

On the Intraday Relation between the VIX and its Futures

Bart Frijns*

Alireza Tourani-Rad

Robert Webb

*Corresponding author. Department of Finance, Auckland University of Technology, Private Bag 92006, 1142 Auckland, New Zealand. Tel. +64 9 921 9999 (ext. 5706); fax. +64 9 921 9940; Email: bfrijns@aut.ac.nz.

On the Intraday Relation between the VIX and its Futures

Abstract

We study the intraday dynamics of the VIX and VXF for the period January 2, 2008 to December 31, 2012. Applying a Vector Autoregression (VAR) model on daily data, we observe some evidence of causality from the VXF to the VIX. However, estimating a VAR using our ultra-high frequency data, we find strong evidence for bi-directional Granger causality between the VIX and the VXF. Overall, this effect appears to be stronger from the VXF to the VIX than the other way around. Impulse response functions and variance decompositions analysis further confirm the dominance of the VXF. Lastly, we show that the causality from the VXF to the VIX has been increasing over our sample period, whereas the reverse causality has been decreasing. This finding suggests that the VIX futures have become increasingly more important in the pricing of volatility. We further document that the VIX futures dominate the VIX more on days with negative returns, and on days with high values of the VIX, suggesting that on those days investors use VIX futures to hedge their positions rather than trading in the S&P 500 index options.

Keywords: VIX, Futures, Vector Autoregressions, Ultra-High Frequency Data

JEL Codes: C11, C13.

1. Introduction

The Chicago Board Options Exchange (CBOE) introduced the CBOE Volatility Index (VIX) in 1993¹. The index has become to act as the benchmark for stock market volatility and, more generally, investor sentiment². The VIX has proven to be very useful in forecasting the future market direction, especially during high volatile periods. Forecasting qualities of the VIX outperform traditional volatility measures based on historical realized volatility and GARCH models (Corrado and Miller, 2005 and Carr and Wu, 2006). However, while the VIX could be used for hedging purposes, it could not easily be traded. Theoretically, it would be possible to replicate a portfolio of the underlying options in the VIX index and maintain the 30-day interpolated maturity; but the costs would be exorbitant. In order to expedite trading in volatility, as well as increasing hedging opportunities, the CBOE introduced futures on the VIX (henceforth referred to as VXF) on March 26, 2004. The VIX futures contracts have become very popular. Due to the existence of a strong negative correlation between S&P 500 index returns and the VIX, investors use them for hedging equity portfolios. Interestingly, VIX futures have proven to be a far more convenient hedging tool than S&P option index (Szado, 2010). In 2006, the CBOE introduced the VIX options. Lastly, the first volatility exchange traded instrument, termed the VXX, was introduced by Barclays in January 2009.

Black (1975) suggests “volatility is crucial” and the informed investors will take advantage of their views on the volatility. They trade in derivatives markets because of the leverage effect

¹ The VIX was originally based on implied volatilities, with 30 days to expiration, of eight S&P 100 at-the-money put and call options (Whaley, 1993). Later on, in 2013, the VIX was expanded to include options based on a broader index, the S&P 500, reflecting a more accurate view of market volatility. The valuation model was also changed to a model-free basis (Britten-Jones and Neuberger, 2003).

² The VIX has also been termed “investor fear gauge” (Whaley, 2000, 2009). Generally, when investors expect the stock market to fall, they will buy S&P put options for portfolio insurance. By doing so, investors push up the option prices and ultimately the level of VIX.

such instruments offer. Therefore, what takes place in the derivative markets has implications for the pricing of financial assets. In more recent years, “volatility” has been considered as an asset class on its own (Grant, Gregory, and Liu, 2007). Several studies report significant diversification benefits of an addition of volatility, long position on VIX derivatives, to equity portfolios (Daigler and Rossie, 2006 and Alexander and Korvilas, 2011), especially during steep equity market declines.

The evolution of the VIX in the past two decades, alongside the subsequent introduction of futures and options on VIX and the accompanying extensive trading therein, indicate a resilient demand by investors for volatility-related financial instruments. At the same time, there has been an ever increasing research, by both academics and practitioners, into many aspects of VIX and its derivatives.

