2019

# The Predictive Power of the VIX Futures Prices on Future Realized Volatility

Siran Zhang

Claremont McKenna College

# The Predictive Power of the VIX Futures Prices on Future Realized Volatility

submitted to
Professor Fan Yu

by
Siran Zhang (Sylvan)

for
Senior Thesis in Financial Economics
Spring Semester 2019
April 29, 2019

*This page is intentionally left blank.*

## ABSTRACT

Many past literatures have examined the predictive power of implied volatility versus that of historical volatility, but they have showed divergent conclusions. One of the major differences among these studies is the methods that they used to obtain implied volatility. The VIX index, introduced in 1993, provides a model-free and directly observable source of implied volatility data. The VIX futures is an actively traded VIX derivative product, and its prices are believed to contain market's expectation about future volatility. By analyzing the relationship between the VIX futures prices and the realized volatilities of the 30-day period that these VIX futures contracts cover, this paper finds that the VIX futures contracts with shorter maturities have predictive power on future realized volatility, but they are upwardly biased estimates. The predictive power, however, decreases as the time to maturity increases. The outstanding VIX futures contracts with the nearest expiration dates outperform GARCH estimates based on historical return data at predicting future realized volatility.

# INTRODUCTION

Everyone wants to predict the future, especially option traders and risk analysts. They want to know what the market in the future will look like, so they can price options more accurately and hedge their risks more efficiently. However, the magical crystal ball that tells everything about the future has not been found yet. Thus, option traders and risk analysts have to rely on estimates to predict the future. Future market volatility is an important thing that they try to predict because it is an essential input for option pricing and a necessary element to for risk forecasting.

The two main choices for predicting future volatility are implied volatility and historical volatility. Implied volatility is the volatility estimate embedded in actively traded options. Because the options are actively traded in the market, the prices of the options reflect the market's expectation of the future. Thus, the volatility estimate implied in these option prices is believed to contain forward looking information about future volatility. Historical volatility estimates are calculated with historical return data. Some examples of historical volatility estimates calculation include moving average of past volatility and GARCH (generalized autoregressive conditional heteroscedasticity) estimates. Because the historical volatility estimate only contains information about the past, many people believe that it is an inferior estimate to the implied volatility estimate.

Gradually, using implied volatility to predict future realized volatility has become a convention. More traders started to quote options with implied volatility and more risk analysts started to use implied volatility to forecast risks. However, does implied volatility contain information about future realized volatility, or is the conventional belief just a false assumption that people failed to recognize? If implied volatility does have predictive power, is its predictive power really superior to that of historical volatility?

Many past studies examined the comparative predicting power of implied volatility and historical volatility, and they have showed divergent conclusions. Canina and Figlewski (1993) studied S&P 100 options and concluded that implied volatility had no correlation with future realized volatility. Flemings (1998), Christensen and Prabhala (1998), and Christensen and Hansen (2002) used the same data, but found that implied volatility contains information about future realized volatility and outperforms historical volatility at predicting future realized volatility. Jorion (1995) and Szakmary et al. (2003) extended the study on implied volatility and historical volatility to asset classes other than stocks, and found that implied volatility is a better forecast than historical volatility. Martens and Zein (2004) used data from different asset classes as well, but they found that the estimate based on high-frequency historical return data can outperform implied volatility at predicting future volatility.

These authors used different methods to obtain implied volatility. Some used the Black-Scholes formula to break down prices into implied volatility. Some used the binomial model to obtain implied volatility. The data selection procedures that they used are different, too. Some only used at-the-money call options, while some included other options trading data as well. The differences in their procedures may partly account for the divergence of their conclusions. The VIX index, introduced in 1993, provides an easily observable source of implied volatility data. Whaley (2008) and Blair et al. (2010) found that the VIX index works well as a predictor of the future volatility. The VIX futures is a common and actively traded derivative product of the VIX index. Nossman and Whilhelmsson (2009) found that VIX futures prices is an upwardly biased estimate of the VIX level in the future. Thus, if the VIX futures can predict the VIX index in the future, and the VIX index can predict future realized volatility, the VIX futures should be able to predict future realized volatility.

This paper seeks to determine whether the prices of the VIX futures have predictive power on future realized volatility. The first hypothesis is VIX futures have predictive power. The rationale underlying the hypothesis is that investors would buy or sell the VIX futures contracts for one of two reasons: either to hedge the risks in their portfolio or to speculate on the future volatility movement. Their expectation of future market volatility is therefore reflected by the trading prices of the VIX future contracts. If they expect the future volatility to be higher, they would bid up the prices of the VIX futures, and if they expect the future volatility to be lower, they would short the VIX futures to drive down the prices. If the market has real predicting power of the future, then the VIX futures prices would be a predictor of future realized volatility.

The second hypothesis is that the VIX futures is an upwardly biased estimate of future volatility. This hypothesis is inspired by Nossman and Whilhelmsson (2009), as they found that VIX futures is an upwardly biased estimate of future VIX level because of the negative risk premium[1]. Moreover, Fleming (1998) found that implied volatility is an upwardly biased estimate of future realized volatility. Combining the findings of the two studies, the VIX futures prices should be upwardly biased estimates of future volatility.

The third hypothesis is that the VIX futures would have better predictive power when it is closer to expiration. The rationale for this hypothesis is that it is more difficult to predict market

[1] Because there is an inverse relationship between volatility level and stock return, many investors choose to long VIX futures to hedge their long positions in the stock market. Thus, the VIX futures contracts work like insurance, and holding these contracts will result in the negative risk premium.

movement too far in the future because there is too much uncertainty. The market today can be very different from the market that is a few months away. Liu (2014) also found that short-term VIX futures are more closely correlated with VIX spot movement.

The fourth hypothesis is that the VIX futures prices are statistically stronger predictor of future realized volatility than historical volatility. The rationale for this hypothesis is that the prices of the VIX futures contracts aggregate market expectation of future market movement. Therefore, they contain forward looking information about future volatility. The rationale is similar to the conventional belief that implied volatility is a better estimate of future volatility than historical volatility.

