

In Search of a Better Volatility

Firas Kotite^a

Clemens Kownatzki^b

Abstract

We derive an alternative volatility index from options on E-Mini S&P 500 futures and compare it with the VIX to see which index could provide a more efficient measure of volatility and risk. VCME, our alternative volatility measure, and the VIX are very similar in price and trend, are quite efficient at forecasting future volatility in the short-term, but lose their effectiveness over longer periods of time. We do not find any meaningful relationship with volatility and future stock returns. However, we propose that VCME may be more attractive to a financial institution seeking responsive risk measures, while on average generating less deviations from actual volatility at any time frame from 1 to 21 days forward.

The research for this paper has been made possible through the generous support of the
CME GROUP FOUNDATION.

Keywords: VIX, Volatility forecasting, Implied Volatility, Risk

^a Claremont Graduate University, firm.kotait@gmail.com

^b Pepperdine University, clemens.kownatzki@pepperdine.edu

Introduction

Motivated by the quest to find an efficient measure of risk, we begin this study examining the effectiveness of the volatility index (VIX) and investigate whether the VIX is a good proxy for risk. This study is an extension of previous work on volatility, which concluded that the VIX consistently over-estimates actual volatility in normal times but it underestimates volatility in times of crises (Kownatzki, 2015). The VIX represents implied volatility from option prices and it is commonly known as the “fear index.” Robert Whaley, its creator, suggests that the VIX is a cost-effective way to hedge risk because it provides a reliable estimate of expected short-term volatility (Whaley, 1993). While the forward-looking aspects of the VIX are an important feature of investor risk expectations, the investor fear gauge (Whaley 2000) may need to be re-examined since these expectations of risk, derived from S&P 500 options prices, are typically inflated (Kownatzki 2016).

Option traders have long recognized that fear, and therefore expected risk, is overstated. This over-estimation of risk translates directly into higher option premiums, given that volatility, the proxy for risk, is perhaps the most critical variable determining the pricing of options. Since volatility is a latent measure that cannot be observed, it needs to be estimated. Market participants derive an implied volatility measure by inverting the Black-Scholes formula (Black and Scholes 1973), and solve for the only un-known variable, volatility. In practice, implied volatility is thusly calculated for each strike price of any tradeable option. Recognizing that the uncertainty over the true value of volatility may be the main factor in over-estimating risk, option traders have an edge and they are given an attractive incentive for their typical go-to strategy - selling, rather than buying options. In essence, option traders feed on the fear of investors just like traditional insurance companies feed on people’s fears over the unknown by inflating option premiums.

Previous research also suggests that the level of the VIX does not provide meaningful answers to the important question of how risk affects future stock returns (Kownatzki, 2015). In an effort to improve upon current risk metrics such as the VIX, we investigate whether an alternative volatility index, derived from options prices on E-Mini S&P 500 futures, could

provide better estimates of future market risk. In this context, we examine if such an index can help us better understand how current market volatility affects future stock returns.

Since the VIX is derived from options prices on the S&P 500 Index, which in itself is not a tradeable instrument, it may not fully reflect actual market sentiment and it may therefore provide inaccurate investor expectations on risk. *Futures* as well as *Options on Futures* may have additional information priced in since futures prices can often have a richer information content, particularly within the important 30-day forward-looking time period associated with the VIX.

Research Questions

Our first task focuses on the model derivations and construction of this alternative VIX measure. We propose a new volatility measure, derived from exchange-traded options on E-Mini S&P 500 futures, a financial instrument created by the CME Group in Chicago. This new volatility measure shall be called VCME.

Given the constraints of the market place and the intricacies of exchange-traded options on futures, we examine whether the results provide reasonable volatility estimates. The bulk of this study therefore focuses on comparisons between the two volatility measures: VIX and VCME

We conduct comparison studies of the two volatility measures by examining the following research questions:

- VIX versus VCME: Which index provides better estimates of future volatility?
- VIX versus VCME: Which index provides better estimates of future S&P 500 returns?

Our study primarily focuses on a direct comparison of the two volatility measures. In addition to comparing the effectiveness of both volatility measures in term future volatility forecasts, we also include inter-temporal effects on future volatility as well as future S&P 500 index returns.

Literature Review

The literature on volatility, the widely-accepted metric of financial risk, is rich but also laden with controversy. As early as Frank Knight (Knight 1921) economists questioned the usefulness of a metric that can be measured when uncertainty, the true nature of risk, was simply not measurable. In order to distance itself from gambling and to associate investment decisions with science rather than pure speculation, quantifiable metrics were critical for the prospects of the finance industry - how else could the typically risk-averse investors be assured that their investments were sound within the given risk parameters.

Among the various risk metrics, implied volatility, represented by the VIX, has been the go-to metric for many financial risk management applications despite the fact that many of the assumptions for these risk models do not hold up to empirical examinations.

The notion that financial asset returns were normally distributed permeated most financial models for decades even though Benoit Mandelbrot revealed the presence of leptokurtosis¹ in empirical return distributions (Mandelbrot 1963). There have been other critics of the normality assumption in return distributions, but it took another four decades before Eugene Fama (1970), the creator of the Efficient Market Hypothesis raised some concerns. He warned that “empirical examinations of asset prices reveal that the problems are serious enough to invalidate most applications of the CAPM” (Fama and French 2004). More recently, empirical examinations of daily stock return distributions suggest that financial return series are heavy-tailed and possibly skewed (Rachev et al 2005). Nevertheless, it took the financial crisis of 2008/09 before academics as well as finance practitioners seriously questioned Gaussian normality, one of the main assumptions in many financial models.