In this paper, we are particularly interested in the dynamic relation between the VIX and its futures. With the introduction of the VIX futures, investors can hedge volatility either using the options on the S&P 500 or by taking positions in the VIX futures. Both the demand for options in the index or the VIX futures provides useful information regarding the market’s expectation of future volatility. An important question that then begs investigation is where information on future volatility enters the market, and which market would lead in terms of incorporating this new information. Prior research has used daily data to (partially) address this question (see e.g. Konstantinidi, Skiadopolous and Tzagkaraki, 2008; Konstantinidi and Skiadopolous, 2011; and Shu and Zhang, 2012). However, inferring informational efficiency and leadership is difficult using daily data as informational asymmetries between the markets could vanish in the data aggregation. In this study, we therefore examine the dynamic relation between the VIX and the near-term VIX futures using ultra-high frequency data (which

currently is the highest frequency at which the VIX is calculated). Sampling at this frequency eliminates issues related to data aggregation, and allows us to obtain a clearer picture on the informational efficiency and leadership in the relation between the VIX and its futures. To our knowledge, we are the first to examine the relation between the VIX and its futures using intraday data.

We study the intraday dynamics of the VIX and VXF for the period January 2, 2008 to December 31, 2012. Considering first the results of a Vector Autoregression (VAR) using daily data, we observe that there is some evidence of Granger causality from the VXF to the VIX, which is in line with Shu and Zhang (2012). However, the fit of this model is poor and the model is left with a considerably high residual correlation of approximately 0.8. Estimating a VAR using our ultra-high frequency data, we find that the residual correlation is negligible, and we find strong evidence for bi-directional Granger causality between the VIX and the VXF. Overall, this effect appears to be stronger from the VXF to the VIX than the other way around. Impulse response functions and variance decompositions analysis confirm the dominance of the VXF. Next, we show that the causality from the VXF to the VIX has been increasing over our sample period, whereas the reverse causality has been decreasing. This result suggests that the VIX futures have become more and more important in the pricing of volatility. We further document that the VIX futures dominate the VIX more on days with negative returns, and on days with high values of the VIX, suggesting that those are the days when investors use VIX futures to hedge their positions rather than trading in the S&P 500 index options.

The remainder of this paper is structured as follows. In section 2, we review some of the relevant literature. Section 3 describes data used in this paper and presents some summary statistics. In section 4, we present our results. Finally, section 6 concludes.

2. Literature

Apart from the literature focusing on the volatility pricing models (e.g, Zhu and Lian, 2012; Lu and Zhu, 2010; Brenner, Shu and Zhang, 2008, Lin, 2007, Zhang and Zu, 2006) and the addition of a long VIX futures position to equity portfolios (Szado, 2010 and Alexander and Korovilas, 2011) there is a limited number of empirical papers investigating the efficiency of the VIX and VIX futures markets. As for the issue of price discovery and causality in the market for volatility, there is one study similar to ours, i.e., Shu and Zhang (2012). In this section, we briefly provide an overview of these studies on the VIX and its futures.

Konstantinidi, Skiaopoulos and Tzagkarakaki (2008) and Konstantinidi and Skiaopoulos (2011) investigate the behavior of the implied volatility indices, for the US and Europe, to assess whether they are predictable. While the authors observe significant predictable patterns in the futures on implied volatility indices, none of these patterns can be exploited through active trading strategies that yield economically significant profits. Hence, from an economic point of view, they cannot reject the efficiency of the volatility futures markets.

Nossman and Wilhelmson (2009) test the expectation hypothesis, using information on the term structure of volatility, to test the efficiency of the VIX futures market. When they allow

for the existence of a volatility risk premium in their analysis, Nossman and Wilhelmsson (2009) cannot reject futures market could predict the future VIX levels correctly.

The paper closest to our study is that of Shu and Zhang (2012). In this paper, authors explicitly examine price discovery between the VIX and the VIX futures. They use daily prices for the period 2004-2009. Shu and Zhang (2012) find that the VIX and the futures are indeed cointegrated and proceed by using a Vector Error Correction model to assess the lead-lag interaction between spot VIX and VIX futures. They find that VIX futures are informative about spot VIX and lead the spot market in a linear error correction model. Overall, they conclude that the VIX futures have some price discovery function.

3. Data and Summary Statistics

We obtain intraday data for the VIX and the futures on VIX from the Thomson Reuters Tick History database (TRTH). We collect data for the period January 2, 2008 to December 31, 2012. All data are collected in tick time, with potential millisecond precision. The CBOE computes intraday values for VIX at approximately 15-second intervals from 8.30 a.m. – 3.15 p.m. Chicago time (note that this time period reflects the normal trading hours for the S&P 500 index options). The time interval between the calculations of the VIX is not exactly 15 seconds, but slightly more. This implies that at the start of the day VIX may be computed at 9:30:15.12, 9:30:30.24, etc. but during the trading day may be reported at, say, 10:30:22.54, etc.