Previously, there are few studies that use the VIX futures prices as the source of implied volatility data. Several studies used the VIX index as the source of implied volatility data, but the VIX index is an indicative index instead of an actively traded product. My paper contributes to existing literatures by using trading data of actively traded VIX futures contracts to examine the predictive power of implied volatility. Moreover, my paper makes incremental contributions to existing literatures by extending the study to a more recent period, as I used data from 2013 to 2019.

## LITERATURE REVIEW

### *Early Researches on Implied Volatility*

The early research on the forecasting power of implied volatility started in the early 1990s, when the option market had been active for more than a decade. Canina and Figlewski (1993) analyzed the data of call options on the S&P 100 index (also known as the OEX index) from 1983 to 1987, which consisted of 17,606 observations. The S&P 100 options were the most actively traded options in the United States when the study was conducted. The authors concluded from the study that implied volatility had virtually no correlation with future realized volatility. They interpreted the conclusion by saying that there were multiple factors that influenced option supply and demand, such as "liquidity considerations, interaction between the S&P 100 option and the S&P 500 index futures contract, investor tastes for particular payoff patterns, and so on" (Canina and Figlewski, 1993, pg. 677). Such factors are not part of the option pricing models, which generally assume frictionless markets. Thus, these factors generated significant noise, which makes implied volatility a poor predictor of future realized volatility. They also tested the forecasting power of historical volatility. Even though historical volatility showed better predictive power, the authors concluded that neither implied volatility nor historical volatility was a meaningful forecast of future volatility for the S&P 100 options.

Fleming (1998) examined the S&P 100 option trading data, but his study showed different results. He used a statistical technique that explicitly accounted for overlapping observations, and he calculated implied volatility with a binomial tree that incorporated factors such as decisions to pay out dividends and decisions to exercise early. The study concluded that implied volatilities of S&P 100 index options are upwardly biased estimators of future volatilities. However, the bias was not economically significant enough to signal the existence of abnormal trading profits.

Christensen and Prabhala (1998) also conducted research on volatility forecast with the trading data of S&P 100 index options, but they included trading data from November 1983 to May 1995, a much longer time horizon than those of previous studies. In contrast to the previous studies, which suggest either implied volatility has virtually no predictive power or it is a biased and inefficient estimator, Christensen and Prabhala found that implied volatility is an unbiased and efficient forecast of future realized volatility, and it outperformed historical volatility in forecasting. They claimed that the results were different because they used longer time series and non-overlapping data with lower sampling frequency. Their method of data collection was different from that of Canina and Figlewski's study in 1993, which used daily option trading data in a shorter time period. Christensen and Prabhala also suggested that the option trading data before the October 1987 market crash were different from the data after. It might help explain why implied volatility was biased in previous work.

Christensen and Hansen (2002) provided new evidence on the relation between implied volatility and future realized volatility by extending the robustness of the previous results to a more recent period. The data that they used were S&P 100 index option trading data, consistent with previous studies. They constructed implied volatility as a trade-weighted average of implied volatilities from both in-the-money (ITM) and out-of-the-money (OTM) options and both puts and calls. It was different from previous studies as they only considered the information content of call options. Then, Christensen and Hansen ran a horse race between implied volatility and historical volatility to compare their predictive powers. The results underscore their conclusion that implied volatility is an efficient forecast of realized return volatility, and that implied volatility is a better predictor of future volatility than historical volatility.

*Broader Studies on Implied Volatility*

The research on the forecasting power of implied volatility has been extended to other asset classes and other financial markets as well. Jorion (1995) examined the predictive power of implied volatility for foreign currency and compared it to that of historical volatility. He derived implied volatilities from Chicago Mercantile Exchange (CME) options on foreign currency futures covering the German deutsche mark, Japanese yen, and Swiss franc, which were the most active contracts on the CME. The data range from 1985 to 1992, with a frequency of daily observation. His conclusion is that implied volatility outperforms historical volatility in predicting future realized volatility. However, he also noticed that implied volatility appears to be biased volatility forecasts.

Szakmary et al. (2003) studied the predictive power of implied volatility with data from 35 futures options market from eight separate exchanges all over the world. They found that implied volatility outperformed historical volatility as a predictor of the subsequent realized volatility in the underlying futures prices over the remaining life of the option. They also concluded that historical volatility contained no economically significant predictive information beyond what was already incorporated in implied volatility. The study contributes to the literature by showing that the predictive power of implied volatility applies more broadly and it is not just limited to the S&P 100 index options and the United States. The analysis was extended to a very broad array of contracts and exchanges, and showed that the futures options markets in general were efficient.

Martens and Zein (2004) compared the predictive power of historical volatility and implied volatility with a new approach, and they found something different. Their study incorporated trading data from different asset classes, including equity (S&P 500 index), foreign exchange (YEN/USD), and commodity (Sweet Crude Oil). They used a fractional integrated autoregressive model, the predictability of long memory realized volatilities computed from squared high-frequency returns can compete with implied volatilities, and in some instances, outperform implied volatilities.

*The VIX as a Data Source for Implied Volatility*

One reason that may account for the differences in the conclusions of the studies discussed above is the model-dependent nature of implied volatility. Because implied volatility is not directly observable, researchers need to obtain implied volatility by breaking down option prices. However, depending on different types of options and different models to break down option prices, the outcomes can be very different. Therefore, the complexity in obtaining implied volatility data makes it difficult to compare the results from different authors.

The VIX, introduced in 1993, was an index that was intended to provide a benchmark of expected short-term market volatility and to provide an index upon which futures and options contracts on volatility could be written (Whaley, 2008). The VIX is calculated in two steps. First, it takes the current market prices for S&P 500 calls and puts for the first and second month expirations as inputs. Then, it calculates the square root of the risk-neutral expectation of S&P 500 variance over the next 30 calendar days, which delivers the expected volatility in annualized percentage points format (Liu, 2014).

Originally, the VIX was based on the prices of S&P 100 options because at the time, the S&P 100 options were the most actively-traded index options in the United States. The original VIX was calculated only with at-the-money index calls and puts because out-of-the-money options had insufficient liquidity. Over the years since the original VIX was introduced, two major changes have taken place. First, options on the S&P 500 index have become the most actively traded options in the United States. Second, the liquidity for out-of-the-money options have gone up because more people seek to buy these options as insurance for their portfolio. Thus, in 2003, the Chicago Board Options Exchange (CBOE) changed the VIX calculation by using S&P 500 options and including out-of-the-money options. The VIX is believed to be a good gauge of people's expectation of the short-term market volatility. For researches on implied volatility, the VIX provides a good source of implied volatility data that is directly observable.