Until that crisis period, most researchers were more occupied with forecasting volatility rather than seriously questioning the main assumptions that allow us to estimate risk. One evidence of that pre-occupation with volatility forecasting comes in the form of an extensive glossary of over 150 GARCH-type models compiled by Tim Bollerslev. Bollerslev’s work on

¹ Kurtosis measures the mass of a distribution’s tails. The kurtosis of a normal distribution is 3. Values above 3 are considered leptokurtic or simply heavy-tailed (Stock & Watson 2011)

dynamic volatility forecasting (Bollerslev 1986) culminated in the creation of Generalized Auto Regressive Conditional Heteroscedasticity (GARCH)². However, Bollerslev lamented the fact that there are so many competing models in what he described as a perplexing ‘alphabet-soup’ of acronyms and abbreviations for these models (Bollerslev 2007).

In addition to GARCH-type models, historical and implied volatility models have also been used to estimate future risk. Blair, Poon and Taylor (2000) ranked the VIX highest for providing the most accurate out-of-sample forecasts for volatility. In an extensive survey of volatility forecasting models Granger and Poon review 93 published papers on various volatility models (Granger and Poon 2003). They found that implied volatility from options using the Black-Scholes model ranked higher than historical volatility and GARCH (Granger and Poon 2003), albeit with sometimes contradicting results.

Studies by Martins and Zein (2002) suggest that volatility forecasts have higher explanatory power at shorter time horizons whereas Christoffersen and Diebold notice a decrease in the accuracy of equity and foreign exchange volatility forecasts as they increase the time horizon from one to ten days. They conclude that volatility forecasts may not be of much value if the horizon of interest is more than ten or twenty days (Christoffersen and Diebold 2000). This directly calls into question the effectiveness of the VIX which provides a volatility forecast of 30 calendar days.

Multiple competing models also raise the question as to the selection of the best model. More importantly, the gravity of the recent crisis suggests that most if not all of the models failed to produce reliable risk estimates that could have prevented such devastating economic repercussions. Robert Whaley argued that the VIX is a cost-effective way to hedge risk (Whaley, 1993) but we are not aware of institutions or investors who have been able hedge their way out of the crisis with the help of the VIX. Conflicting results from competing models and volatility

² Whereas ARCH specified conditional variance as a linear function of past sample variances only, the GARCH process allowed lagged conditional variances to enter into the process as well. GARCH estimates variance by two distributed lags, one on past squared residuals and a second one on lagged values of the variance itself. The simplest GARCH process is GARCH (1,1) given by:

$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_t h_{t-1}$ where, $\alpha_0 > 0$, $\alpha_1 \geq 0$ and $\beta_t > 0$ (Bollerslev 1986)

estimates also encourage us to find an alternative to existing volatility models. Our work in the following chapters is an attempt to generate a new volatility index with more reliable risk estimates.

Data & Methodology

We obtain data from ivolatility.com for put and call options on E-mini S&P 500 Futures contracts. Our data set comprises daily end-of-day settlement prices from July 2011 to July 2016 based on the National Best Bid and Offer (NBBO³). The total dataset includes over 8 million observations for call and put prices since for each given settlement date, there are several futures and options expirations and thus, thousands of strike prices. For instance, on our starting date, 11 June 2011, there are about 4,000 individual settlement prices relevant to our calculations. These consist of the strikes and put/call prices for options expiring on 7/15/2011, 7/22/2011, 7/29/2011, 8/5/2011, 8/19/2011, 8/31/2011 and 9/16/2011 based on the underlying futures contracts expiring on 9/16/2011 in addition to strikes for options expiring on 9/30/2011, 10/21/2011, 11/30/2011 and 12/16/2011 for underlying futures contracts expiring on 12/16/2011.

Our new volatility measure, VCME, is derived from the 30-day expected volatility of the S&P 500 futures using a similar methodology as described in the CBOE White Paper for the calculation of the VIX (CBOE 2014). We use Bi-Weekly options to determine the two main components of the VCME measure that are the near and next-term options with an expiration date between 23 to 50 days. Our use of a 50-day upper bound differs from the VIX's 37 day upper bound due to the nature of the E-Mini Futures that led to multiple cases where we could not match a next-term option to a near-term contract. We use Stata to calculate the different variables associated with the volatility measure. We assume an interest rate of 0.0025 for the front and 0.004 for the back contract based on an average yield of T-Bills for our time periods observed which is similar to the process outlined by CBOE (CBOE 2014). We then follow the methodology used in the calculation of the VIX to get the following:

$$VCME = 100 * \sqrt{\frac{t_1}{t_M} * \frac{t_2 - t_M}{t_2 - t_1} * \sigma_1 + \frac{t_2}{t_M} * \frac{t_2 - t_M}{t_2 - t_1} * \sigma_2} \quad (1)$$

³ NBBO is the lowest ask price and the highest bid price available to investors. U.S. securities brokers are required to guarantee price execution based on NBBO.

Where,

T_1 = Number of minutes to the settlement of the Near-Term Options

T_2 = Number of minutes to the settlement of the Next-Term Options

T_m = Number of minutes in 30 days.

$$\sigma^2 = \frac{1}{t_1} * 2e^{r_1 * t_1} * \sum \left(P_i \frac{(S_n - S_{n-1} + S_{n+1} - S_n)}{2 S_i^2} \right) - \left(e^{r_1 * t_1} * \frac{Call - Put}{S_M^2} \right)^2$$

S_n = Current strike price

P_i = Corresponds to the option price; of a call if $S_i > S_{ATM}$; and of a put if $S_i < S_{ATM}$. If $S_i = S_{ATM}$ then an average between the put and call prices is used.

