

2011

# Quantifying the Variance Risk Premium in VIX Options

Reed M. Hogan  
*Claremont McKenna College*

---

## Recommended Citation

Hogan, Reed M., "Quantifying the Variance Risk Premium in VIX Options" (2011). *CMC Senior Theses*. Paper 147.  
[http://scholarship.claremont.edu/cmc\\_theses/147](http://scholarship.claremont.edu/cmc_theses/147)

This Open Access Senior Thesis is brought to you by Scholarship@Claremont. It has been accepted for inclusion in this collection by an authorized administrator. For more information, please contact [scholarship@cuc.claremont.edu](mailto:scholarship@cuc.claremont.edu).

**CLAREMONT McKENNA COLLEGE**

**QUANTIFYING THE VARIANCE RISK PREMIUM  
IN VIX OPTIONS**

SUBMITTED TO

PROFESSOR AMIR BARNEA

AND

DEAN GREGORY HESS

BY

REED McCARTY HOGAN

FOR

SENIOR THESIS

SPRING 2011

04/25/2011

## **Abstract**

This thesis uses synthetically created variance swaps on VIX futures to quantify the variance risk premium in VIX options. The results of this methodology suggest that the average premium is -3.26%, meaning that the realized variance on VIX futures is on average less than the variance implied by the swap rate. This premium does not vary with time or the level of the swap rate as much as premiums in other asset classes. A negative risk premium should mean that VIX option strategies that are net credit should be profitable. This thesis tests two simple net credit strategies with puts and calls, and finds that the call strategy is profitable while the put strategy is not.

## Table of Contents

Acknowledgements .....	4
1. Introduction .....	5
2. Literature Review .....	6
2.1 Importance of measure VRP in VIX options .....	9
3. Methodology .....	10
3.1 Swap Rate .....	11
3.2 Realized Variance .....	13
3.3 Variance Risk Premium .....	14
4. Data and Results .....	14
4.1 Summary of Results .....	15
4.2 Is the premium time-varying? .....	17
5. Trade Strategy .....	18
6. Conclusion .....	20
Bibliography .....	22

## Acknowledgements

This thesis was made possible by the guidance and support of many people. I would first like to thank my thesis advisor, Professor Amir Barnea, who advised me throughout the duration of this project. I would also like to express my gratitude to Bill Luby, founder of the blog *Vix and More*, and Jared Woodard, principal of *Condor Options*. Their writing first interested me in the volatility asset class, and their readiness to respond to questions and offer assistance is not only admirable, but also made this project possible. I would also like to extend special thanks to Peter Carr, who offered guidance with regard to the implementation of the methodology used in this thesis. Lastly, I would like to thank the many family and friends who read through countless drafts, listened to new ideas, and supported me throughout the entire process.

## 1. Introduction

The VIX index is a measure of the 30-day forward looking implied volatility for the S&P 500, and is an important tool to gauge perceived short-term risk in the stock market. Until recently, this index was merely a reference tool for investors rather than an actively traded asset. However, the index's strong negative correlation with the S&P 500 made it a very attractive hedging tool, and the increasing demand to trade volatility in the market led to the rapid development of the volatility asset class. In March of 2004, the Chicago Board Options Exchange (CBOE) began listing VIX futures contracts, and in February of 2006 VIX options were introduced. This paper uses VIX options to synthetically replicate a variance swap on VIX futures, which will quantify the average difference between realized and implied volatility in the VIX options. While this difference – referred to as the variance risk premium – has been measured for other asset classes, to this writer's knowledge it has not been measured in VIX options in an academic context.

The average difference between realized and implied volatility for many assets is typically negative, meaning that on average volatility purchasers realize less volatility than the volatility implied by the contract price. This explains why the phenomenon is usually referred to as a "premium". It is important for option market participants to understand the sign of this premium and also some arguments for its existence. One common explanation is the negative correlation between volatility and returns, which makes volatility assets powerful portfolio risk management tools. Option buyers should be willing to pay a premium for these contracts if the volatility they purchase serves to reduce their overall portfolio variance. This explanation relates the premium to

purchasing insurance, because buyers are paying an additional cost to manage overall portfolio risk. Another explanation for the existence of this premium is the skewness of returns on variance contracts. Volatility of returns is mean reverting, but prone to very large positive spikes. This means that a long position in a variance swap typically loses smaller amounts of money, but is very profitable when volatility spikes. For the sellers of variance swaps, regular premiums compensate them for the risk of an infrequent but very large loss in times of market uncertainty.