Thus, the question of the predictive power of implied volatility on future realized volatility can be examined with the VIX data. Blair et al. (2010) used daily observations of the VIX index as implied volatility data and compared its forecasting power with historical volatility. They found that nearly all relevant information of the forecast was provided by the VIX index, which showed that implied volatility had better predictive power. They further included historical volatility of intraday returns to forecast future realized volatility, but the evidence for incremental forecasting information was insignificant.

Robert Whaley (2008), the inventor of VIX, ran an informal test on VIX data and the results showed an affirmative answer to the predictive power of implied volatility. In his test, he computed the 50%, 75%, and 95% expected ranges of S&P 500 rate of return for a one-month period with the level of VIX at the beginning of the month. Then, he computed the actual rate of return of the S&P 500 over the month. His results showed that 34.7% observations fell outside the 50% range, 7.3% fell outside the 75% range, and 1.1% fell outside the 95% range. He concluded that VIX "works reasonably well as a predictor of the expected of stock index movements" (Whaley, 2008, pg. 11).

*The VIX Futures*

The VIX was an indicative index until the CBOE launched VIX futures contracts in May 2004 and VIX options contracts in February 2006. Since then, people can directly trade on the VIX level, and the prices of the VIX derivatives closely reflect investors' expectation of the market volatility. In addition to the VIX futures and options, there have been several other tradeable financial products that link to the VIX index, offering investors with broad market access to trade on volatility. The focus of this essay will be on the VIX futures because they are the simplest VIX derivatives product.

The VIX futures and the VIX spot are closely related. The S&P 500 VIX Short-Term Futures Index measures the return from a rolling long position in the nearest and second nearest VIX futures contracts. “Nearest” means that the contract has the closest expiration date. Liu (2014) found that the S&P 500 VIX Short-Term Futures Index has an 88.15% correlation with the VIX spot. However, she points out that the VIX futures do not tract the VIX spot movement perfectly due to the characteristics of the futures market.

Nossman and Wilhelmsson (2009) studied the VIX futures’ predictability of the VIX level in the future. They mentioned that because the VIX index is negatively correlated with the S&P 500 index, the VIX futures price should contain a negative risk premium, and therefore should be an upward biased estimate of the VIX level in the future. Their results confirmed their hypothesis on the negative risk premium in the VIX futures, and suggested that the VIX futures can predict the VIX index in the future very well.

## DATA

### *Data Collection*

The VIX futures trading data is from the CBOE website. The available trading data for VIX Futures range from 1/2/2013 to present. I use the monthly VIX futures contracts because they have very little overlap in the 30-day period that they cover for expected volatility. For any trading day, there are 9 outstanding monthly VIX futures contracts with different time to maturity. Each monthly VIX futures contract implies the volatility of a 30-day period starting from its expiration date. I select the monthly VIX futures close prices from 1/2/2013 to 2/13/2019 because this range provides an observable *ex post* realized volatility of the 30-day period for any VIX futures contract with the nearest distance to maturity.

To calculate realized volatility, I use daily trading data of S&P 500, which is the underlying asset for the VIX. To match the coverage of all the monthly VIX futures contract that I selected, I use the close price of S&P 500 from 1/2/2013 to 3/15/2019 and calculate the daily return. Then, I calculate the annualized realized volatility of a 30-day period following each trading day with the following formula:

$$\sigma_{realized} = \sqrt{\frac{252}{N} * \sum_{t=1}^{N} (r_t - \bar{r})}$$

where $\sigma_{realized}$ is the annualized realized volatility, $N$ is the number of trading days in the 30-day period following the trading day of observation, 252 is the number of trading days in a calendar year, $r_t$ is the daily return for day t in the period, and $\bar{r}$ is the average daily return of the period. With the formula, I can obtain the annualized realized volatility for each trading day from 1/2/2013 to 2/13/2019, which matches the range of our monthly VIX futures sampling.

Table 1 presents the descriptive statistics of the monthly VIX futures prices and the annualized realized volatility for all 1540 trading days within the selected time frame. Note that the $9^{th}$ nearest expiring VIX futures contract only has 1176 observations, which is different from all other contracts with 1540 observations. The lower number of observations is due to the lack of trading volume of the contracts with the longest maturity. On some trading days there are no trades at all, and thus these days have provided no observation of prices.

Starting with the means for the VIX futures prices with different distance to maturity, I find that the mean price increase when the time to maturity increases. It suggests that the VIX futures contracts with longer maturity dates are more expensive than those with shorter maturity dates. It also suggests that the prices of VIX futures are falling through time. These two characteristics show that the market for the VIX futures is in contango, which describes the market where future prices are positively correlated to their times to maturity.

Table 1:
Descriptive Statistics

| | Observations | Mean | Standard Error | Standard Deviation | Kurtosis | Skewness | Avg. Daily Volume |
|---|---|---|---|---|---|---|---|
| *Annualized Realized Volatility* | 1540 | 11.8036 | 0.1414 | 5.5500 | 1.2487 | 1.2383 | |
| *1st Nearest Expiring VIX Futures Contract* | 1540 | 15.5116 | 0.0823 | 3.2277 | 2.1418 | 1.2767 | 99241.2182 |
| *2nd Nearest Expiring VIX Futures Contract* | 1540 | 16.2647 | 0.0646 | 2.5361 | 0.9675 | 0.7956 | 80622.3708 |
| *3rd Nearest Expiring VIX Futures Contract* | 1540 | 16.8893 | 0.0575 | 2.2554 | 0.7017 | 0.5754 | 24722.1922 |
| *4th Nearest Expiring VIX Futures Contract* | 1540 | 17.3782 | 0.0520 | 2.0406 | 0.2165 | 0.4084 | 12196.0909 |
| *5th Nearest Expiring VIX Futures Contract* | 1540 | 17.7909 | 0.0486 | 1.9083 | 0.1238 | 0.3268 | 7225.5617 |
| *6th Nearest Expiring VIX Futures Contract* | 1540 | 18.1560 | 0.0458 | 1.7960 | 0.0330 | 0.2713 | 4289.4409 |
| *7th Nearest Expiring VIX Futures Contract* | 1540 | 18.5101 | 0.0443 | 1.7374 | -0.1405 | 0.2578 | 2495.3494 |
| *8th Nearest Expiring VIX Futures Contract* | 1540 | 18.7876 | 0.0421 | 1.6538 | -0.1661 | 0.2897 | 696.1006 |
| *9th Nearest Expiring VIX Futures Contract* | 1176 | 18.8265 | 0.0507 | 1.7396 | 22.3225 | -1.8632 | 165.6735 |