Sample Calculation

The following example uses an actual S&P 500 future, sampled on July 5th 2016, to illustrate the VCME calculation used in this paper. We selected the near-term option that expires on July 29th 2016 and the next-term option of August 5th 2016, to satisfy the 23 to 50-day window required for the calculation. In this case the expiration periods are 24 for near-term and 31 days for next-term.

Sigma is calculated using the strike price with the smallest absolute difference in Call and Put prices from the closing price of the S&P 500 futures contract on each specific day. In addition, we eliminate options with a bid or ask price of zero, as well as truncate the series to exclude instances where the call or put are less than 5 cents on 2 consecutive days. In this case, the ATM strike is 2085 and the series only includes prices between 2025 and 2270. Once the ATM strike is determined, we calculate sigma 1 for the near-term and sigma 2 for the next-term option using the formulas listed above, where P_{ATM} is 27.25 and 31.25 for the near and next-term options respectively. This leads to a Sigma 1² of 0.028 and a Sigma 2² of 0.024 and a VCME of 15.75.

$$\sigma_1^2 = \frac{1}{0.0658} * 2e^{0.0025 * 0.0658} * 0.000926 - \left(e^{0.0025 * 0.0658} * \frac{26.75 - 27.75}{2085} \right)^2 \quad (2)$$

$$VCME = 100 * \sqrt{\frac{0.0658}{0.0849} * \frac{0.0849 - 0.0821917}{0.0849 - 0.0658} * 0.0281 + \frac{0.0849}{0.0821917} * \frac{0.0849 - 0.0821}{0.0849 - 0.0658} * 0.0243} \quad (3)$$

Where,

$T_1 = 0.0658$, $T_2 = 0.0849$, $T_m = 0.0821$

Near-Term	24 days		Next-Term	31 days	
Strike	Call	Put	Strike	Call	Put
2025	70	72.25	2025	73.5	75.5
2030	66	68	2030	69.5	71.5
2035	61.75	64	2035	65.5	67.5
2040	58	60	2040	61.75	63.5
2045	54	56	2045	58	59.75
2050	50.25	52.25	2050	54.25	56
2055	46.5	48.25	2055	50.5	52.25
2060	43	44.75	2060	47	48.5
2065	39.25	41	2065	43.5	45
2070	36	37.5	2070	40	41.5
2075	32.5	34.25	2075	36.75	38.25
2080	29.5	31	2080	33.5	35
2085	26.75	27.75	2085	30.75	31.75
2090	23.75	24.75	2090	27.75	28.75
2095	21	22	2095	25	26
2100	18.5	19.25	2100	22.25	23.25
2105	16	16.75	2105	19.5	20.5
2110	13.5	14.5	2110	17.25	18
2115	11.5	12.25	2115	15	15.75
2120	9.75	10	2120	12.75	13.5
2125	8	8.5	2125	10.75	11.75
2130	6.5	7	2130	9	10
2135	5	5.75	2135	7.5	8.25
2140	4.1	4.5	2140	6.25	7
2145	3.15	3.55	2145	5	5.75
2150	2.4	2.8	2150	4.1	4.55
2155	1.8	2.25	2155	3.2	3.65
2160	1.3	1.75	2160	2.5	2.95
2165	0.95	1.35	2165	1.95	2.35
2170	0.7	1.05	2170	1.5	1.9
2175	0.5	0.9	2175	1.15	1.55
2180	0.4	0.75	2180	0.85	1.25
2185	0.3	0.65	2185	0.7	1.05
2190	0.2	0.55	2190	0.55	0.9
2195	0.2	0.5	2195	0.45	0.8
2200	0.3	0.5	2200	0.35	0.7
2205	0.1	0.45	2205	0.3	0.65
2210	0.1	0.45	2210	0.25	0.6
2215	0.1	0.4	2215	0.2	0.55
2220	0.1	0.4	2220	0.2	0.5
2225	0.05	0.4	2225	0.15	0.5
2230	0.05	0.4	2230	0.15	0.5
2235	0.05	0.4	2235	0.1	0.45
2240	0.05	0.35	2240	0.1	0.45
2245	0.15	0.35	2245	0.1	0.45
2250	0.15	0.35	2250	0.1	0.45
2255	0.15	0.2	2255	0.1	0.4
2260	0.05	0.35	2260	0.1	0.4
2265	0.05	0.35	2265	0.1	0.4
2270	0.1	0.35	2270	0.05	0.4

Table 1 – Sample of near-term and next-term option prices on E-Mini S&P 500 futures.

While this is an illustrative example, as stated previously, our final dataset includes around 4 million observations which allows us to calculate 1,257 daily VCME measures, ranging from values of 9 to 53, with an average of 17.20.

Model Specifications

The main variables in our study are defined as follows: *spx*, *vix*, *vcme* and *r_vol* represent the daily closing prices of the S&P 500 Index, the VIX, VCME (our alternative VIX measure) and Realized (historical) Volatility, respectively. *lnr1_spx*, *lnr5_spx*, *lnr10_spx* and *lnr21_spx* are the 1-, 5-, 10- and 21-day forward-looking log returns of the S&P 500.