## 2. Literature Review

Different methods have been used to analyze this premium in options. Bakshi and Kapadia (2003) used delta-hedged S&P 500 options to examine the sign of the premium<sup>1</sup>. By purchasing options and delta-hedging with the underlying security, one can effectively hedge out all risk other than volatility risk, so the payoff from these positions determines whether the realized volatility was greater or less than the implied volatility of the contract. These payoffs therefore can be used as a measure of the direction of the premium. They find that a strategy of delta-hedging a long position in S&P 500 options underperforms zero, implying that option prices in general include a premium over the expected volatility during the life of the contract. They also find that the return of this strategy decreases when the options are further out of the money and when volatility is lower. While this method successfully argues that the premium is negative, it cannot quantify the premium. Each option contract has a different vega – a

---

<sup>1</sup> Bakshi, Gurdip, and Nikunj Kapadia. "Delta-Hedged Gains and the Negative Market Volatility Risk Premium." *Review of Financial Studies* 16.2 (2003): 528-66

measure of the sensitivity of option prices to changes in volatility – which would affect the size of the premium for different option contracts. Therefore, the payoff of a delta-hedged strategy would not only be affected by the premium, but also by the individual option contract’s sensitivity to the realized variance.

While it is important to understand the sign of the premium, market participants would benefit greatly from quantifying the size of this premium. Carr and Wu (2009) were the first to do this analysis by creating synthetic variance swap rates derived from option prices<sup>2</sup>. The payoff from a variance swap is the difference between the realized variance of an asset and a predetermined forward-looking variance benchmark, called the swap rate. The average of these payoffs precisely measures the variance risk premium, and is therefore a very useful tool to quantify the premiums. Data on variance swaps is not readily available, however, because these contracts are traded over-the-counter (OTC) rather than through a public marketplace. By synthetically creating these contracts, this data can be replicated for any asset. Carr and Wu (2009) use this methodology to measure the volatility risk premium in 5 indices and 35 individual stocks. They find that there is a strong negative premium in the S&P 500 and 100, as well as with the Dow Jones Industrial Index. The individual stocks show more variation in the size of the premiums, and Carr and Wu (2009) attribute this to a “common stochastic variance risk” for which the market demands a premium. To test this theory, they regress the individual stock’s return variance on the S&P 500 variance to find a “variance beta”. The results

---

<sup>2</sup> Carr, Peter, and Liuren Wu. "Variance Risk Premiums." *Review of Financial Studies* 22.3 (2009): 1311-341.



show that stocks with higher variance betas had larger premiums, supporting their explanation for the differences in premium size.

Trolle and Schwartz (2010) adopted this methodology to measure the premium in energy commodities, specifically crude oil and natural gas<sup>3</sup>. While they find significant premiums in both commodities, the premium is larger for crude oil. Their paper also studies the seasonality of these premiums, showing that the premiums are higher for natural gas in winter, but are not statistically different for crude oil. They also examine whether the relationship between the variance swap return and the underlying security return is non-linear. The motivation for this exercise is that the negative relationship between index volatility and index return in the S&P500 is almost entirely driven by the strong negative correlation when index return is negative<sup>4</sup>. To test this with commodities, they regress the returns of the underlying futures on the returns of the variance swap, but first use only negative returns on the futures and then only the positive returns. They find that the relationship is in fact non-linear. For natural gas, there is a statistically significant positive relationship when returns on the futures are positive, but are weakly negative when returns are negative. For the S&P 500, there is a strong negative relationship when returns are negative, but no real relationship when returns are positive. Crude oil does not exhibit any strong relationships. Using these results, Trolle

---

<sup>3</sup> Trolle, Anders, and Eduardo Schwartz. "Variance Risk Premium in Energy Commodities." *Journal of Derivatives* 17.3 (2010): 15-32.

<sup>4</sup> Figlewski, Stephen and Wang, Xiaozu "Is the 'Leverage Effect' a Leverage Effect?" (2008), Working paper, NYU Stern School of Business.

and Schwartz (2010) argue that the return profile for natural gas is like a call option, while the S&P 500 returns resemble a put.

## **2.1 Importance of measuring VRP in VIX options**

While previous papers have analyzed this premium in various asset classes, this methodology has yet to be applied to the volatility asset class. This thesis aims to do that by measuring the variance risk premium in VIX options. Since the inception of VIX derivatives in 2004 (futures) and 2006 (options), the trade volume in these contracts grew substantially. *Figure 1* shows the growth since 2006 in VIX options, in terms of call and put volume, and VIX futures<sup>5</sup>. The values for 2011 are extrapolations of the realized numbers through April.

