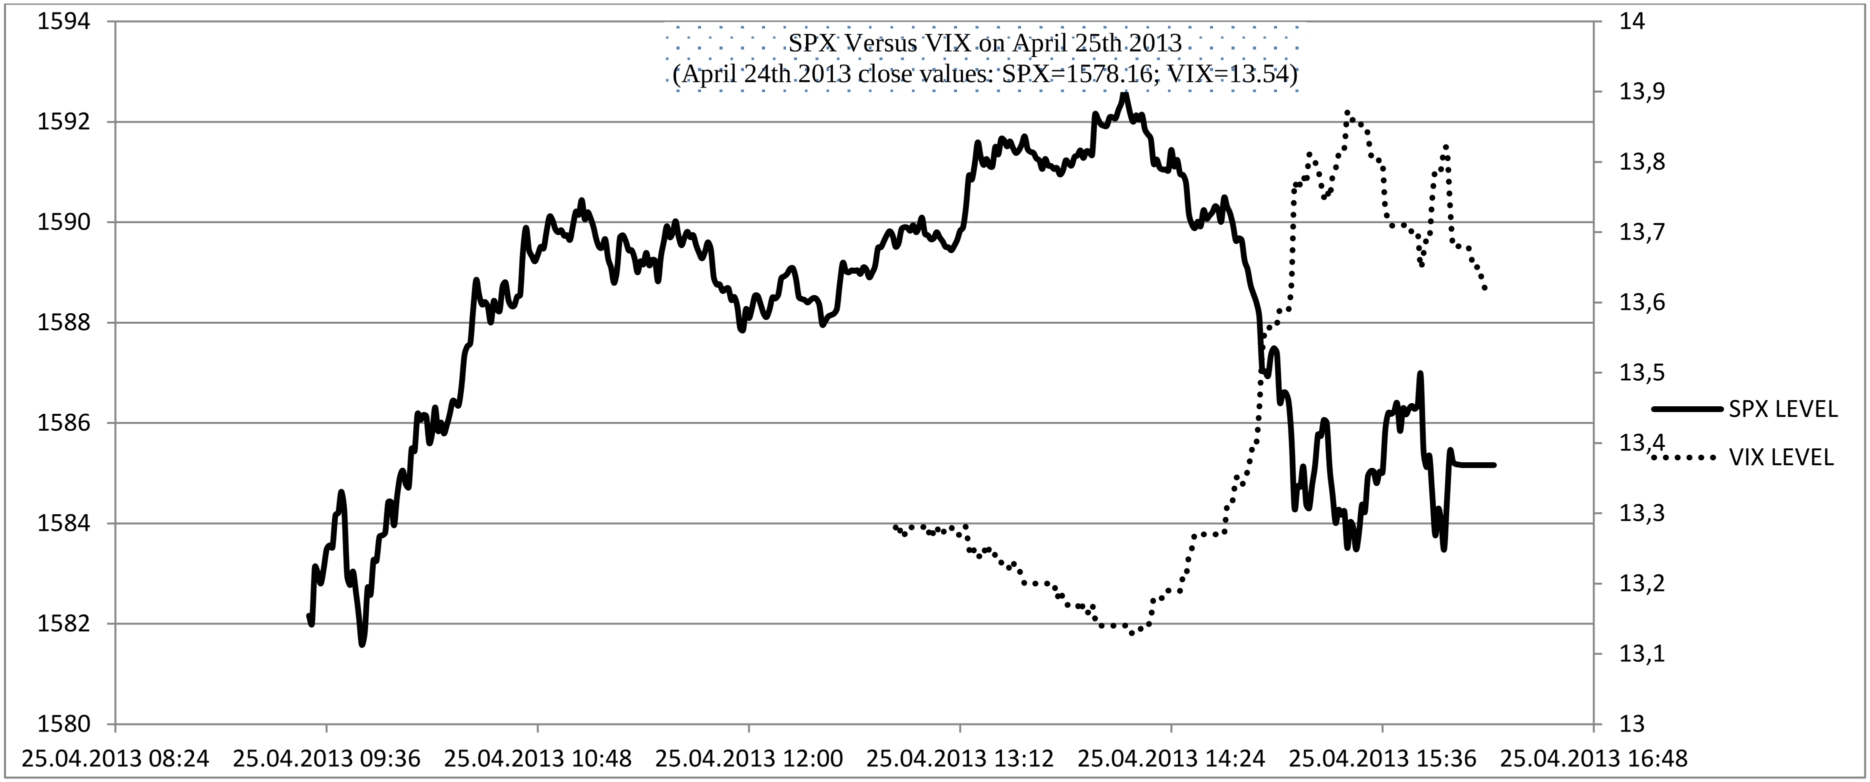

If we use the VAR model estimates and knowing of a 0.7% change, we can calculate that this change will result in a level of 13.32 for the VIX; if we use the VECM model estimates and knowing of about 11 points of an increase in the SPX level, the calculation will result in approximately 13.29 for the VIX. These are quite good estimates considering that the VIX opened at 13.28.

To trade this, one could not have bought options on the VIX, as they are also exclusively traded on the CBOE. One could, however, trade across the street at the CME group with futures on the VIX or options on ETN's that represent the VIX. In addition one could have traded the ETN's themselves (and maybe even take an additional risk with leveraged ETN's if one strongly believes in his/her bet).

Conclusion and Future Research

This paper is the first to examine the causality between the S&P 500 (SPX) and the VIX. Previous research has focused on the predictive power of the VIX, while ignoring the circular mechanism inherited in the calculation of the VIX – a measure of implied volatility, which is based on the prices of options on the SPX. Since the VIX is not a standalone model but a model derived from observed prices (of the options and the underlying – the SPX), it seems a fair assumption that some aspect of causality exist between the SPX and the VIX.

The paper postulate two main hypothesis: (1) VIX "granger cause" SPX; (2) SPX "granger cause" VIX; (3) and a secondary hypothesis – a bi-directional causality relationship between the SPX and the VIX. Using intraday, minute bar time series for the SPX and the VIX, the granger causality test is in favor of the secondary hypothesis. All other estimation models and analyses, however, support the first hypothesis. Thus, we conclude that even though a bi-directional causality relationship might exit, the empirical evidence shows that the impact of the VIX on the SPX is negligible (to none existing) while the impact of the SPX on the VIX is not only significant but also persistent over a relatively long-period of time.