Days are measured in trading days and the time intervals have been chosen to reflect 1-day, 1-week, 2-week and 1-month calendar periods. Log returns are calculated as $\ln(SPX_{t+1} / SPX_t)$ for 1-day returns and $\ln(SPX_{t+21} / SPX_t)$ for 21-day returns. Similarly, *d1_vix*, *d5_vix*, *d10_vix*, *d21_vix* are the 1-, 5-, 10-, and 21-day forward looking differences calculated as $vix_{t+n} - vix_t$ and *d1_vcme*, *d5_vcme*, *d10_vcme*, *d21_vcme* are the 1-, 5-, 10-, and 21-day forward looking differences calculated as $vcme_{t+n} - vcme_t$. *r_vol* is based on the 21-day period standard deviation of daily log returns and then annualized by multiplying with the square root of 365.

Preliminary Statistics

Our three main variables of interest in terms of comparisons are *vix*, *vcme* and *r_vol*. Their summary statistics are as follows:

variable	Obs	Mean	Std. Dev.	Min	Max
<i>vix</i>	1257	.1741054	.0591713	.1032	.48
<i>vcme</i>	1257	.1719703	.0601138	.0987011	.539649
<i>r_vol</i>	1257	.1704685	.0855814	.058	.5947

Table 2 – Summary statistics of *vix*, *vcme* and *r_vol*.

We note that the mean of the *vix* is higher than *vcme* which in turn is higher than *r_vol*. By contrast, the *vix* shows the lowest standard deviation at about 6.9% while *r_vol* has the highest standard deviation at about 8.6%. The summary statistics of our measure *vcme* appear to be somewhere between the *vix* and *r_vol*.

A visual comparison of the time series plots of vix and r_vol reveals that the two measures are very similar, although there are a few periods when vcme appears to be higher than the vix. Still, the two variables are highly correlated at 0.9924.

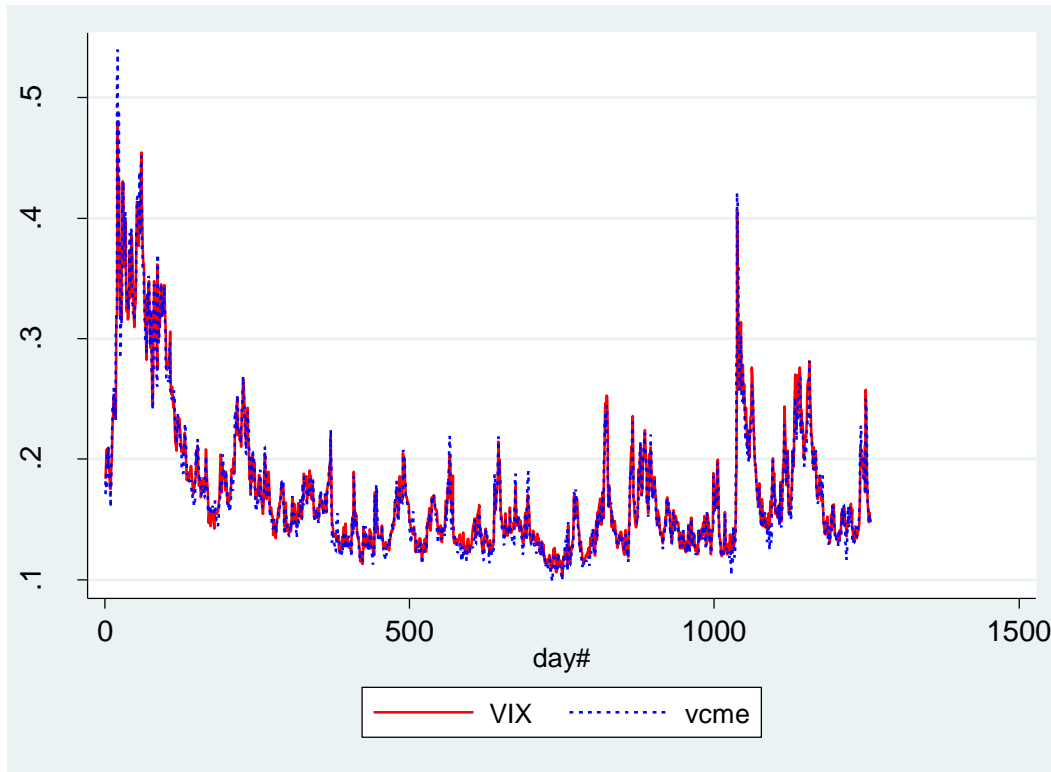


Figure 1 – VIX versus VCME.

In addition, we examine the correlation with r_vol and find that actual volatility shows a much lower correlation with the two measures vix and vcme. Further, the 21-day lagged values of vix and vcme now see a further decline and correlations drop to only about 63%.

	vix	vcme	r_vol		L21. vix	L21. vcme	r_vol
vix	1.0000			vix	1.0000		
vcme	0.9924	1.0000		L21. vcme	0.9925	1.0000	
r_vol	0.8323	0.8251	1.0000	L21. r_vol	0.6353	0.6306	1.0000

Table 3 – Correlations of vix, vcme and r_vol.

Table 4 – Correlations of 21-day lags vix, vcme and r_vol.

Moreover, the histograms of the differences in vix and vcme suggest that their respective changes are not normally distributed but instead show leptokurtic behavior which extends from 1-day to 21-day periods.




Figure 2 – Histogram of 1-day VIX changes




Figure 3 – Histogram of 1-day VCME changes




Figure 4 – Histogram of 21-day VIX changes




Figure 5 – Histogram of 21-day VCME changes

Non-Gaussian distribution of 1-day changes is confirmed by skewness values of 1.44 and 1.59 for vix and vcme respectively.

Financial time series data typically exhibit auto-correlation, non-stationarity and heteroskedasticity. In our dataset, the daily closing prices of spx, vix and vcme are clearly auto-

correlated, however, the log returns of spx and differences in vix and vcme do not exhibit autocorrelation.