Active product innovation in the volatility asset class also indicates a growing demand for volatility products. One prominent example of this is the development of 16 new exchange traded products designed to capture different parts of the VIX futures terms structure. *Figure 2*, taken from Bill Luby's blog *VIX and More*, is an excellent visual representation of these new products<sup>6</sup>. The x-axis measures what part of the VIX futures term structure the product focuses on, and the y-axis shows the leverage of the product. Black tickers represent ETFs while the rest are ETNs, and the red "O" represents indexes for which options are available. Another recent development in the volatility space is the creation of volatility indexes for individual equities. On January

---

<sup>5</sup> Data comes from the CBOE website, [www.CBOE.com](http://www.CBOE.com).

<sup>6</sup> Luby, Bill. "Now Sixteen Volatility ETPs, Four of Which Are Optionable." [Vixandmore.blogspot.com](http://Vixandmore.blogspot.com), 26 Jan. 2011. Web. 15 Apr. 2011.

5<sup>th</sup>, 2011, the CBOE announced that they would be applying the VIX methodology on five popular equities -- Apple, Amazon, IBM, Google, and Goldman Sachs<sup>7</sup>. It is clear from these recent developments that market participants are more aware of volatility as an asset class of its own and are demonstrating a growing interest to transact in these products.

With more people participating in the VIX options and futures market, and the volatility asset class more broadly, it is important to better understand the pricing of these contracts. As this asset class develops further, there will certainly be interesting pair trading strategies among the different products and the underlying assets, and market participants who understand the pricing of the products will be in a better position to participate in these opportunities. It is also equally important for investors looking to hedge portfolio returns with the growing variety of volatility products to understand the relative pricing of these contracts. One aspect of this understanding is to quantify the variance risk premium on VIX options, which to this author's knowledge has not been previously measured in an academic scenario.

### **3. Methodology**

In this section, I will start by describing how the synthetic swap rate is generated, then I will discuss the calculation of the realized variance on the underlying security, and finally I will compare the average difference of the two through a dollar return measure and a continuously compounded excess return measure. These return measures are used

---

<sup>7</sup> "CBOE to Apply VIX Methodology to Individual Equity Options." (5 Jan. 2011). Web.

in other academic papers that quantify this premium, and provide common terms in which to compare the relative sizes of this premium<sup>8</sup>.

### 3.1 Swap Rate

The payoff from a variance swap is determined from the difference between the realized variance of the underlying security during the life of the contract and a swap rate determined at the initiation of the contract. A long position in a variance swap is a bet that the realized variance of the underlying security will be higher than the determined swap rate, and has a payoff equal to the realized variance minus the predetermined swap rate multiplied by a notional dollar amount. When variance swaps are initiated, there is no cash payment to either side. Assuming absence of arbitrage, this implies that the swap rate is a conditional risk-neutral expectation of the future variance of the underlying security<sup>9</sup>. Using this assumption, Carr and Wu (2009) derived a method to accurately calculate theoretical swap rates using a basket of out-of-the-money (OTM) puts and calls on the underlying asset<sup>10</sup>. The resulting formula is

$$SR(t, T) = \frac{2}{B(t, T)(T - t)} \times \left( \int_0^{F(t, T)} \frac{P}{X^2} dx + \int_{F(t, T)}^{\infty} \frac{C}{X^2} dx \right) \quad [1]$$

---

<sup>8</sup> Compared to values used in Carr and Wu (2009) and Trolle and Schwartz (2010)

<sup>9</sup> Carr, Peter, and Liuren Wu. "Variance Risk Premiums." *Review of Financial Studies* 22.3 (2009): 1311-341.

<sup>10</sup> Carr, Peter, and Liuren Wu. "Variance Risk Premiums." *Review of Financial Studies* 22.3 (2009): 1311-341.

where  $SR(t, T)$  is the swap rate initiated at time  $t$  for a contract that matures at time  $T$ ,  $B(t, T)$  is the time  $t$  price of a zero-coupon treasury bill that matures at time  $T$ ,  $P$  is the time  $t$  price of an OTM European put expiring at time  $T$  with a strike of  $X$ , and  $C$  is the time  $t$  price of an OTM European call expiring at time  $T$  with a strike of  $X$ .