The VIX futures (VXF) were first listed on the CBOE futures exchange on March 6, 2004. The contracts use the VIX as the underlying and use a multiplier of \$1,000. The contracts trade on CBOE Direct, the electronic trading platform of the CBOE. The minimum tick size of these contracts is 0.01 index points. The CBOE lists up to 9 near-term serial months and five months on the February quarterly cycle. The regular trading hours for the VXF are from 8:30 a.m. to 3:15 p.m., the same intraday period over which the VIX is computed. Trading of the contracts terminates on the last day before the final settlement date. We collect tick-by-tick trade and quote data for all contracts traded during our sample period. However, to construct a continuous series, we splice together the nearest-term contracts, which are rolled over on the day when the trading volume of the second nearest-term contract exceed the trading volume of the nearest term contract.³ Since the contracts trade on an electronic market we collect the bid and ask quotes for these contracts and compute the midpoint from these quotes. This frequency and sample period provides us with a total of 1,987,254 observations for each series.

As the VIX is computed approximately every 15 seconds, we sample VIX and the VXF at a 15 second frequency.⁴ In addition, we exclude the first and last five minutes of the trading day, to avoid any noise that may be due to opening and closing affecting our results.

³In some cases the nearest-term contract remains the most active one until the settlement date of the contract. In this case, we roll the contract over on the day before the last trading day, to avoid any price distortions that may be due to final settlement of the contracts.

⁴Since VIX is computed at slightly more than 15 seconds, we construct 3 series, starting on the whole minute and 5 and 10 seconds after the minute. By constructing these three series, we aim to minimize the impact of stale values of VIX that could affect our results. Note that we only report the results for the sampling interval that starts on the whole minute. Other sampling intervals yield similar results and are available upon request.

INSERT FIGURE 1 HERE

In Figure 1, we provide a time series plot of the data. The plot shows that the VIX was relatively low at the start of our sample period but swiftly increased and doubled from early August 2007 onwards, which can be attributed to the European Sovereign Debt Crisis. For the remainder of the period the VIX stayed relatively high, due to the continuation of the global financial crisis, and was decreasing only at a very slow rate. The VXF resembles the pattern of the VIX closely, but as expected the VXF is generally below the VIX when the VIX is high and is above the VIX when the VIX is low (see also Shu and Zhang, 2012).

In Table 1, we present summary statistics on the VIX and the VXF, respectively. In the first two columns of Table 1, we report the summary statistics for the levels of the VIX and the VXF. Over our sample period, the VIX was on average 25.83, while the VXF was slightly higher at 26.43. As can be seen from maximum and minimum values, the VIX has more extreme values than the VXF. This is also reflected in the standard deviation, which is higher for the VIX than for the VXF. Both series have positive skewness, as can be expected and have excess kurtosis. The persistence, at the 15 second frequency, is extremely high and the first order autocorrelation is not discernibly different from 1.00. Finally, when conducting a unit root test on both series, we observe that we can reject the presence of a unit root for the VIX at the 1% level, while for the VXF we can only reject the unit root at the 10% level.

INSERT TABLE 1 HERE

The last two columns of Table 1 report summary statistics for the first difference of the (log) VIX and VXF. These first differences only include changes during the trading day, and hence exclude the overnight change. On average, the changes in the VIX and VXF are nearly zero as could be expected at these are ultra-high frequencies data. Again, as we observed in the levels, ΔVIX has more extreme changes than ΔVXF as can be seen from the maximum and minimum values, and from the standard deviations. Interestingly, the skewness of ΔVIX is negative at -0.35, whereas the skewness of ΔVXF is nearly zero. Although both series display excess kurtosis, the excess kurtosis in ΔVIX is much higher than that in ΔVXF . The first order autocorrelation in both series is negative, at a value of -0.084 for ΔVIX and -0.118 for ΔVXF , suggesting that there is stronger negative autocorrelation in the VIX futures. Finally, the ADF statistics, for the difference series, suggest that we can strongly reject the presence of a unit root in both series.