One explanation of this phenomenon is the inverse relationship between stock returns and market volatility. When market returns are high, market volatility tends to be low. When market volatility rises, the market tends to fall. Thus, people would long VIX products to hedge their long positions in the stock market. Taking a long position in the VIX futures will result in negative risk premiums because they work like insurance on the stock investments. This finding is consistent with the conclusion of the Nossman and Wilhelmsson (2009) study about the VIX futures.

The daily trading volume data of the VIX futures contracts provides insights about the behavior of participants in this market. We can see that for the nearest and second nearest VIX futures contracts, the average daily trading volume are 99241.2182 and 80622.3708 respectively, which suggest that they are very liquid financial products. Then, the daily trading volume dropped significantly for longer maturity contracts. The average daily trading volume for the $5^{th}$ nearest VIX futures contracts is below 10000, and the average daily trading volume for the $9^{th}$ nearest VIX futures contracts is less than 200. One explanation is that the speculators only trade the VIX futures contracts with shorter distance to maturity because their expectation of market movement cannot go too far in the future. They have more relevant information for the near future than for a few months from the present day. Investors who trade longer maturity VIX futures contracts may have hedging mandates for their portfolio, meaning that they have to purchase a certain amount of the VIX futures contracts with specific maturities no matter what the prices are.

*Testing the Predictive Power of the VIX Futures*

The expiration dates and settlement values of the VIX futures contracts are complicated. The expiration date for a VIX futures contract is the Wednesday that is 30 days prior to the third Friday of the calendar month immediately following the month in which the contract expires.

These expiration dates match with the start dates of 30-day periods before the constituent S&P 500 options expire. Within the selected time frame for our study, the only exception is the March 2014 VIX futures contract, which has an irregular expiration date. The March 2014 contract expired on Tuesday, March 18, 2014 because the April 2019 S&P 500 option expiration is on Thursday, April 19, 2014 due to Good Friday holiday. The settlement prices of VIX futures contracts are calculated with prices of the aforementioned constituent S&P 500 options. For example, for the VIX futures contract that expired on 11/21/2018, its settlement value was calculated using S&P 500 options that expired 30 days later on 12/21/2018.

Table 2:
Corresponding realized volatility of selected VIX futures contracts

| | *Corresponding realized volatility of selected VIX futures contracts* |
|---|---|
| Mean | 11.94188784 |
| Standard Error | 0.659770119 |
| Standard Deviation | 5.675557161 |
| Kurtosis | 1.741533167 |
| Skewness | 1.281367916 |
| Range | 28.1779611 |
| Minimum | 3.713390979 |
| Maximum | 31.89135208 |
| Observations | 74 |

After calculations, I am able to obtain realized volatility of 30-day periods immediately following the expiration dates of 74 VIX futures contracts. The first contract is the January 2013 contract, and the last contract is the February 2019 contract. An exhaustive list of the 74 VIX futures contracts can be found in the Appendix section. Table 2 shows the summary statistics for these realized volatility calculations, which is very similar to that of the annualized realized volatility in Table 1. The only difference is that the standard error in Table 2 is larger because there are fewer observations. However, the realized volatility calculations on the 30-day period basis have little overlap, which make the data more independent from each other.

To test the predictive power of the VIX futures on future realized volatility, I will use the following linear model for regression

$$\sigma_{realized}(T) = \alpha + \beta * P_{VIX\ Futures}(t, T) + \varepsilon$$

where $P_{VIX\ Futures}(t, T)$ denotes the close price at $t$ of a VIX futures contract that expires at $T$, and $\sigma_{realized}(T)$ is the annualized realized volatility in percentage points for the 30-day period immediately following $T$.

For the predictive power to be significant for the VIX futures, we want $\beta$ to be nonzero under statistical significance. If the price of VIX futures is an unbiased estimator of future realized volatility, then $\alpha = 0$ and $\beta = 1$. However, my hypothesis is that the price of VIX futures to be an upwardly biased estimator, so we would want a combination of $\alpha$ and $\beta$ to make the VIX futures prices consistently higher than the future realized volatility.

To test whether shorter maturity contracts have more predictive power than longer maturity contracts, we observe the close prices of all outstanding VIX futures contracts on the trading days right before the expiration dates of the nearest contracts. For example, on Tuesday, January 16, 2018, which is one trading day ahead of the expiration date of the January 2018 VIX futures contract, we observe the close prices of the nearest January 2018 contract, the second nearest February 2018 contract, and all the way to the farthest September 2018 contract. Then, I regress the realized volatility calculations for each VIX futures contract on their observed prices.