This methodology is inadequate for VIX options, however, because VIX options contracts have a multiplier of 100 while VIX future contracts have a multiplier of 1,000. To avoid this issue, I substituted equation [2] into equation [1].

$$\left( \int_0^{F(t,T)} \frac{P}{X^2} dx + \int_{F(t,T)}^{\infty} \frac{C}{X^2} dx \right) = (T - t) \int_{-\infty}^{\infty} n(y) \sigma^2(y) dy \quad [2]$$

Where

$$y = d_2(X) = \left( \frac{-\ln \left( \frac{X}{F(t,T)} \right)}{\sigma(X)\sqrt{(T-t)}} - \frac{\sigma(X)\sqrt{(T-t)}}{2} \right) \quad [3]$$

for a call and  $n(y)$  represents the standard normal density function. This method, first developed by A. Matytsin (private communication) and first published by Gatheral (2006)<sup>11</sup>, takes a Gaussian weighted average of implied variance at every level of the newly defined moneyness function  $Y$ <sup>12</sup>. This method accounts for the multiplier issue

---

<sup>11</sup> Gatheral, Jim, and Nassim Taleb. *The Volatility Surface: A Practitioner's Guide*. John Wiley & Sons, 2006.

<sup>12</sup> Carr, Peter, and Roger Lee. "Volatility Derivatives." *Annual Review of Financial Economics* 1 (2009): 319-39.

because it is in terms of implied variances, rather than dollar prices, of options at different moneyness levels.

Formula [3] requires implied variances at every level of moneyness. I approximate these values by calculating Black implied volatilities – which uses the forward price instead of the spot price as an input into the Black-Scholes formula – from the available market data and linearly interpolating the implied volatilities between these points. I approximated the integral over infinite levels of moneyness by truncating the integral at the outermost market data points. This method underestimates the swap rate by a marginal amount, because implied variances outside these points have a negligible contribution to the total value of the integral due to the very small normal density weight. Wolfram Mathematica processed the value of the integrals.

### 3.2 Realized Variance

The realized variance of the contract is determined by using the following formula

$$RV(t, T) = \frac{365}{(T - t)} * \sum_{i=1}^{(T-t)} \left( \frac{F(t + 1, T) - F(t, T)}{F(t, T)} \right)^2 \quad [4]$$

where  $F(t+1, T)$  is the time  $t+1$  price of a future that expires at time  $T$ , and  $F(t, T)$  is the time  $t$  price of a future that matures at time  $T$ . This is the same measure of variance used in Carr and Wu (2009).

### 3.3 Variance Risk Premium

For every day in the data set, the difference between the realized variance and the swap rate for the nearest maturity contract of at least 12 days to maturity was calculated. Option contracts close to maturity can exhibit uncharacteristic pricing, so filtering out these maturities is important for calculating accurate swap rates. The maturities range from 12 to 40 days. Using the same measures as Carr and Wu (2009) and Trolle and Schwartz (2010) in order to have comparable results, I calculated a dollar return

$$Return = (RV - SR) \times \$100$$

and also a continuously compounded excess return.

$$Return \% = \ln \left( \frac{RV}{SR} \right)$$

## 4. Data and Results

VIX option data come from *Market Data Express*, and zero-coupon US Treasury Bill information came from Bloomberg. The original options price data set was filtered on the basis of volume, time to maturity, and availability of at least four market prices per day per maturity. Option prices with low volume have larger bid-ask spreads, and are therefore less accurate, so option contracts with a volume smaller than 10 were not used. Option contracts with short times to maturity can exhibit peculiar pricing tendencies, so contracts with less than 12 days to maturity were not used. Finally, because the formula for the synthetic swap rate involves an approximation of Black implied volatilities by

interpolation, it is important that there are multiple market price observations per day per maturity. Trade days when there were less than four market price observations at different moneyness levels were not used. Additionally, in order to ensure that the truncation of the integral was not a large source of error, days with less than 150 interpolated implied volatilities were not used. Option prices that met these requirements were determined using the average of the bid and ask, and prices for the futures were determined using daily settlement prices.

Passing the original 283,202 individual option contracts through the specified filters, the final dataset includes 1,107 daily observations of the variance risk premium. *Table 1* shows a summary of the data set as various filters limited the sample size.

**Table 1 – Summary of Data Set**

Number of original options	283,202
Number of options after volume filter	92,419
Number of options with rest of filters	14,668
Average number of options per swap rate	13.25
Number of days with a swap rate	1,107

#### 4.1 Summary of Results

*Figure 3* shows a time series of the swap rates and the corresponding realized variance of that contract. This graph depicts how realized variance spikes, followed by a lagged reaction in swap rates, and then a normalization of both levels. *Figure 4* shows the difference of these two values, representing the variance risk premium of the contract initiated at that date. It is clear from this figure that the variance risk premium remains negative most of the time, but infrequently is very positive.



*Table 2* shows summary statistics of the realized variance during the life of the contract, the theoretical swap rate at the initiation of the contract, and the variance risk premium. The relatively large average realized variance of VIX futures suggests that the volatility of volatility is substantial. The wide range of realized volatilities – from 7.5% to 419.7% on an annualized basis – and the positive skewness of 2.402 supports the argument that the variance risk premium exists as a way to compensate sellers of VIX options for the infrequent, but very large, volatility spikes. It is also interesting to compare the standard deviation of the realized variance and the swap rate. The larger standard deviation in the realized variance suggests that swap rates do not react in the same magnitude to realized variance spikes, which can be attributed to the forward looking characteristic of the swap rate and the mean reverting characteristic of volatility.