4. Results

4.1 Daily Analysis

To assess the relationship between the VIX and the VIX Futures, we start by conducting our analysis at a daily frequency as in Shu and Zhang (2012). Since the summary statistics show that there is weak evidence of a unit root in the VIX futures, we compute first differences of the log of the volatility series. We then estimate the following VAR:

$$\begin{cases} \Delta VIX_{t+1} = \alpha_1 + \phi_1(L)\Delta VIX_t + \varphi_1(L)\Delta VXF_t + \varepsilon_{1t+1} \\ \Delta VXF_{t+1} = \alpha_2 + \phi_2(L)\Delta VIX_t + \varphi_2(L)\Delta VXF_t + \varepsilon_{2t+1} \end{cases}, \quad (1)$$

where $\phi_1(L)$, $\varphi_1(L)$, $\phi_2(L)$, and $\varphi_2(L)$, are polynomials in the lag operator of identical length. Daily data suggests an optimal lag length of one using the Schwartz Information Criterion (SIC). Hence, we estimate Equation (1) as a VAR(1).

In Table 2, we report the regression results for the VAR(1) as well as Newey-West adjusted t-statistics in parentheses. For ΔVIX_t , we find a negative and significant coefficient on ΔVIX_{t-1} suggesting that there is negative autocorrelation in changes in the VIX at the daily frequency. We also find a positive coefficient for ΔVXF_{t-1} , significant at the 5% level, suggesting that at the daily frequency the VIX futures have some predictive value for the changes in the VIX. For ΔVXF_t , we find that there is no statistical evidence for predictability of the changes in the VIX futures, neither lagged changes in the VIX or the VXF can be used to predict these changes. Although we find some evidence of predictability for ΔVIX , we note that the R^2 's of the regressions are quite low at 1.82% and 0.18% for ΔVIX_t and ΔVXF_t , respectively. The contemporaneous correlation between the residuals in the VAR is quite high at 0.813, suggesting a strong correlation between the series. Granger causality tests, reported in Panel B, confirm the findings of the coefficients, i.e., there is Granger Causality from ΔVXF to ΔVIX , but not the reverse. Overall, the results of our daily analysis are in line with Shu and Zhang (2012).

4.2 Intraday Analysis

The daily analysis reveals some evidence for predictability of changes in the VIX. However, this analysis also revealed a very high contemporaneous correlation between the changes in the VIX and the VXF. Part of this contemporaneous correlation may be due to data aggregation. Sampling at higher frequencies may reduce the contemporaneous correlation and reveal more lead-lag dynamics. We therefore estimate the VAR in Equation (1) using intraday ultra-high-frequency data sampling at a 15 second frequency. Instead of estimating Equation (1) as one big VAR using 1,987,254 observations, we estimate the model every day in the sample, so that we obtain a daily series of coefficients and statistics.

First, we determine the optimal lag length of the VAR by computing the SIC for 1 lag up to 10 lags every day. Then, we compute the average SIC over all days in the sample. We find that the average SIC is lowest for a VAR with three lags. Hence, we estimate all coefficients and compute all statistics based on daily models using three lags.

In Panel A of Table 3, we report the results for the intraday VAR(3) model. We report coefficients and indicate significance using asterisks (based on Newey-West corrected standard errors obtained from the time series of coefficients). In brackets, we report the 2.5% and 97.5% percentile values from the time series of coefficients.

When we consider the dynamics of ΔVIX , we find evidence of some negative autocorrelation, with the first and second lag significant at the 5% level and values of -0.047

and -0.020, respectively. The coefficient for the third autoregressive lag is positive at 0.005 and although very small, it is significant at the 10% level. For the coefficients on the lagged values of ΔVXF , we find that all three coefficients are positive and significant at the 1% level, with values of 0.172, 0.119, and 0.058, for lags one, two and three, respectively. This provides evidence that there is some predictability for the changes in the VIX based on lagged changes in the VXF. The R^2 of this regression, 9.26%, is considerably higher than that for the regression using daily data, suggesting that much more of the variation in the intraday changes in the VIX can be explained by past information. For the changes in the VXF, we note that lagged values of ΔVIX have a positive and significant effect on changes in the VXF for all three lags. This finding suggests that there is some predictability of changes in the VXF based on lagged values of ΔVIX . This contrasts the findings that we observed for the daily estimation. We also find significant evidence for negative autocorrelation in the changes in the VIX futures at this frequency, with all three lags yielding negative and significant coefficients. Again, compared with the daily VAR, the R^2 of the intraday VAR is considerably higher at 4.12%, but still lower for the regression for ΔVIX . At this high level of data frequency, we also observe that the contemporaneous correlation in the changes in VIX and VXF is close to zero (on average 0.0019, with 2.5% and 97.5% percentile values at -0.0001 and 0.0058, respectively). This finding indicates the presence of high positive contemporaneous correlation between daily changes in VIX and VXF is driven by data aggregation.