Table 3:
Predictive power of contracts with different distance to maturity

| | **Observations** | **α** | **β** | **Standard Error** | **t Stat** | **P-value** | **Lower 95%** | **Upper 95%** | **R Square** |
|---|---|---|---|---|---|---|---|---|---|
| *1st Nearest Expiring VIX Futures Contract* | 74 | -0.9466 | 0.8478 | 0.1472 | 5.7596 | 0.0000 | 0.5543 | 1.1412 | 0.3154 |
| *2nd Nearest Expiring VIX Futures Contract* | 73 | 4.5717 | 0.4666 | 0.2372 | 1.9673 | 0.0531 | -0.0063 | 0.9394 | 0.0517 |
| *3rd Nearest Expiring VIX Futures Contract* | 72 | 11.6529 | 0.0214 | 0.2766 | 0.0774 | 0.9386 | -0.5302 | 0.5730 | 0.0001 |
| *4th Nearest Expiring VIX Futures Contract* | 71 | 16.0413 | -0.2377 | 0.3131 | -0.7591 | 0.4504 | -0.8622 | 0.3869 | 0.0083 |
| *5th Nearest Expiring VIX Futures Contract* | 70 | 16.5180 | -0.2568 | 0.3389 | -0.7576 | 0.4513 | -0.9330 | 0.4195 | 0.0084 |
| *6th Nearest Expiring VIX Futures Contract* | 69 | 20.8077 | -0.4931 | 0.3562 | -1.3842 | 0.1709 | -1.2042 | 0.2180 | 0.0278 |
| *7th Nearest Expiring VIX Futures Contract* | 68 | 22.8156 | -0.5920 | 0.3785 | -1.5641 | 0.1226 | -1.3476 | 0.1637 | 0.0357 |
| *8th Nearest Expiring VIX Futures Contract* | 67 | 18.8092 | -0.3640 | 0.4042 | -0.9004 | 0.3712 | -1.1713 | 0.4433 | 0.0123 |
| *9th Nearest Expiring VIX Futures Contract* | 66 | 22.5897 | -0.5552 | 0.4241 | -1.3090 | 0.1952 | -1.4025 | 0.2921 | 0.0261 |

The results of the regression are in Table 3. They are consistent with the hypothesis that the VIX futures contract with shorter distance to maturity has better predictive power on realized volatility. For the nearest VIX futures contract, the $\alpha$ is -0.9466 and the $\beta$ is 0.8478, which shows that its price is a good estimate of future realized volatility. The close-to-zero p-value suggests the result is statistically significant. The R square is 0.3154, which shows that over 31% of the change in future realized volatility can be explained by the price of the nearest expiring VIX futures contract.

For the second nearest VIX futures contract, there is still some marginal statistical significance in its predicting power. The p-value is just slightly above 0.05. The R square is significant smaller than that of the nearest contract, suggesting that very little change in the future realized volatility can be explained by the price of the second nearest VIX futures contract. The rest of the results have no significance at all, which show that longer maturity VIX futures contracts have virtually no predicting power on the future realized volatility.

Given the results in Table 3, I focus on examining the predictive power of the nearest VIX futures contract. I treat the different contracts as cross-sectional data, and group all the observations for the nearest VIX futures contracts by the number of trading days to their maturity. For example, if we are examining the predicting power of the nearest VIX future contract with 5 trading days to maturity, we would observe the close price of January 2013 contract on 1/9/2013, February 2013 contract on 2/6/2013, March 2013 contract on 3/13/2013, and all the way to February 2019 contract on 2/6/2019. With such grouping method, the prices of all 74 nearest contracts have equal distance to expiration in every regression.

Table 4
Predictive power of the nearest VIX futures contract with different distance to maturity

| Days Before Expiration | Observations | α | β | Standard Error | t Stat | P-value | Lower 95% | Upper 95% | R Square |
|---|---|---|---|---|---|---|---|---|---|
| *1* | 74 | -0.9466 | 0.8478 | 0.1472 | 5.7596 | 0.0000 | 0.5543 | 1.1412 | 0.3154 |
| *2* | 74 | 0.3171 | 0.7567 | 0.1482 | 5.1057 | 0.0000 | 0.4613 | 1.0521 | 0.2658 |
| *3* | 74 | 2.1573 | 0.6340 | 0.1583 | 4.0055 | 0.0001 | 0.3185 | 0.9496 | 0.1822 |
| *4* | 74 | 2.6965 | 0.5963 | 0.1659 | 3.5933 | 0.0006 | 0.2655 | 0.9271 | 0.1521 |
| *5* | 74 | 2.6441 | 0.5984 | 0.1749 | 3.4217 | 0.0010 | 0.2498 | 0.9470 | 0.1399 |
| *6* | 74 | 2.6994 | 0.5989 | 0.1855 | 3.2292 | 0.0019 | 0.2292 | 0.9686 | 0.1265 |
| *7* | 74 | 5.3064 | 0.4272 | 0.1611 | 2.6524 | 0.0098 | 0.1061 | 0.7483 | 0.0890 |
| *8* | 74 | 3.2592 | 0.5658 | 0.2101 | 2.6935 | 0.0088 | 0.1471 | 0.9846 | 0.0915 |
| *9* | 74 | 3.7875 | 0.5260 | 0.2130 | 2.4697 | 0.0159 | 0.1014 | 0.9505 | 0.0781 |
| *10* | 74 | 4.6529 | 0.4690 | 0.2045 | 2.2938 | 0.0247 | 0.0614 | 0.8766 | 0.0681 |
| *11* | 73 | 5.0565 | 0.4464 | 0.2223 | 2.0083 | 0.0484 | 0.0032 | 0.8897 | 0.0538 |
| *12* | 73 | 4.6255 | 0.4744 | 0.2209 | 2.1471 | 0.0352 | 0.0338 | 0.9150 | 0.0610 |
| *13* | 73 | 3.9529 | 0.5156 | 0.2212 | 2.3309 | 0.0226 | 0.0745 | 0.9567 | 0.0711 |
| *14* | 73 | 5.1777 | 0.4354 | 0.2213 | 1.9675 | 0.0530 | -0.0059 | 0.8766 | 0.0517 |
| *15* | 73 | 4.9171 | 0.4514 | 0.2117 | 2.1328 | 0.0364 | 0.0294 | 0.8735 | 0.0602 |
| *16* | 73 | 5.8604 | 0.3892 | 0.2047 | 1.9012 | 0.0613 | -0.0190 | 0.7974 | 0.0484 |
| *17* | 71 | 2.3933 | 0.6023 | 0.2190 | 2.7507 | 0.0076 | 0.1655 | 1.0392 | 0.0988 |
| *18* | 69 | 2.5743 | 0.5862 | 0.2346 | 2.4983 | 0.0149 | 0.1179 | 1.0545 | 0.0852 |

Table 4 presents the results of the regressions. We can see that all the regressions show statistical significance except for 14 and 16 trading days as the distance to maturity, which has p-values of 0.0530 and 0.0613 respectively. These regressions are still marginally significant. The R square statistic declines as the time to maturity increases, which show that higher predictive power for the VIX future contracts with shorter maturity holds true within a 30-day period.