Two main findings are imperative for trading or hedging purposes. First, the VIX time series is autocorrelated over a long period of time, implying that the VIX has "permanent" market impact (as it is referred to in the market microstructure literature), while the SPX market impact seems to be more "transitory". This finding should be further investigated in future research to assess "permanent" and "transitory" market impact for both the SPX and the VIX (using for example the methodology developed by Almgren, Thum, Hauptmann and Li (2005), Almgren and Chris (2000) and Almgren (2003)). The examination will include different tradable vehicles of the VIX – such as ETN's or future contracts.

Second, the SPX shock stays in the VIX system longer, which suggests that one might "predict" the direction (if not, to a certain degree, the magnitude) of the VIX in the next ten minutes simply by observing the current movement of the SPX. It would be worthwhile to further investigate whether the SPX indeed has a predictive power, which can then be applied for investment decision. The analysis should include the following: (1) partitioning on different market conditions; (2) partitioning on different macroeconomic conditions; (3) investigating the impact of news on the significant and persistence of the shock; (4) analyzing the shock components, in an attempt to decompose it to a transitory versus permanent shock – such decomposition might be useful in risk analysis and risk management.

References

Aboniemi Katja, Modeling and Forecasting the VIX Index, working paper, 2008

Almgren Robert, Optimal Execution with Non-linear Impact Functions and Trading-Enhanced Risk, *Applied Mathematical Finance*, 10, 2003, 1-18

Almgren Robert and Neil Chriss, Optimal Execution of Portfolio Transactions, *Journal of Risk*, 3 Winter 5-39, 2000/2001

Almgren Robert, Hauptmann Emmanuel, Li Hong and Thum Chee, Equity Market Impact, *Risk*, July 2005

Anderson T.G., Bollerslev T., Christoffersen P.F. and Diebold F.X., Volatility Forecasting, Working paper, Cambridge 2005

Badi H. Baitagi, Econometrics, Springer, Third Edition, 2002

Barndorff-Nielsen, Ole E, Peter Hansen, Asger Lunde and Neil Shephard , Realized kernel in Practice: trades and quotes, *Econometrics Journal*, 2009

Bekaert Geert and Marie Hoerova, The VIX Premium and Stock Market Volatility, Working Paper, 2013