*Table 3* shows the dollar return and continuously compounded excess return of a long position in the theoretical variance swaps. The dollar return terms are derived from multiplying the difference between the realized variance and the swap rate by a notional value of \$100, suggesting that the average return from investing \$100 dollars in these swaps is -\$3.26. The continuously compounded excess return values are derived from taking the natural log of the realized variance divided by the swap rate, and represents the percentage payoff from continuously purchasing these contracts. It is valuable to compare these results to variance risk premiums derived in other academic papers using a similar methodology in order to better understand the relative size of the premium. Carr and Wu (2009) measure this premium in the S&P500 and the Dow Jones Industrial Index, among many other individual stocks, deriving dollar returns of -\$2.74 and -\$2.58 and log returns of -0.66 and -0.61 respectively. Trolle and Schwartz (2010) also measure

this premium in crude oil and natural gas futures, and derive a dollar premium of -\$3.58 and -\$2.96 and a log return premium of -.265 and -.223 respectively. The premium values derived in this paper – a dollar return of -\$3.26 and a log return of 0.311— fit the range of premiums previously measured, and seem to have characteristics closer to the premium in commodity futures.

## 4.2 Is the premium time-varying?

Similar to the approach of Carr and Wu (2009) and Trolle and Schwartz (2010), I performed two regressions to measure whether the variance risk premium is time-varying. Because of the overlap in observations of realized variance between consecutive contracts, I used Newey-West T-statistics with a lag of 40, corresponding to the longest maturity in the dataset<sup>13</sup>. The first regression, seen below, uses the absolute levels of the realized variance and the swap rate.

$$RV = a + B(SR) + \epsilon$$

Using a null hypothesis representing a premium that does not vary with time implies that  $a = 0$  and  $B = 1$ . The second regression, seen below, is similar but uses log terms instead.

$$\ln(RV) = a + B \ln(SR) + \epsilon$$

Assuming constant variance risk premium, the null hypothesis is again  $a = 0$  and  $B = 1$ .

Results for both regressions can be found in *Table 4*. The coefficient on the Swap Rate variable is 0.924 in the first regression, and 0.874 in the log return regression.

These values are close to 1, suggesting that the premium only slightly varies with time

---

<sup>13</sup> Newey, Whitney, and Kenneth West. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55.3 (1987): 703-08.

and the level of the Swap Rate. These results differ from the regression results for different assets. It can be inferred from a comparison of these results that the variance risk premium in VIX options is less time varying and is less affected by the value of the swap rate.

## **5. Trade Strategy**

The negative variance risk premium suggests that variance realized during the life of a variance swap is on average smaller than the variance implied by the price of the contract. Since these hypothetical swaps are derived from individual implied volatilities of option contracts and the realized volatility of the option's underlying future contract, a negative risk premium in a variance swap should mean that option strategies that are net credit should be profitable, on average. This paper tests this claim by running simple net credit trading strategies.

The four trading strategies measured in this paper can be summarized in *Table 2* below. The numbers in the buy and sell columns correspond to the position of the option in a list of the OTM options. For instance, sell 1 and buy 2 means that the contract closest to the ATM option is sold and the next furthest OTM option contract is bought. These options are OTM relative to the future price maturing at the same time as the option contract. The options have been filtered so that the bid-ask spread is no larger than 0.35, in order to prevent the spread from deteriorating the premium collected. These strategies do not involve both puts and calls because daily volume in these two types of contracts is typically uneven, with puts being more actively traded when spot VIX is high and calls being more actively traded when spot VIX is low.

The column labeled “% Return” is the average of the annualized percentage return from each contract in the sample set. The return for each contract is measured by dividing the payoff of the position – the net premium collected plus the payoffs of each individual option position – by the margin requirement required to hold this position, which is the difference between the strikes x 100. It is interesting to note that this simple net selling strategy is profitable for calls, which has an average annualized return of up to 16.55%, while a similar strategy implemented with puts has an average annualized return as low as -69.87%. One reason this strategy may be more profitable with calls is that volume in this market is driven by retail investors who wish to purchase negative correlation with market returns to hedge portfolio returns, which is most easily done by purchasing calls. This could translate into a richer premium in calls, making these products more attractive to sell. The methodology used in this paper cannot separate the variance risk premium between calls and puts, but this might be an interesting area of future research.