We observe ADF⁴ t-statistics for spx of -1.012 for 1-day lags and -1.355 for 21-day lags both of which suggest that the daily prices of spx display a non-stationary stochastic trend. By contrast, the t-statistics of vix (-3.219), vcme (-3.515) as well as their 21-day differences in d21vix (-10.717) and d21vcme (-10.972) clearly reject the null hypothesis. Similarly, we cannot find evidence of a non-stationary stochastic trend in 1-day and 21-day log returns of spx with t-statistics of -24.586 and -9.609 respectively.

Heteroskedasticity is present when regressing r_vol on vix and vcme but it is also present when we regress 1-day to 21-day log returns of spx on either vix or vcme. To correct for heteroskedasticity, we use ‘White-corrected’ robust standard errors for all time-series regressions to follow. We now have to assess to what extent these initial insights help us answering our two main research questions.

VIX versus VCME: Which index provides better estimates of future volatility?

To address our first research question, we initially run a set of regressions using the following model specifications:

$$r_vol_t = \beta_0 + \beta_1 vix_t + \varepsilon_t \quad (4)$$

$$r_vol_t = \beta_0 + \beta_1 vcme_t + \varepsilon_t \quad (5)$$

We also examine how the lagged values of vix and vcme change the forecasting efficiency of future realized volatility. For the 1-, 5-, 10- and 21-day lagged values of the vix and vcme, our regressions take on the following format:

$$r_vol_t = \beta_0 + \beta_1 vix_{t-n} + \varepsilon_t \quad (6)$$

⁴ The Augmented Dickey-Fuller (ADF) test examines the presence of a stochastic trend. Under the null hypothesis of this test, Y_t has a stochastic trend, i.e. a unit root, whereas the alternative hypothesis suggests that Y_t is stationary. Critical values are -2.570 (10%), -2.860 (5%) and -3.430 (1%).

$$r_vol_t = \beta_0 + \beta_1 vcme_{t-n} + \varepsilon_t \quad (7)$$

Tables 5 and 6 show the regression summaries for our variables of interest. The two volatility measures have very similar coefficients, all highly significant at the 1% level.

VARIABLES	(1) r_vol	(2) r_vol	(3) r_vol	(4) r_vol	(5) r_vol
vix	1.204*** (0.0408)				
L.vix		1.222*** (0.0390)			
L5.vix			1.254*** (0.0368)		
L10.vix				1.222*** (0.0345)	
L21.vix					0.918*** (0.0302)
Constant	-0.0391*** (0.00644)	-0.0424*** (0.00615)	-0.0480*** (0.00583)	-0.0423*** (0.00547)	0.0103** (0.00504)
Observations	1,257	1,256	1,252	1,247	1,236
R-squared	0.693	0.714	0.751	0.711	0.404

Robust standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 5 – Regressions of r_vol on vix and its lagged values.

VARIABLES	(1) r_vol	(2) r_vol	(3) r_vol	(4) r_vol	(5) r_vol
vcme	1.175*** (0.0438)				
L.vcme		1.194*** (0.0410)			
L5.vcme			1.225*** (0.0371)		
L10.vcme				1.194*** (0.0345)	
L21.vcme					0.897*** (0.0304)
Constant	-0.0315*** (0.00686)	-0.0348*** (0.00641)	-0.0404*** (0.00581)	-0.0350*** (0.00543)	0.0159*** (0.00506)
Observations	1,257	1,256	1,252	1,247	1,236
R-squared	0.681	0.703	0.740	0.702	0.398

Robust standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 6 – Regressions of r_vol on vcme and its lagged values.

Interestingly with both variables, R^2 increases from just under 0.7 at contemporaneous effects to about 0.75 for 5-day lagged values. At that one-week time period, in both cases about 75% of the variations in r_vol are explained by the model. For periods of 10-day lags, forecasting efficiency reverts back to just above 70%. However, for our longest period of interest, the 21-day lags, we see a dramatic reduction in R^2 to only about 40%. For time horizons of 1-5 days, our results are in contrast with Christoffersen and Diebold (2000) who suggest the accuracy of volatility forecasts decreases as they increase the time horizon. However, our findings are in line with their results suggesting that the efficiency of volatility forecasts beyond the 10-day time horizon rapidly decreases. In these initial regressions, the vix has slightly higher R^2 values which suggests that the level of the vix has a modestly higher forecasting efficiency over our volatility measure.

We further examine the forecasting efficiency of the two volatility measures by investigating their n-day differences with r_vol . To do so, we create an algorithm that determines the first two moments of n-day differences between vix and r_vol as well as $vcme$ and r_vol . Our findings are summarized in the following graphs below.




Figure 6 – n-day differences between VIX and actual volatility




Figure 7 – n-day differences between VCME and actual volatility

The most obvious take-away from here is also in line with earlier studies. As we increase the time-horizon, the average differences between implied volatility (vix, vcme) and actual volatility increase steadily. In addition, these results confirm that the 5-day time period seems to be a turning point when the variation of differences increases as well. Although standard deviations of n-day differences are generally higher for vcme, we find that average differences between vcme and r_vol are lower than those of the vix. From 1- to 5-day periods average vix differences are more than twice as high as average vcme differences. At the 21-day time horizon, average vix differences are 0.007 whereas average vcme differences are only 0.0049. A risk-seeking investor may be more akin to use our volatility measure to fine-tune volatility forecasts using average n-day differences. Here, our measure has an advantage over the vix since n-day differences for vcme are substantially lower than those of the vix, especially for the shorter time horizons when vcme differences are only about half of those for the vix, albeit with the caveat that the swings in vcme prices may be slightly larger than those of the vix.