*VIX Futures vs. Historical Volatility*

To construct a forecast of future volatility with historical data, I use the generalized autoregressive conditional heteroscedasticity (GARCH) model to account for volatility clustering and time-varying volatility in time series financial data. The GARCH model is an extension of an autoregressive conditional heteroscedasticity (ARCH) model. The basic ARCH model consists of two equations: the mean equation and the variance equation:

Mean equation: $y_t = \beta + e_t$

Variance equation: $h_t = \alpha + \alpha_1 * e_{t-1}^2$

The mean equation is a linear regression function that contains a constant and some explanatory variables. It describes the behavior of the mean of the time series data. In the basic model, the mean function only contains an intercept, $\beta$. The variance equation the behavior of the error variance. The error of the regression is normal and heteroskedastic. The variance of the current period's error, $h_t$, depends on information of the previous period, $e_{t-1}$.

The GARCH model adds lags of the variance to the ARCH model. A GARCH (1,1) model would have the variance equation of

$$h_t = \delta + \alpha_1 * e_{t-1}^2 + \beta_1 * h_{t-1}$$

where $\delta + \alpha_1 * e_{t-1}^2$ is the standard ARCH variance equation, and $\beta_1 * h_{t-1}$ is one lag of the error variance itself. Because the mean equation only contains an intercept of the mean, the error variance equation essentially provides a prediction model.

Table 5:
The Comparison between GARCH estimates and VIX futures prices

| | **β (VIX Futures)** | **β (GARCH Vol)** | **R Square** |
|---|---|---|---|
| 1 Day Before Expiration | 0.8478*** | | 0.3154 |
| | | 0.5845*** | 0.1628 |
| | 1.0959*** | -0.2874 | 0.3277 |
| 2 Day Before Expiration | 0.7567*** | | 0.2658 |
| | | 0.6855*** | 0.1623 |
| | 0.9218*** | -0.2248 | 0.2706 |
| 3 Day Before Expiration | 0.6340*** | | 0.1822 |
| | | 0.5530** | 0.0794 |
| | 0.7231*** | -0.1581 | 0.1851 |

*** 1% significance level
** 5% significance level
* 10% significance level

To predict future volatility based on historical return data and the GARCH (1,1) model, I use a five-year period of S&P 500 trading data that is n-days ahead of the expiration date of each contract. Then, I construct the variance equation to generate out-of-sample prediction of the n-days ahead daily variance. Then, I multiply the daily variance by 252 and take the square root of the product to obtain the n-days ahead annualized GARCH volatility prediction. For example, the Jan 2013 contract covers the 30-day period following 1/16/2013, and for one-day ahead prediction, I used the five-year trading data from 1/15/2008 to 1/15/2013 to construct the out-of-sample prediction. To generate two days ahead prediction, I use the data from 1/14/2008 to 1/14/2013 to construct the out-of-sample prediction. The GARCH estimates of future volatility are directly compared with the VIX futures prices of the nearest contract that are observed n-days ahead of the expiration date.

Table 5 presents the results of the regressions of one-days ahead to three-days ahead GARCH volatility and VIX futures prices as independent variables and future realized volatility as dependent variables. The results show that when GARCH volatility is the only variable in the regression, there is statistical significance. It suggests that the GARCH volatility has predictive power of future realized volatility. The values of R Square are 0.1628, 0.1623, and 0.0794 respectively for one day ahead, two days ahead, and three days ahead, respectively. The R square of GARCH volatility is lower than that of VIX futures.

When both the GARCH volatility estimate and the VIX futures prices are included as independent variables to explain the change in the future realized volatility, the coefficient of the VIX futures clearly dominates that of the GARCH volatility. The coefficients of the GARCH volatility are not statistically significant in the regressions as well. It suggests that the GARCH volatility estimate has no additional information than what is already included in the VIX futures prices. It is consistent with my hypothesis that VIX futures prices have more predictive power than historical volatility.

*Limitations of the Study*

There are a few caveats to note about the data and the methods that are used for my study. First, the observations of VIX futures prices are between January 2013 and February

2019. The time frame does not cover the period of the Great Financial Crisis. The rationale for the selection of the time frame is to avoid sharp changes in volatility. The study does not include trading data for the period before the Great Financial Crisis as well. If the range of data selection was extended to when the VIX futures were initially available, the predictive power of the VIX futures could be different before and after the Great Financial Crisis.

Second, for the GARCH estimation, I used a five-year period of S&P 500 return data to construct the prediction model. The reason for that is to stay consistent with most studies that employ the GARCH model to predict stock market volatility. However, when using a five-year period for GARCH forecasting, I had to include S&P 500 return data during the 2008-2009 Great Financial Crisis for predicting the realized volatility in 2013 and 2014. The high volatility of stock market returns during the Great Financial Crisis might distort the GARCH model. However, the level of distortion should be insignificant. For example, for the realized volatility of January 2013, the GARCH prediction contains information of the five-year period from January 2008 to January 2013. The error variance of the one-day ahead out-of-sample prediction is 0.0000691, which corresponds to an annualized GARCH volatility prediction of 13.1959%. If I exclude the data in 2008 and 2009, and only use the data of the three-year period from January 2010 to January 2013, the error variance of the one-day ahead out-of-sample prediction is 0.0000634, which corresponds to an annualized GARCH volatility prediction of 12.6399%. The difference in the volatility prediction is around 4.4%, which is not very significant.

Moreover, in the study, I used monthly VIX futures data, but there are also weekly VIX futures data that is available. The weekly VIX futures contracts expire on Wednesdays when no monthly VIX futures contracts expire. I did not use the weekly VIX futures data because the periods of VIX that they cover overlap with each other, which makes them not independent data from each other. However, the weekly VIX futures contracts can be arranged in ways that they do not overlap with each other. For example, one can take all the VIX futures contracts that expire on the first Wednesday of a month, and test their predictive power of the future realized volatility. With such arrangement, the data become independent and there's no overlapping issue. The results of the study with weekly VIX futures contracts should yield the same results as the study that use monthly VIX futures contracts, but it will add more observations to the regressions to make the results more robust.