INSERT TABLE 3 HERE

In Panel B of Table 3, we report the results for Granger causality tests, where we report the average Granger causality statistic over the sample period and report the percentages of significant Granger causality statistics at conventional significance levels. When we consider the Granger causality from ΔVXF to ΔVIX , we find that the average coefficient is equal to 102.80, thus giving very strong statistical evidence for a causal relation running from ΔVXF to ΔVIX (note that the 1% critical value for this statistic is 11.30). When we consider the percentage of days that we find a significant effect of ΔVXF on ΔVIX , we find that between 92.37% (85.45%) of the days there is significant causal effect measured at the 10% (1%) significance level. When we investigate the causal effect of ΔVIX on ΔVXF , we find an average test statistic of 14.31, which, although lower than the causality statistic of ΔVXF on ΔVIX , is still highly significant. Next, we compute the percentage of days that there is a causal effect of ΔVIX on ΔVXF . It is observed that on 63.99% (42.05%) of the days there is significant evidence for causality running from ΔVIX to ΔVXF measured at the 10% (1%) level. Overall, the Granger Causality tests reveal that there is significant evidence for reverse causality, but the effect appears to be stronger running from ΔVXF to ΔVIX than vice versa.

Given that the contemporaneous correlation between ΔVIX and ΔVXF is nearly zero, we can easily estimate impulse response functions and compute variance decomposition assuming that both series are contemporaneously uncorrelated.⁵ In Figure 2, we plot the impulse response functions for 10 steps ahead, where we apply a unit shock to the residuals of each series. Panels A and B of Figure 2 show the responses to a unit shock in ΔVIX . As can be

⁵Note that this additional analysis is difficult to conduct in a meaningful way using daily data, due to the high contemporaneous correlations observed in daily data.

seen, this shock has some impact on changes in the VIX, but dies out after about four periods. Panel B shows that this shock leads to a small increase in the VXF after which it decreases and again dies out after about four periods. Overall, a unit shock to ΔVIX does not lead to a change of more than 0.1 (in absolute terms) in the VXF. Panels C and D show the responses of a unit shock in ΔVXF . Considering Panel D first, we note that a unit shocks to ΔVXF leads to a drop in the VXF after about two periods, and dies out after about 5 periods. The response of ΔVIX to a shock in ΔVXF (Panel C), shows a positive reaction in ΔVIX after two periods and dies out again after about 5 periods. Overall, the impulse response analysis shows that there is bidirectional spillover between ΔVIX and ΔVXF , however, the response in the VIX to a shock in the VXF is substantially larger than the response of the VXF to a shock in the VIX.

The last rows of Table 3 report the Variance Decomposition, where we decompose the variance of ΔVIX (ΔVXF) that is due to either ΔVIX or ΔVXF . When we decompose the variance of ΔVIX , we find that 94.53% of the variance comes from ΔVIX , whereas 5.47% originates from changes in the VXF. Vice versa, the variance of changes in the VXF is for 3.54% due to changes in the VIX and the remaining 96.46% is due to its own changes. In line with the Granger causality tests and the impulse response functions, this result suggests that the changes in the VXF have a greater influence on the changes in the VIX than the other way around.

4.3 Time-variation in Granger Causality

The next question we address is whether we observe any time variation in the intraday causal relation between ΔVIX and ΔVXF . In Table 4, we report Granger causality statistics per year during our sample period. The first column of Table 4 reports the causality statistics from ΔVXF to ΔVIX . We note that since 2008, there is an upward trend in the causality statistics, indicating that causality from ΔVXF to ΔVIX has becoming stronger. However, we also observe a slight drop off in the causality statistic in the last year in the sample in 2012. For the causality in the opposite direction, we observe a downward trend in the statistic going from an average of 25.896 in 2008 to 7.151 in 2012, which is just below the 5% significance level. Hence the causal effect of ΔVIX on ΔVXF seems to have died off over time.

INSERT TABLE 4 HERE

In the next two columns of Table 4, we report the average statistics for the variance decomposition per year over our sample period. The first of these columns reports the percentage of variance of ΔVIX that is attributable to ΔVXF . Again, we note an upward trend over time going from 2.62% in 2008 to 10.04% in 2011, after which it declines to 5.86% in 2012. The percentage of variance of ΔVXF attributable to ΔVIX shows a less clear picture, which starts at 4.66% in 2008 and declines to 3.06% in 2012.