These crude trading strategies were selected on the basis of simply testing the idea that being in a net short position in VIX options would be profitable because of the negative variance risk premium. All of the strategies were implemented indiscriminately on the sole basis of having a narrow enough bid-ask spread with the intention of enforcing an unbiased selection process while trading on the thesis mentioned above. Dynamic selection of option positions and management of those positions, along with implementing more complex strategies, would likely increase return.

**Table 2 – Summary of Trade Strategies**

<b>Calls</b>				<b>Puts</b>			
<b>Sell</b>	<b>Buy</b>	<b>% Return</b>	<b># Contracts</b>	<b>Sell</b>	<b>Buy</b>	<b>% Return</b>	<b># Contracts</b>
1	2	6.25%	1079	1	2	-69.87%	998
1	3	16.55%	1079	1	3	-43.37%	897

\* "% Return" is an average of the annualized return from each contract

## **6. Conclusion**


This paper synthetically creates variance swap data on VIX futures to infer the variance risk premium in VIX options. Measuring this premium is important for the growing number of market participants in VIX options and futures contracts, as well as the volatility asset class as a whole. The premium is measured in terms of dollar return with a \$100 notional, and also in a continuously compounded excess return term. In both cases, the premium is negative, with a dollar return of \$-3.26 and a log return of -.312, which suggests that being short a variance swap on VIX futures is profitable. The size of this premium fits the range of premiums measured in other academic papers using a similar methodology, but is more negative than the premium in the S&P 500 index. Regressing the realized variance on the calculated swap rate shows that the premium is less time varying than other premiums in other assets classes, and is only slightly affected by the value of the swap rate. A negative variance risk premium in VIX futures should also imply that net credit VIX option positions are profitable. This paper finds that a net credit position selling calls is profitable, with an average annualized return as high as 16.55%, while a similar position that sells puts instead has an average annualized return as low as -69.87%. While this thesis hypothesizes that selling calls is more profitable

because the premium is richer in these options due to the insurance affect, analyzing this difference in return could be an interesting area of future research.