To summarize the findings of our first research questions, we conclude that the vix and vcme have very similar regression results. Forecasting efficiency of future volatility is highest at the 5-day time interval where approximately 75% of the variations in future volatility can be explained by each of the models. However, there is a stark decline at the 21-day time interval

when R^2 drops to only about 40%. The regressions show a slight edge of the vix over vcme. Still, our measure appears to have smaller differences with future realized volatility at all time horizons, making vcme a viable alternative for an implied volatility measure. From a risk management perspective, particularly for an institution that is more capable of managing slightly larger standard deviations, our measure may be more attractive, if the n-day differences method fits within their overall risk management models.

VIX versus VCME: Which index provides better estimates of future S&P 500 returns?

While an accurate volatility forecast is desirable and, in the case of some financial institutions a daily required procedure, we submit that it may be more beneficial to find out how implied volatility measures relate to stock returns. This leads to our second research question and we use similar models to examine the effects of vix and vcme on S&P 500 log returns. Initially, we test how the levels of vix and vcme relate to stock returns. The regressions take on the following format:

$$\ln r_{n_spx_t} = \beta_0 + \beta_1 vix_{t-n} + \varepsilon_t \quad (8)$$

$$\ln r_{n_spx_t} = \beta_0 + \beta_1 vcme_{t-n} + \varepsilon_t \quad (9)$$

	(1)	(2)	(3)	(4)
VARIABLES	$\ln r_{1_spx}$	$\ln r_{5_spx}$	$\ln r_{10_spx}$	$\ln r_{21_spx}$
vix	-0.0284*** (0.00968)	-0.113*** (0.0180)	-0.179*** (0.0220)	-0.288*** (0.0269)
Constant	0.00531*** (0.00153)	0.0216*** (0.00282)	0.0348*** (0.00346)	0.0579*** (0.00422)
Observatio	1,256	1,252	1,247	1,236
R-squared	0.029	0.102	0.155	0.213

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 7 – Regressions of future spx log returns on vix.

	(1)	(2)	(3)	(4)
VARIABLES	lnr1_spx	lnr5_spx	lnr10_spx	lnr21_spx
vcme	-0.0294*** (0.00993)	-0.115*** (0.0180)	-0.177*** (0.0224)	-0.284*** (0.0266)
Constant	0.00543*** (0.00156)	0.0216*** (0.00280)	0.0341*** (0.00348)	0.0566*** (0.00411)
Observatio	1,256	1,252	1,247	1,236
R-squared	0.032	0.108	0.156	0.212

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 8 – Regressions of future spx log returns on vcme.

We find that both sets of regressions have statistically significant results at the 1% level. Nevertheless, we observe an interesting change in terms of the goodness of fit of these models. Whereas in our earlier regressions, we observed a decline in R^2 values when we increased the time horizon. Here we notice quite the opposite. At 1-day time horizons, only about 3% of the variations in stock returns can be explained by the models. However, as we increase the time horizon, R^2 increases to about 21%. We also notice that our measure has a slight edge over vix, albeit the differences are not showing a meaningful improvement of the goodness of fit.

In addition, we test the impact of changes in vix and changes in vcme on contemporaneous and inter-temporal SPX returns. To test contemporaneous effects, the models are as follows:

$$\lnr_n_spx_t = \beta_0 + \beta_1 d_n_vix_t + \varepsilon_t \quad (10)$$

$$\lnr_n_spx_t = \beta_0 + \beta_1 d_n_vcme_t + \varepsilon_t \quad (11)$$

VARIABLES	(1) lnr1_spx	(2) lnr5_spx	(3) lnr10_spx	(4) lnr21_spx
d1vix	-0.514*** (0.0181)			
d5vix		-0.570*** (0.0161)		
d10vix			-0.597*** (0.0172)	
d21vix				-0.648*** (0.0169)
Constant	0.000356** (0.000144)	0.00176*** (0.000303)	0.00354*** (0.000405)	0.00758*** (0.000564)
Observations	1,256	1,252	1,247	1,236
R-squared	0.735	0.741	0.720	0.704

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 9 – Regressions of future spx log returns on n-day vix changes.

VARIABLES	(1) lnr1_spx	(2) lnr5_spx	(3) lnr10_spx	(4) lnr21_spx
d1vcme	-0.434*** (0.0184)			
d5vcme		-0.499*** (0.0170)		
d10vcme			-0.552*** (0.0176)	
d21vcme				-0.614*** (0.0174)
Constant	0.000363** (0.000160)	0.00180*** (0.000335)	0.00357*** (0.000433)	0.00760*** (0.000591)
Observations	1,256	1,252	1,247	1,236
R-squared	0.675	0.682	0.681	0.675

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 10 – Regressions of future spx log returns on n-day vcme changes.

As before, all variables are significant at the 1% level. For contemporaneous effects, the models testing changes in vix on spx returns appear to have a better fit than vcme. Similar to our earlier studies, the 5-day time interval performs best wherein 74% of the changes in spx returns

can be explained by the model. *vcme* fairs slightly worse with an R^2 of 68.2%, however, the 5-day window seems to be the optimal interval within the time periods we examined.