To provide a visual representation of the time variation in the Granger Causality between ΔVIX and ΔVXF , we plot the 10-day moving average of the log of the ratio of the Granger

causality statistics in Figure 3, i.e. $\log(GC_{\Delta VXF}/GC_{\Delta VIX})$, where $GC_{\Delta VXF}$ is the Granger causality statistic of causality from ΔVXF to ΔVIX , and $GC_{\Delta VIX}$ is the Granger causality statistic of causality from ΔVIX to ΔVXF . From this plot, we observe that there is an increase in this ratio. At the start of the sample period the ratio is less than zero, suggesting that the causality from ΔVIX to ΔVXF is stronger. However, this ratio quickly becomes positive. Overall, there is an upward trend in this ratio.

INSERT FIGURE 3 HERE

In order to explain what drives the changes and increase in this ratio, we conduct the following regression analysis, i.e.

$$\log\left(\frac{GC_{\Delta VXF}}{GC_{\Delta VIX}}\right)_t = \alpha + \beta trend_t + \gamma_1 R_SP_t + \gamma_2 \log(VIX)_t + \gamma_3 \log\left(\frac{FutVol}{OptVol}\right)_t + \varepsilon_t, \quad (2)$$

where $\log\left(\frac{GC_{\Delta VXF}}{GC_{\Delta VIX}}\right)_t$ is the daily ratio of the Granger causality statistics, $trend_t$ captures the time trend that we observed in the ratio, R_SP_t is the return on the S&P 500 index on day t , $\log(VIX)_t$ is the log of the VIX on day t , and $\log\left(\frac{FutVol}{OptVol}\right)_t$ is the log of the ratio of daily

volume traded in the near-term VIX futures relative to the daily volume traded in the near-term S&P 500 index options on day t .⁶

In Table 5, we report the results for Equation (2), where we include variables step-by-step. We report all coefficients and Newey-West corrected t-statistics in parentheses. First, we estimate the regression only including the time trend. The results, in the first column of Table 5, show that this time trend is highly significant, and this regression produces an adjusted R^2 of 28.43%. This provides a clear evidence of an upward trend in the informativeness of the VIX futures over the VIX. In the second column of Table 5, we add the returns on the S&P500 index. The coefficient on these returns is negative and significant at the 1% level. This finding suggests that VIX futures are more informative on days when returns are negative and could be due to the increased hedging using VIX futures on days with negative returns. The adjusted R^2 increases slightly relative to the model that only includes the time trend to 29.28%.

INSERT TABLE 5 HERE

Next, we include the log of the VIX. We find that the coefficient for this term is positive and significant at the 1% level. Hence this result suggests that when the VIX is high, the informativeness of the VXF relative to the VIX increases. Again, this could be due to the

⁶Note that we use a detrended version of this ratio as there is a strong positive upward trend in this variable. We have also conducted the analysis with the ratio of daily volume traded in the near-term VIX futures relative to the volume traded in the near-term S&P500 index put options. The results for this analysis are nearly identical to the ones reported here.

increased hedging activity in using VIX futures, when uncertainty (VIX) in the market is high. The adjusted R^2 of this regression is 31.40%, suggesting that the VIX is more informative for the causality ratio than the returns on the S&P 500. In the next column, we include the ratio of volume traded in the VIX futures relative to the volume traded in the S&P 500 options. One reason why we could expect an increase in the informativeness of the VXF is that it is more efficient for investors to use the VIX futures to hedge their positions than the options on the S&P500 (Szado, 2009) and most likely more and more investors are using the futures. We observe, however, that the ratio is insignificant in this regression, showing that an increase in activity in the VIX futures relative to the options does not explain the ratio of causality. Finally, we estimate the regression model where we include both the returns on the S&P 500 and log of the VIX. Prior literature has shown that there is a strong and negative relation between these two variables (e.g. Whaley, 2000) and hence we need to include both in a single regression to determine whether the negative relation between the returns on the S&P500 and the relative informativeness of the VXF is driven by the VIX, and vice versa. When we include both variables, we observe that both maintain their sign and significance, suggesting that both returns on the S&P500 and the VIX are informative for the causality between the VIX and its futures.