## CONCLUSION

In this paper, I test four hypotheses. The first hypothesis is that the prices of the VIX futures has predictive power on future realized volatility. This hypothesis is tested by the regressions on the future realized volatility and VIX futures prices. I find that the nearest expiring VIX futures contract prices have statistically significant predictive power for the future realized volatility.

The second hypothesis is that the price of the VIX futures is an upwardly biased estimator of future realized volatility. The regression results of the nearest VIX futures contracts show positive values of constants and coefficients that are less than unity. The combinations of coefficients and constants cannot directly tell whether the nearest VIX futures prices are upwardly biased estimators of future realized volatility. However, we can look at the mean of the nearest VIX futures with different remaining time to maturity. Figure 1 shows the mean prices of the nearest VIX futures prices with respect to the number of days to maturity. All the prices within 18 trading days ahead of expiration have statistical significant predictive power of the future realized volatility. We can see that the prices of the nearest VIX futures prices are consistently higher than the mean of the realized volatility, which suggests that the VIX futures prices are upwardly biased estimators of future realized volatility. Thus, the results are consistent with the second hypothesis.

Figure 1:
Nearest VIX futures mean prices vs. days to maturity

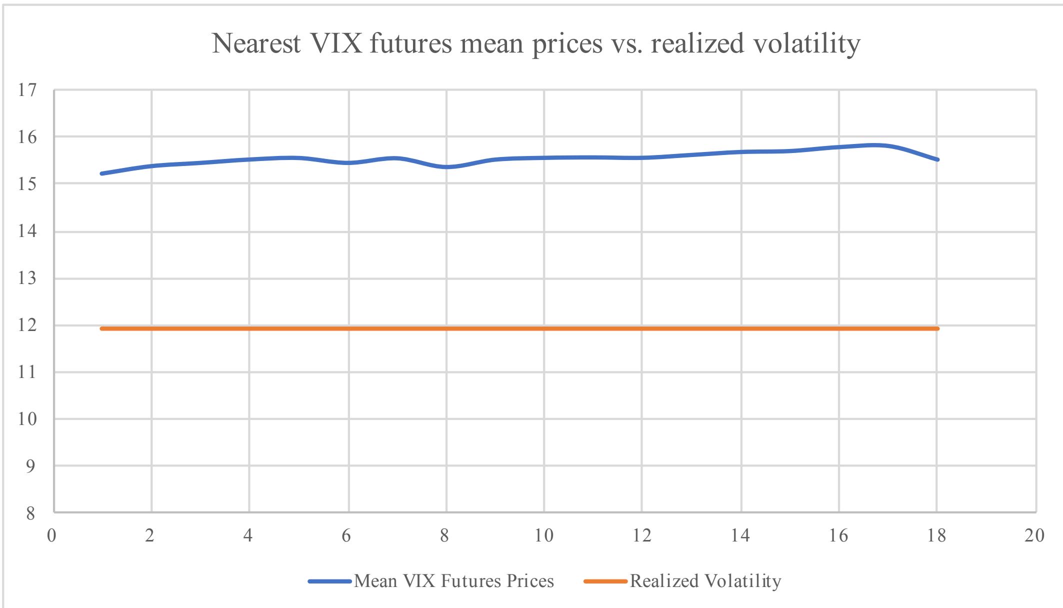


The third hypothesis is that the longer maturity VIX futures have less predictive power of future volatility than the shorter maturity VIX futures. Based on the regressions of outstanding VIX futures prices with different expiration dates, I find that only the nearest VIX futures contract shows statistical significance in terms of predicting future volatility. The second nearest VIX futures contract has some marginal statistical significance in predicting future volatility. The VIX futures with longer maturity have no significance at all.

Even with the same expiration dates, the predictive power is higher when the number of days ahead of expiration is lower. In the analysis that examines the predictive power of the nearest contract with different numbers of days ahead of maturity, I find that the statistical significance and R square decreased when the remaining life of the contract is longer. Therefore, the third hypothesis is validated with both VIX futures with different expiration dates and VIX futures with same expiration dates but different remaining life.

The fourth hypothesis is that the VIX futures prices are statistically stronger predictors of future volatility than historical volatility. Based on the comparison between the predictive power of the VIX futures prices and that of historical volatility, I show that the VIX futures dominates historical volatility in term of predictive power of future volatility. The historical volatility contains nearly no additional information other than the information content of the VIX futures about future volatility.

Appendix I: The 74 VIX futures contracts and their corresponding realized volatility