While understanding these contemporaneous relationships is important, inter-temporal effects may be much more relevant when it comes to formulating certain trading and/or risk management strategies. Therefore, we examine inter-temporal effects of changes in implied volatility on subsequent *spx* returns. To test inter-temporal effects, the models are as follows:

$$\ln r_{5_spx_t} = \beta_0 + \beta_1 d1_vcme_t + \varepsilon_t \quad (12)$$

$$\ln r_{10_spx_t} = \beta_0 + \beta_1 d1_vcme_t + \varepsilon_t \quad (13)$$

$$\ln r_{21_spx_t} = \beta_0 + \beta_1 d1_vcme_t + \varepsilon_t \quad (14)$$

$$\ln r_{10_spx_t} = \beta_0 + \beta_1 d5_vcme_t + \varepsilon_t \quad (15)$$

$$\ln r_{21_spx_t} = \beta_0 + \beta_1 d5_vcme_t + \varepsilon_t \quad (16)$$

$$\ln r_{21_spx_t} = \beta_0 + \beta_1 d10_vcme_t + \varepsilon_t \quad (17)$$

For brevity, we only show the model specifications for *vcme* but the same models are applied to *vix* as well.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	<i>lnr5_spx</i>	<i>lnr10_spx</i>	<i>lnr21_spx</i>	<i>lnr10_spx</i>	<i>lnr21_spx</i>	<i>lnr21_spx</i>
<i>d1vcme</i>	-0.338*** (0.0725)	-0.236** (0.112)	-0.139 (0.122)			
<i>d5vcme</i>				-0.344*** (0.0431)	-0.261*** (0.0562)	
<i>d10vcme</i>						-0.376*** (0.0364)
Constant	0.00185*** (0.000566)	0.00361*** (0.000756)	0.00809*** (0.00103)	0.00356*** (0.000687)	0.00794*** (0.00100)	0.00780*** (0.000945)
Observations	1,252	1,247	1,236	1,247	1,236	1,236
R-squared	0.091	0.027	0.004	0.196	0.057	0.160

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 11 – Inter-temporal regressions of future *spx* log returns on *n*-day *vcme* changes.

VARIABLES	(1) lnr5_spx	(2) lnr10_spx	(3) lnr21_spx	(4) lnr10_spx	(5) lnr21_spx	(6) lnr21_spx
d1vix	-0.396*** (0.0792)	-0.294** (0.116)	-0.172 (0.129)			
d5vix				-0.407*** (0.0379)	-0.302*** (0.0516)	
d10vix						-0.407*** (0.0351)
Constant	0.00185*** (0.000564)	0.00360*** (0.000754)	0.00808*** (0.00103)	0.00355*** (0.000673)	0.00792*** (0.00100)	0.00781*** (0.000939)
Observations	1,252	1,247	1,236	1,247	1,236	1,236
R-squared	0.098	0.033	0.006	0.229	0.065	0.173

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 12 – Inter-temporal regressions of future spx log returns on n-day vix changes.

Testing inter-temporal effects, we can now see how the results might be more relevant in an applied setting. For short-term inter-temporal effects, the 1-day changes in implied volatility have very little if any relevance on future stock returns 5-days, 10-days and 21-days out. While the effects on 5-day to 10-day spx returns are still significant, the further out our estimates go, the smaller the coefficients and at the 21-day window, the model is no longer significant.

For 5-day changes in vcme, the subsequent 10-day spx return window has the highest R^2 with almost 20% but as we increase the time interval to subsequent 21-day returns, the goodness of fit measure decreases drastically to only 6.5%.

Inter-temporal vix regressions show very similar results albeit with slightly higher R^2 values. Similar to our earlier studies, we find that there is limited effectiveness in forecasting efficiency beyond the 10-day window. Using these simple models, we find observing changes in implied volatility over a 5-day time horizon gives us the best goodness of fit when it comes to stock returns 10 days out with the caveat however, that only about 20% (22.9% for vix) of the variation in future stock returns are explained by 5-day changes in vcme. Overall, the vix appears to have a small edge over vcme when it comes to the forecasting efficiency of subsequent stock returns with these direct comparisons.

Assessing the findings of our second research questions, we observe rather similar results in a direct comparison of the two implied volatility measures. We notice that the level of vcme has a modest edge over the vix in terms of forecasting future stock returns up to five days out. For future 10-day and 21-day returns, both volatility measures show nearly identical results. We are encouraged by the fact that our measure fares slightly better than the vix for up to five days out. In addition, we notice that the goodness of fit increases when we lengthen the time period returns so that at 21-day returns over 21% of the changes in spx returns can be explained by the model.

We also note that changes in our implied volatility measures are all highly significant in terms of contemporaneous spx returns but these concurrent relationships may not be all that helpful for a risk manager. When we examine inter-temporal effects, we notice that our beta coefficients are much smaller and R^2 values, particularly for effects of 1-day changes, lead us to question the effectiveness of a return forecast. In terms of finding an optimal forecasting window, we notice that 5-day changes in vix have the highest impact on 10-day spx returns at an R^2 of 22.9% ahead of our vcme measure at 19.6%.

CONCLUSIONS

We began this journey with the goal of finding a more efficient measure of volatility and risk that did not have to rely on the VIX. Our findings were quite interesting, in that we showed that both measures followed similar trends and patterns. These measures were highly efficient at forecasting short-term future volatility, being able to explain up to 75% of that variation. While both measures lose their effectiveness over longer periods of time, the average n-day differences between VCME and actual volatility are consistently lower than those of the VIX at any of the given time frames.

When it comes to forecasting stock returns, we see an interesting reversal of effects where longer time windows, of returns and changes in VCME, have a higher explanatory power. Both measures fare relatively well for contemporaneous effects, with highly significant results, specifically at the 5-day time interval, with more modest inter-temporal effects, specifically for shorter 1-day windows, as well as windows over 10-days.