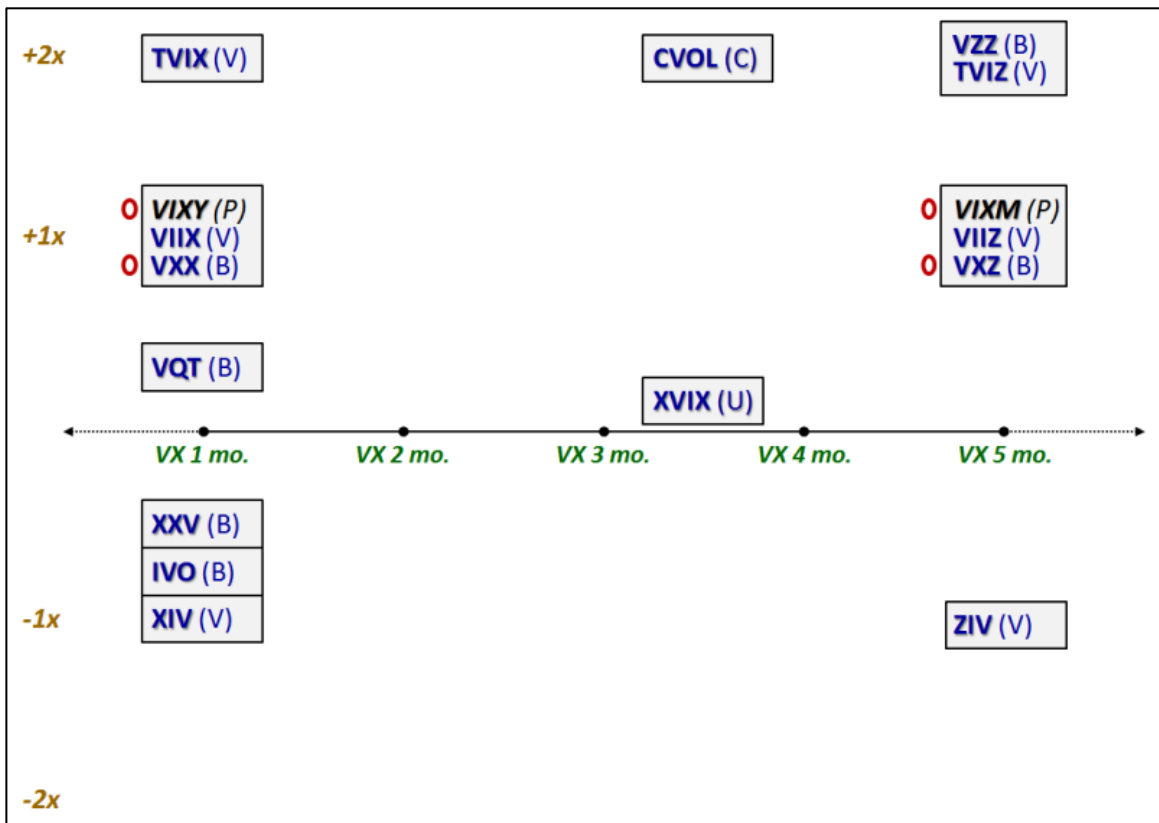
## Bibliography

- Bakshi, Gurdip, and Nikunj Kapadia. "Delta-Hedged Gains and the Negative Market Volatility Risk Premium." *Review of Financial Studies* 16.2 (2003): 528-66.
- Carr, Peter, and Liuren Wu. "A Tale of Two Indices." *Journal of Derivatives* 13.3 (2006): 13-29.
- Carr, Peter, and Liuren Wu. "Variance Risk Premiums." *Review of Financial Studies* 22.3 (2009): 1311-341.
- Carr, Peter, and Roger Lee. "Volatility Derivatives." *Annual Review of Financial Economics* 1 (2009): 319-39.
- "CBOE to Apply VIX Methodology to Individual Equity Options." (5 Jan. 2011). Web.
- Figlewski, Stephen and Wang, Xiaozu "Is the 'Leverage Effect' a Leverage Effect?" (2008), Working paper, NYU Stern School of Business.
- Gatheral, Jim, and Nassim Taleb. *The Volatility Surface: A Practitioner's Guide*. John Wiley & Sons, 2006.
- Luby, Bill. "Now Sixteen Volatility ETPs, Four of Which Are Optionable." Weblog post. 26 Jan. 2011. Web. 15 Apr. 2011.
- Newey, Whitney, and Kenneth West. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55.3 (1987): 703-08.
- Trolle, Anders, and Eduardo Schwartz. "Variance Risk Premium in Energy Commodities." *Journal of Derivatives* 17.3 (2010): 15-32.
- Woodard, Jared. "Options and the Volatility Risk Premium." *Financial Times*. Web. 6