5. Conclusion

In this paper, we examine the intraday dynamics of the VIX and VXF for the period January 2, 2008 to December 31, 2012. In line with Shu and Zhang (2012), we observe some evidence of Granger causality from the VXF to the VIX at a daily level. However, at the intraday level, we find strong evidence for bi-directional Granger causality between the VIX

and the VXF. Overall, this effect appears to be stronger from the VXF to the VIX than the other way around, which is confirmed by impulse response functions and variance decompositions. We document that the causality from the VXF to the VIX has been increasing over our sample period, whereas the reverse causality has been decreasing. This suggests that the VIX futures have become progressively more important in the pricing of volatility. We further find that the VIX futures dominate the VIX more on days with negative returns, and on days with high values of the VIX. These findings indicate that on those particular days investors use VIX futures to hedge their positions rather than trading in the S&P 500 index options. Overall, our results suggest that the VIX futures are informationally dominant over the VIX in reflecting future volatility.

References

Alexander, C. and Korvilos, D. (2011). The Hazards of volatility diversification, ICMA Centre Discussion Paper in Finance, No. DP2011-04.

Black, F. (1975). Fact and Fantasy in the Use of Options, *Financial Analysts Journal*, 3, 36-41.

Brenner, M., Shu, J., and Zhang, J. (2008). The market for volatility trading: VIX futures, working paper, New York University, Stern School of Business.

Britten-Jones, M., and Neuberger, A. (2000). Option Prices, Implied Price Processes, and Stochastic Volatility, *Journal of Finance*, 55, 839-866.

CBOE (2009). VIX: The CBOE Volatility Index, White Paper, Chicago Board Options Exchange, (Available at www.cboe.com/VIX.)

Carr, P. and Wu, L. (2006). A tale of two indices, *Journal of Derivatives*, 13, 13-29

Corrado, C., and Miller, T. (2005). The Forecast Quality of CBOE Implied Volatility Indexes. *Journal of Futures Markets*, 25, 339-373.

Fernandes, M., Medeiros, and Scharth, M. (2007). Modeling and Predicting the CBOE Market Volatility, Working Paper, Queen Mary University.

Grant, M., M. K. Gregory, and Lui, J. (2007). Volatility as an Asset, Goldman Sachs Global Investment Research. November.

Konstantinidi, E, Skiadopolous, G., and Tzagkaraki, E. (2008). Can the evolution of implied volatility be forecasted? Evidence from European and US implied volatility indices, *Journal of Banking and Finance*, 32, 2401-2411.

Konstantinidi, E, and Skiadopolous, G. (2011). Are VIX Futures Prices Predictable? An Empirical Investigation, *International Journal of Forecasting*, 27, 543-560.

Lin, Y. N. (2007). Pricing VIX Futures: Evidence from Integrated Physical and Risk Neutral Probability Measures, *Journal of Futures Markets*, 27, 1175-1217.

Lu, Z. and Zhu, Y. (2010). Volatility Components: The Term Structure Dynamics of VIX Futures. *Journal of Futures Markets*, 30, 230-256.

Mixon, S. (2007). The Implied Volatility Term Structure of Stock Index Options. *Journal of Empirical Finance*, 14, 333-354.

Nossman, M. and Wilhelmson, A. (2009). Is the VIX Futures Market Able to Predict the VIX Index? A Test of the Expectation Hypothesis, *The Journal of Alternative Investment*, Fall, 54-67

Shu, J. and Zhang, J. (2012). Causality in the VIX futures market. *Journal of Futures Markets* 32, 24-46.

Szado, E., (2009). VIX futures and options – A case study of portfolio diversification, *Journal of Alternative Investments*, Fall, 68-85

Toikka, M., E.K. Tom, S. Chadwick, and M. Bolt-Christmas (2004). Volatility as an Asset? CSFB Equity Derivatives Strategy, February 26.

Zhang, J., and Zhu, Y., (2006). VIX Futures, *Journal of Futures Market*, 26, 521-531.

Whaley, R. (1993). Derivatives on Market Volatility: Hedging Tools Long Overdue, *Journal of Derivatives*, 1 (Fall): 71-84.

Whaley, R. (2000). The Investor Fear Gauge: Explication of the CBOE VIX. *Journal of Portfolio Management*, 26, 12-17

Whaley, R. (2009). Understanding the VIX. *Journal of Portfolio Management*, 35, 98-105.

Zhu, S.-P., and Lian, G.-H. (2012). An Analytical Formula for VIX Futures and its Applications. *Journal of Futures Markets*, 32, 166-190.

