How Good is the VIX as a Predictor of Market Risk?

Clemens Kownatzki
Pepperdine University

Volatility is a metric widely used to estimate financial risk. The VIX is an index derived from S&P 500 options prices designed to estimate the market's expected 30-day volatility. Robert Whaley, the creator of the VIX, argued that it provided a cost-effective way to hedge risk but we question Whaley's underlying assumption in this paper. We examine the VIX and implied volatility as a proxy for risk. Our studies show that the VIX consistently over-estimates actual volatility in normal times but it underestimates volatility in times of market crashes and crises making it unsuitable for many risk-management applications.

INTRODUCTION

The concept of risk has fascinated financial economists for a long time. *Volatility*, often measured as the standard deviation of historical returns, has been used as a proxy for risk at least since Markowitz (1952). The forward-looking concept *implied volatility,* a by-product of the ground-breaking work of Black-Scholes (1973) and Merton (1973), has been the preferred metric for risk-modeling more recently and is considered a more efficient predictor of risk. In 1993, the Chicago Board Options Exchange introduced a *Volatility Index*, also known as the VIX[1] (CBOE 2009). Since the recent crisis of 2008/09, financial news media routinely report the VIX along with stock market indices.

The crisis also exposed limitations of accepted risk metrics, including the VIX, during periods when prices fluctuate much more than anticipated by standard risk models. Recent research shows that extreme price changes, referred to as "Black Swan" events (Taleb 2007), are a regular feature of markets[2]. Indeed, the financial crisis sparked a race among researchers to explain why periods of extreme price changes, or 'outliers', occur far more often than generally accepted models would suggest[3].

The objective of this paper is an empirical examination of *implied* and *historical volatility*. Specifically, we analyze the ability of the VIX to estimate future risk. We will show that the VIX systematically overestimates actual volatility in non-crisis periods, and it underestimates volatility in times of financial crises which suggests that the VIX is unsuitable for many risk management applications. We will conclude that the expectation of future uncertainty, implied by the VIX, is overstated. Our studies suggest that option traders have known all along that implied volatility derived from option prices is over-priced, which is why they are typically net-sellers of options.

This paper is organized as follows: Section two reviews the existing literature and provides theoretical perspectives on volatility and the VIX. Section three discusses the statistical properties of the VIX. Section four is an empirical examination of the VIX as a predictor of actual volatility. In section five, we summarize our findings, draw conclusions from our results and close with recommendations for further research.

LITERATURE REVIEW

Modern Finance Theory has used a number of metrics to measure risk. Harry Markowitz (1952) and James Tobin (1958) associated risk with the variance of portfolio returns over time. Markowitz (1952) showed how a risk-averse investor could use estimates of expected return and standard deviation to choose an optimal portfolio. Based upon Markowitz's work, Treynor (1961;1962), Sharpe (1964), Lintner (1965a,b) and Mossin (1966), developed the Capital Asset Pricing Model (CAPM) which defines a capital market in equilibrium. Treynor's work inspired Fischer Black to develop an early version of the Black-Scholes formula based on the CAPM model[4]. With the help of Myron Scholes and Robert Merton, Fischer Black transformed this early version into one of the most widely used models in finance Black-Scholes-Merton (BSM) model. Option pricing models had been used by option traders (Haug and Taleb, 2009; Mixon, 2008) at least a century prior to the seminal papers of Black/Scholes (1973) and Merton (1973). But it may have been the elegance and simplicity of the closed-form solution of the BSM formula that led to the rapid adoption of their model by practitioners as well as academics.

Even though Benoit Mandelbrot showed that the presence of leptokurtosis[5] in empirical return distributions was indisputable (Mandelbrot 1963), many finance models still held on to the assumption that financial asset returns were normally distributed. With the assumptions of general equilibrium theory and rational agents who know the probability distributions of future outcomes, mainstream economics and finance maintained that risk could be securitized and thereby priced and hedged[6]. But modern finance theory has had its share of critics too. Eugene Fama, the father of the Efficient Market Hypothesis (Fama 1970), 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 recent studies (Rachev et al 2005) suggest that there is little evidence in support of a normal distribution for most financial assets. Their empirical examinations conclude that financial return series are heavy-tailed and possibly skewed. Still, Gaussian normality, albeit in log-return space, survived as an essential part of modern finance a few more decades, but the financial crisis of 2008/09 caused many academics as well as practitioners to re-evaluate the main assumptions of their models.