**Figure 1 – VIX option and VIX future volume**

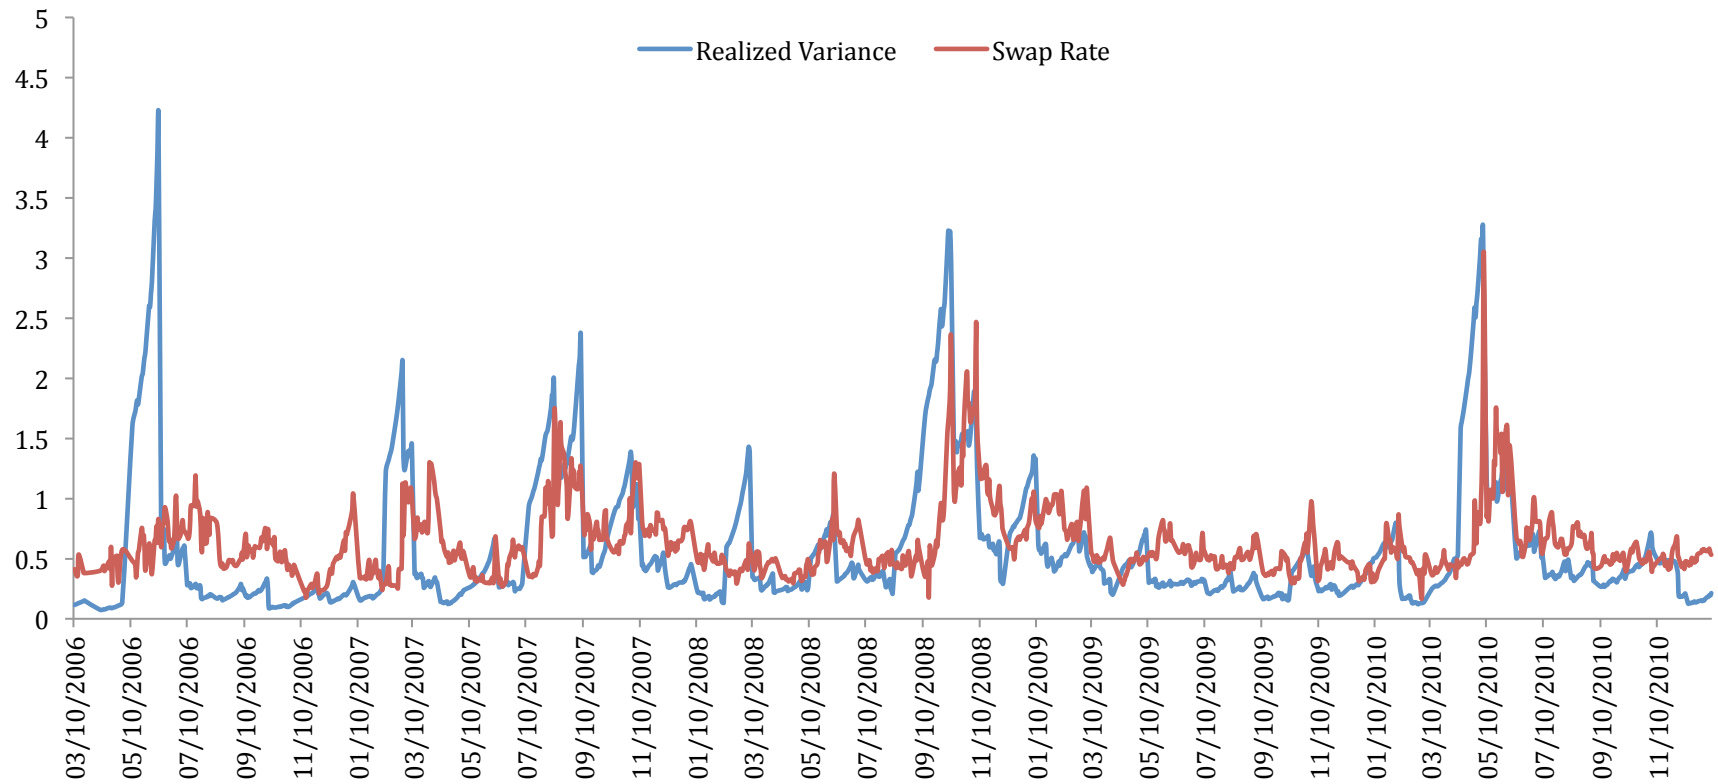


**Figure 2 – Mapping VIX futures ETPs**






**Figure 3 – Realized Variance versus Swap Rate Time Series**



**Figure 4 – Variance Risk Premium Time Series**



**Table 2 – Summary Statistics of RV, SR, and VRP**

	<i>Realized Variance</i>	<i>Swap Rate</i>	<i>VRP</i>
Mean	0.595	0.628	-0.033
Standard Error	0.018	0.009	0.016
Median	0.371	0.551	-0.181
Standard Deviation	0.595	0.295	0.530
Sample Variance	0.355	0.087	0.281
Kurtosis	6.595	8.888	7.830
Skewness	2.402	2.306	2.488
Range	4.122	2.860	4.379
Minimum	0.075	0.165	-1.011
Maximum	4.197	3.025	3.367
Observations	1107	1107	1107
Confidence Level(95.0%)	0.035	0.017	0.031

**Table 3 – Summary Statistics Returns**

	<i>Dollar Return</i>	<i>Log Excess Return</i>
Mean	-3.260	-0.312
Standard Error	1.593	0.021
Median	-18.145	-0.385
Standard Deviation	53.003	0.707
Sample Variance	2809.269	0.500
Kurtosis	7.830	0.570
Skewness	2.488	0.667
Range	437.871	4.344
Minimum	-101.132	-1.989
Maximum	336.740	2.355
Observations	1107	1107
Confidence Level(95.0%)	3.126	0.042

**Table 4 - Regression Results**

	$RV = a + B(SR) + \epsilon$			$\ln(RV) = a + B \ln(SR) + \epsilon$		
	a	B	$R^2$	a	B	$R^2$
VIX	0.015 (0.2)	0.924 (7.54)	0.2092	-0.38 (-3.41)	0.874 (6.49)	0.1951
<i>Regression Results taken from Carr and Wu (2009)</i>						
SPX	0.01 (-1.416)	0.455 (-4.596)	0.262	-0.891 (-2.593)	0.919 (-0.684)	0.378
DJX	0.013 (1.524)	0.443 (-4.046)	0.19	-1.21 (-2.859)	0.781 (-1.467)	0.253
NDX	-0.023 (-1.329)	0.995 (-0.042)	0.571	-0.17 (-1.233)	1.06 (0.876)	0.672
<i>Regression Results taken from Trolle and Schwartz (2010)</i>						
CRUDE	0.04 (4.652)	0.539 (-7.87)	0.313	-0.819 (-6.078)	0.72 (-4.088)	0.322
NG	0.076 (2.936)	0.691 (-4.589)	0.314	-0.317 (-3.734)	0.919 (-1.173)	0.45

\*The values in parenthesis are Newey-West T-Statistics with a lag of 40