Bommarito Michael, Top Twenty VIX Proxy ETFs/CEFs, *Seeking Alpha*, 2007

Brenner Menachem, Jinghong Shu and Jin Zhang, The New Market for Volatility Trading, *The Journal of Future Markets*, Vol. 30 No. 9, 809-833, 2010

Brownless, C T and G M Gallo, Financial Econometrics at Ultra-High Frequency: data handling concerns, *Computational Statistics and Data Analysis*, 51, 2232-2245, 2006

Burnham K.P. and Anderson D.R, Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach, Springer, Second Edition, 2002

Burnham K.P. and Anderson D.R, Multimodel Inference: Understanding AIC and BIC in Model Selection, *Sociological Methods and Research*, 33, 2004

Carr peter and Liuren Wu, A Tale of Two Indices, *The Journal of Derivatives*, Spring 2006

Cipollini Alessandro and A. Manzini, Can the VIX Signal Market Direction? An Asymmetric Dynamic Strategy, Working paper, April 2007

Chung San-Lin, W. Tsai, Y.Wang and P.Weng, The Information Content of the S&P 500 Index and VIX Options on the Dynamics of S&P 500 Index, Working paper, 2011

Robert Daigler and Zhiguang Wang, The Performance of VIX Option Pricing Models: Empirical Evidence Beyond Simulation, Working Paper, 2009

Doran S. James, R. S. Goldberg and E.I. Ronn, A Simple Model for Time –Varying Expected Returns on the S&P 500 Index, Working paper, 2008

Enders Walter, Applied Econometric Time Series, Willey, Third Edition 2009

Falkenbry, T N, High Frequency Data Filtering, unpublished technical report.

Figuerola-Ferretti Isabel and I.Paraskevopoulos, Pairing Market and Credit Risk, Working paper, March 2011

Gayawan Ezra and Reuben Ipinypmi, A comparison of Akaike, Schwartz and R square Criteria for Model Selection Using Some Fertility Models, *Australian Journal of Basic and Applied Sciences*, 3(4): 3524-3530, 2009

Giot P., Relationship between Implied Volatility Indexes and Stock Index Returns. Are Implied Volatility Indexes Leading Indicators? , *Jornal of Portfolio Management*, 2005

Granger C.W.L. and Poon S.H., Forecasting Volatility in Financial Markets: A Review, *Journal of Economic Literature*, Vol. XLI, 2003

Greene William, Econometric Analysis, 7th Edition, Prentice Hall, 2011

Hacker R. Scott and Hatemi-J Abdulnasser, Tests for Causality between Integrated Variables Using Asymptotic and Bootstrap Distributions: Theory and Application, *Applied Economics*, Vol. 38, Issue 13, 2006

Hamilton D. James, Time Series Analysis, Princeton University Press, 1994

Hao Jinji and Jin Zhang, GARCH Option Pricing Models, the CBOE VIX and Variance Risk Premium, *Journal of Financial Econometrics Advance Access*, Vol. 0, No. 0, 1-25, Jan 2013

Heydon T., Fear and Greed in Global Asset Allocation, *Journal of Investing*, Vol. 9, Issue 1, 2000

Lui Qaing and Gaoxiu Qiao, Variance Risk Premium and VIX Pricing: A Simple GARCH Approach, Working Paper, 2012

Madhavan Ananth, Market Microstructure: A Survey, *Journal of Financial Markets*, Vol. 3, 205-258, 2000

Marsh Terry and P.Pfleiderer, The Relation between Fixed Income and Equity Return Factors, Stanford Rock Center for Corporate Governance, January 2010

Nossman Marcus and A. Wihelmsson, Is the VIX Futures Market Able to Predict the VIX Index? A Test of the Expectation Hypothesis, Working paper, 2008

Oxford-Man Institute Realized Library Documentation

Ozair Merav, An Investigation of the Relationship between the VIX and the S&P 500 Index, Working Paper, 2011

Swensen David, High-Yield Bonds and Their Correlation to Equities, *the Educated Investor*, December 2008

Robert Whaley, Understanding VIX, Working Paper, 2008

Yang Y., Can the Strengths of AIC and BIC be Shared?, *Biometrika*, 92, 2005

The CBOE Volatility Index – VIX, Chicago Board Options Exchange (CBOE), 2009

Investigating the VIX-S&P 500 Correlation, *Seeking Alpha*, 2008