| Contract Name | Covering Period | Corresponding Realized Volatility | Contract Name | Covering Period | Corresponding Realized Volatility |
|---|---|---|---|---|---|
| *CFE_F13_VX* | Jan-13 | 7.7868 | *CFE_G16_VX* | Feb-16 | 13.6639 |
| *CFE_G13_VX* | Feb-13 | 11.3476 | *CFE_H16_VX* | Mar-16 | 9.9825 |
| *CFE_H13_VX* | Mar-13 | 14.5644 | *CFE_J16_VX* | Apr-16 | 10.1460 |
| *CFE_J13_VX* | Apr-13 | 9.2986 | *CFE_K16_VX* | May-16 | 8.2289 |
| *CFE_K13_VX* | May-13 | 16.3073 | *CFE_M16_VX* | Jun-16 | 19.3108 |
| *CFE_M13_VX* | Jun-13 | 13.3108 | *CFE_N16_VX* | Jul-16 | 5.5304 |
| *CFE_N13_VX* | Jul-13 | 8.1196 | *CFE_Q16_VX* | Aug-16 | 12.5343 |
| *CFE_Q13_VX* | Aug-13 | 10.1783 | *CFE_U16_VX* | Sep-16 | 8.9658 |
| *CFE_U13_VX* | Sep-13 | 13.1117 | *CFE_V16_VX* | Oct-16 | 10.2742 |
| *CFE_V13_VX* | Oct-13 | 8.9767 | *CFE_X16_VX* | Nov-16 | 7.7361 |
| *CFE_X13_VX* | Nov-13 | 9.7124 | *CFE_Z16_VX* | Dec-16 | 6.3024 |
| *CFE_Z13_VX* | Dec-13 | 8.4058 | *CFE_F17_VX* | Jan-17 | 5.9936 |
| *CFE_F14_VX* | Jan-14 | 15.9926 | *CFE_G17_VX* | Feb-17 | 6.9866 |
| *CFE_G14_VX* | Feb-14 | 9.9109 | *CFE_H17_VX* | Mar-17 | 6.1694 |
| *CFE_H14_VX* | Mar-14 | 13.3196 | *CFE_J17_VX* | Apr-17 | 8.8621 |
| *CFE_J14_VX* | Apr-14 | 7.8112 | *CFE_K17_VX* | May-17 | 4.7561 |
| *CFE_K14_VX* | May-14 | 5.4044 | *CFE_M17_VX* | Jun-17 | 7.5841 |
| *CFE_M14_VX* | Jun-14 | 8.2530 | *CFE_N17_VX* | Jul-17 | 8.3085 |
| *CFE_N14_VX* | Jul-14 | 11.5280 | *CFE_Q17_VX* | Aug-17 | 8.7991 |
| *CFE_Q14_VX* | Aug-14 | 5.6850 | *CFE_U17_VX* | Sep-17 | 3.7134 |
| *CFE_U14_VX* | Sep-14 | 16.5404 | *CFE_V17_VX* | Oct-17 | 5.8362 |
| *CFE_V14_VX* | Oct-14 | 7.2341 | *CFE_X17_VX* | Nov-17 | 6.8192 |
| *CFE_X14_VX* | Nov-14 | 15.3995 | *CFE_Z17_VX* | Dec-17 | 6.6328 |
| *CFE_Z14_VX* | Dec-14 | 16.5577 | *CFE_F18_VX* | Jan-18 | 24.4698 |
| *CFE_F15_VX* | Jan-15 | 14.2624 | *CFE_G18_VX* | Feb-18 | 14.2141 |
| *CFE_G15_VX* | Feb-15 | 12.4235 | *CFE_H18_VX* | Mar-18 | 23.2806 |
| *CFE_H15_VX* | Mar-15 | 10.6589 | *CFE_J18_VX* | Apr-18 | 10.7382 |
| *CFE_J15_VX* | Apr-15 | 11.0944 | *CFE_K18_VX* | May-18 | 9.1380 |
| *CFE_K15_VX* | May-15 | 9.7103 | *CFE_M18_VX* | Jun-18 | 9.9015 |
| *CFE_M15_VX* | Jun-15 | 13.3012 | *CFE_N18_VX* | Jul-18 | 8.0228 |
| *CFE_N15_VX* | Jul-15 | 15.8811 | *CFE_Q18_VX* | Aug-18 | 6.3831 |
| *CFE_Q15_VX* | Aug-15 | 31.8914 | *CFE_U18_VX* | Sep-18 | 17.3613 |
| *CFE_U15_VX* | Sep-15 | 17.2271 | *CFE_V18_VX* | Oct-18 | 20.6753 |
| *CFE_V15_VX* | Oct-15 | 13.8993 | *CFE_X18_VX* | Nov-18 | 22.8626 |
| *CFE_X15_VX* | Nov-15 | 16.8770 | *CFE_Z18_VX* | Dec-18 | 28.5519 |
| *CFE_Z15_VX* | Dec-15 | 20.4642 | *CFE_F19_VX* | Jan-19 | 12.2062 |
| *CFE_F16_VX* | Jan-16 | 21.3366 | *CFE_G19_VX* | Feb-19 | 8.9738 |

**REFERENCE**


Blair, Bevan J, et al. "Forecasting S&P 100 Volatility: The Incremental Information Content of Implied Volatilities and High-Frequency Index Returns." *Handbook of Quantitative Finance and Risk Management*, 2010, pp. 1333–1344., doi:10.1007/978-0-387-77117-5_88.

Canina, L., and S. Figlewski. "The Informational Content of Implied Volatility." *Review of Financial Studies*, vol. 6, no. 3, 1993, pp. 659–681., doi:10.1093/rfs/6.3.659.

Christensen, B.J., and N.R. Prabhala. "The Relation between Implied and Realized Volatility." *Journal of Financial Economics*, 1998, pp. 125–150.

Christensen, Bent Jesper, and Charlotte Strunk Hansen. "New Evidence on the Implied-Realized Volatility Relation." *The European Journal of Finance*, vol. 8, no. 2, 2002, pp. 187–205., doi:10.1080/13518470110071209.

Fleming, Jeff. "The Quality of Market Volatility Forecasts Implied by S&P 100 Index Option Prices." *Journal of Empirical Finance*, vol. 5, no. 4, 1998, pp. 317–345., doi:10.1016/s0927-5398(98)00002-4.

Jorion, Philippe. "Predicting Volatility in the Foreign Exchange Market." *The Journal of Finance*, vol. 50, no. 2, 1995, pp. 507–528., doi:10.1111/j.1540-6261.1995.tb04793.x.

Liu, Berlinda. *Identifying the Differences Between VIX Spot and Futures*. 2014, us.spindices.com/documents/education/practice-essentials-difference-between-vix-spot-futures.pdf.

Martens, Martin P.e., and Jason Zein. "Predicting Financial Volatility: High-Frequency Time-Series Forecasts Vis-a-Vis Implied Volatility." *SSRN Electronic Journal*, 2002, doi:10.2139/ssrn.301382.

Nossman, Marcus, and Anders Wilhelmsson. "Is the VIX Futures Market Able to Predict the VIX Index? A Test of the Expectation Hypothesis." *The Journal of Alternative Investments*, vol. 12, no. 2, 2009, pp. 54–67., doi:10.3905/jai.2009.12.2.054.

Szakmary, Andrew, et al. "The Predictive Power of Implied Volatility: Evidence from 35 Futures Markets." *Journal of Banking & Finance*, vol. 27, no. 11, 2003, pp. 2151–2175., doi:10.1016/s0378-4266(02)00323-0.

Whaley, Robert E. "Understanding VIX." *SSRN Electronic Journal*, 2008, doi:10.2139/ssrn.1296743.