Our main take-away from this study is that while both measures do behave similarly, there is a clear attractiveness, for institutions, to using our measure for the following reasons: Higher volatility might suggest that our measure is more responsive to changes in underlying market conditions. At the same time, vcme had consistently lower average n-day differences to actual volatility than those observed by the vix. This would suggest that a financial institution seeking responsive risk measures, while on average generating less deviations from actual volatility, might find our volatility measure more beneficial.

In closing, we suggest some extensions to this study as follows: Given that our dataset captures a time frame of an unambiguous bull market, we propose to do a follow-up study that includes daily data from the financial crisis period of 2008/09 which unfortunately was not available from our data vendor. Examining how the two volatility measures compared during the crisis period may give additional insights to help distinguish the pros and cons of these risk measures.

Similarly, we would like to extend our methodology to intra-day data. Again, this would necessitate the purchase of data from a different vendor and a drastic increase of computing requirements. Lastly, we would like to examine other market variables in addition to implied volatility measures with a view to generate better volatility forecasts but also to improve predictions on underlying asset returns.

ACKNOWLEDGEMENTS

We would like to thank John Rutledge, Bluford Putnam and Kassie Davies for the support and encouragement for our research.

Funding: This work was supported by the CME GROUP FOUNDATION, 20 South Wacker Drive, Chicago, Illinois 60606

REFERENCES

- Black, Fischer**, 1989, How We Came Up with the Option Formula, *The Journal of Portfolio Management*, Vol. 15, pp. 4-8.
- Black, Fischer and Scholes, Myron**, 1973, The Pricing of Options and Corporate Liabilities, *The Journal of Political Economy*, Vol. 81, No. 3 (May - Jun., 1973), pp. 637-654
- Blair, B.J., S. Poon, and S.J. Taylor**, 2000, Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns, *Journal of Econometrics*, 105, 5–26.
- Bollerslev, T.**, 1986, Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics*, 31, 307-327.
- Bollerslev, T.**, 2007 Glossary to ARCH (GARCH). *Duke University and NBER*.
- CBOE**, 2014, THE CBOE VOLATILITY INDEX® - VIX®, Chicago Board Options Exchange.
- CBOE**, How is VIX Calculated? Step-by-Step:
<http://cfe.cboe.com/education/vixprimer/About.aspx>
- Christoffersen, Peter F. and Diebold Francis X.**, 2000, How Relevant is Volatility Forecasting for Financial Risk Management? *Review of Economics and Statistics*, Vol. 82, No. 1 pp. 12-22
- Fama, Eugene F.** 1970. Efficient Capital Markets: A Review of Theory and Empirical Work, *Journal of Finance*. 25:2, pp. 383-417.
- Fama, Eugene F. and French, Kenneth R.**, 2004, The Capital Asset Pricing Model: Theory and Evidence, *The Journal of Economic Perspectives*, Vol. 18, No. 3, pp. 25-46
- Giot, Pierre**, 2002, Implied Volatility Indices as Leading Indicators of Stock Index Returns? CORE Discussion Paper No. 2002/50.
- Granger, Clive W. J. and Poon, Ser-Huang**, 2003, Forecasting Volatility in Financial Markets: A Review, *Journal of Economic Literature* Vol. XLI (June 2003) pp. 478–539
- Hull, John C.**, 2012, Options Futures & other Derivatives, 8th edition, *Prentice Hall*
- Knight, Frank H.**, 1921, Risk, Uncertainty, and Profit, *Library of Economics and Liberty*
- Kownatzki, Clemens**, 2015, “Examination of implied volatility as a proxy for financial risk” *The Claremont Graduate University*, Publication Number 3717039
- Kownatzki, Clemens**. 2016, “How Good is the VIX as a Predictor of Market Risk?” *Journal of Accounting and Finance* Vol. 16(6)

Mandelbrot, Benoit, 1963, The Variation of Certain Speculative Prices, *The Journal of Business*, Vol. 36, No. 4 (Oct., 1963), pp. 394-419

Martens, Martin and Zein, Jason, 2002, Predicting Financial Volatility: High-Frequency Time-Series Forecasts Vis-a-Vis Implied Volatility (February 21, 2002). Available at SSRN: <http://ssrn.com/abstract=301382>

Merton, Robert C., 1976, Option Pricing When Underlying Stock Returns are Discontinuous, *Journal of Financial Economics* 3 (January-February 1976): 125-144

Merton, Robert C., 1973, Theory of Rational Option Pricing, *Bell Journal of Economics and Management Science* 4, no. 1: 141-183. (Chapter 8 in Continuous-Time Finance.)

Natenberg, Sheldon, 1994, Option Volatility & Pricing, *McGraw-Hill* pp 65-66

Rachev, Svetlozar T. and Stoyanov, Stoyan V. and Biglova, Almira and Fabozzi, Frank J., 2005, An Empirical Examination of Daily Stock Return Distributions for U.S. Stocks, Studies in Classification, *Data Analysis, and Knowledge Organization*, pp 269-281

Stock, James H. and Watson, Mark W., 2011, Introduction to Econometrics 3rd edition, *Addison-Wesley*

Whaley, Robert E., 1993, Derivatives on Market Volatility: Hedging Tools Long Overdue, *Journal of Derivatives* 1 (Fall 1993), pp. 71—84.

Whaley, Robert E., 2000, The Investor Fear Gauge, *Journal of Portfolio Management* 26, pp. 12-17.

Whaley, Robert E., 2009, Understanding the VIX, *The Journal of Portfolio Management*