Table 1. Descriptive Statistics

	Levels		Differences	
	VIX	VXF	Δ VIX	Δ VXF
Mean	25.83	26.43	-4.14e-06	-8.22e-07
Median	22.60	24.11	0.00	0.00
Max.	88.05	69.26	0.197	0.054
Min.	13.30	15.08	-0.212	-0.049
Standard Deviation	11.03	9.17	0.00130	0.000996
Skewness	1.944	1.647	-0.353	0.0415
Kurtosis	7.218	5.687	1785	33.91
ρ_1	1.00	1.00	-0.084	-0.118
ADF	-3.64***	-2.78*	-121.27***	-450.65***

Table 2. Daily VAR Analysis

Panel A: Estimation Results		
	ΔVIX_t	ΔVXF_t
C	-0.0001	-0.0002
ΔVIX_{t-1}	-0.206***	-0.050
ΔVXF_{t-1}	0.149**	0.047
R^2	0.0182	0.0018
Panel B: Granger Causality Tests		
	Causality from ΔVXF to ΔVIX	Causality from ΔVIX to ΔVXF
	4.84**	2.12

Table 3. Intraday VAR Analysis

Panel A: Estimation Results		
	ΔVIX	ΔVXF
constant	-4.80e-06*** [-6.11e-05, 5.82e-05]	-1.07e-06 [-4.60e-05, 5.24e-05]
ϕ_{11}	-0.047*** [-0.486, 0.169]	0.056*** [-0.036, 0.174]
ϕ_{12}	-0.020*** [-0.295, 0.124]	0.033*** [-0.053, 0.127]
ϕ_{13}	0.005* [-0.154, 0.128]	0.017*** [-0.074, 0.096]
φ_{11}	0.172*** [0.015, 0.503]	-0.142*** [-0.289, 0.003]
φ_{12}	0.119*** [-0.010, 0.391]	-0.066*** [-0.166, 0.028]
φ_{13}	0.058*** [-0.033, 0.203]	-0.034*** [-0.108, 0.036]
R^2	0.0926 [0.0057, 0.3147]	0.0412 [0.0075, 0.0954]
Panel B: Granger Causality and Variance Decomposition		
Mean	Causality from ΔVXF to ΔVIX	Causality from ΔVIX to ΔVXF
	102.80	14.31
Perc. exceeding 10% Crit. Val.	92.37%	63.99%
Perc. exceeding 5% Crit. Val.	90.38%	57.15%
Perc. exceeding 1% Crit. Val.	85.45%	42.05%
Variance Decomposition	ΔVIX	ΔVXF
due to ΔVIX	94.53%	3.54%
due to ΔVXF	5.47%	96.46%

Table 4. Causality and Variance Decomposition by Year

	Causality from ΔVXF to ΔVIX	Causality from ΔVIX to ΔVXF	VD ΔVIX due to ΔVXF	VD ΔVXF due to ΔVIX
2008	39.722	25.896	2.62%	4.66%
2009	51.552	17.622	2.77%	2.80%
2010	108.320	12.903	6.08%	3.49%
2011	220.061	7.870	10.04%	3.68%
2012	94.556	7.151	5.86%	3.06%

Table 5. Regression Results for the Causality Ratios

	$\log\left(\frac{GC_{\Delta VXF}}{GC_{\Delta VIX}}\right)_t$	$\log\left(\frac{GC_{\Delta VXF}}{GC_{\Delta VIX}}\right)_t$	$\log\left(\frac{GC_{\Delta VXF}}{GC_{\Delta VIX}}\right)_t$	$\log\left(\frac{GC_{\Delta VXF}}{GC_{\Delta VIX}}\right)_t$	$\log\left(\frac{GC_{\Delta VXF}}{GC_{\Delta VIX}}\right)_t$
α	0.0366 (0.260)	0.0258 (0.18)	-3.592*** (-3.43)	0.00366 (0.259)	-3.398*** (-3.21)
$trend_t$	0.0027*** (12.21)	0.0027*** (12.28)	0.0032*** (12.94)	0.0027*** (12.16)	0.0032*** (12.82)
R_SP_t		-0.105*** (-4.10)			-0.081*** (-03.18)
$\log(VIX)_t$			1.041*** (3.37)		0.983*** (3.155)
$\log\left(\frac{FutVol}{OptVol}\right)_t$				0.0458 (0.43)	
$R^2(adj)$	0.2843	0.2928	0.3140	0.2839	0.3187

Figure 1. Time Series Plot of VIX and VXF