Long before the development of modern portfolio theory, Frank Knight differentiated between risk and uncertainty (Knight 1921). He dismissed risk metrics such as standard deviation and suggested that the things we can measure were not helpful in understanding uncertainty, the true measure of risk. Although Frank Knight's work was mainly conceptual, his theories influenced many economists including the Nobel Laureates Friedman, Stigler and Buchanan. (Nobelprize.org 2014). Quantifiable metrics, however, were indispensable if the world of finance was to maintain an appearance of safety based on sound science. The soundness of that science, in view of the most recent financial crisis, has been the subject of much debate. We will contribute to the debate by examining whether one of these risk metrics, implied volatility, serves as a reliable proxy for risk.

Not long after the stock market crash of October 1987, Menachem Brenner and Dan Galai proposed the concept of a volatility index that could help formulate risk-management strategies (Brenner and Galai 1989). In 1992, the Chicago Board Options Exchange (CBOE) retained Robert Whaley to develop a volatility index derived from traded option prices on the S&P 100 Index (OEX). Whaley suggested that the VIX provided a "reliable estimate of expected short-term market volatility" (Whaley 1993) and remarked that the main attraction was its forward-looking nature, "measuring volatility that investors expect to see" (Whaley 2009).

Volatility has a number of properties that sets it apart from other financial instruments. Numerous studies confirmed a negative correlation between the VIX and equity returns (Brenner and Galai 1989; Anderson T. G. et al 2001; Granger and Poon 2003, Whaley 2009). Other researchers documented *mean reversion* (Hafner 2003; Bali et al 2006; Wong and Lo 2008; Fouque et al 2008), *asymmetric volatility* (Whaley 2009; Aboura and Wagner 2015), and *volatility clustering* (Cont 2005). The notion of volatility as an asset class was recognized as well (Grant et al 2007; Fieldhouse 2012).

The 1996 Basle Accord required banks to calculate Value at Risk (VaR)[7], which made volatility forecasting a "compulsory risk-management exercise for banks" (Granger and Poon 2003). Robert Engle

(Engle 1982) relaxed the assumption that expected volatility was constant and treated volatility as a process instead (Chen et al 2005). His seminal work[8] on Auto Regressive Conditional Heteroscedasticity (ARCH) was the first in a series of dynamic volatility forecasting models. The Generalized Auto Regressive Conditional Heteroscedasticity model (GARCH)[9], a direct successor of ARCH, was developed by Tim Bollerslev (Bollerslev 1986), a student of Robert Engle (Engle 2004). GARCH (1,1) is considered the most popular model for many financial time series (Granger and Poon 2003) but there are many versions today. Tim Bollerslev compiled an extensive glossary of over 150 GARCH-type models (Bollerslev 2007) but the severity of the recent financial crisis suggests that none of these models produced a reliable predictor of risk.

Granger and Poon provide a concise review of volatility forecasting models in a survey that aggregates 93 published papers examining the performance of different volatility models (Granger and Poon 2003). Although they were getting mixed results among various competing models, they ranked implied volatility from options using the Black-Scholes model ahead of Historical volatility and ahead of GARCH (Granger and Poon 2003). Still earlier studies by Blair, Poon and Taylor (2000) ranked the VIX highest in terms of providing the most accurate out-of-sample forecasts. Martins and Zein (2002) demonstrate that volatility forecasts have higher explanatory power at shorter time horizons while Christoffersen and Diebold (2000) suggest that the accuracy of equity and foreign exchange volatility forecasts decreases rapidly from 1 to 10 days forecast horizon. While they agree that short-horizon volatility is highly forecastable, their results indicate 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).

Whaley (1993) argued that the VIX is a cost-effective way to hedge risk because it provides a reliable estimate of expected short-term volatility. However, conflicting results from numerous studies give us motivation to question the robustness of the underlying theoretical framework. Our work in the following chapter will lead us to question Whaley's underlying assumption.[10]

STATISTICAL PROPERTIES OF THE VIX

In 1993, the Chicago Board Options Exchange (CBOE) introduced a contract for the volatility Index (VIX) based on the S&P 100 Index (OEX) with the intent to measure the market's expected 30-day implied volatility (CBOE 2009). OEX options represented about 75% of the volume of Index options traded at the time, which might explain why the narrower OEX was chosen over the broader S&P 500 (Whaley 1993). Using a different methodology derived from options on the more widely-watched S&P 500 Index (SPX), a new VIX contract was introduced in 2003 (CBOE 2009). The new VIX was now based on fair value of future variance instead of the old derivation from the Black-Scholes equation. The old VIX was subsequently renamed to VXO. VIX and VXO remained highly correlated at over 98% despite their different derivations. The derivation and generalized formula of the VIX calculation[11] is provided by CBOE (CBOE 2009). In 2014, the CBOE made a further change in the VIX to include weekly options on the S&P 500, reflecting the increased trading volume of weekly options expirations.

To get a first sense of how the VIX has evolved over time, we reference *Figure 1* showing daily closing prices of the S&P 500 (SPX) as well as the VIX and VXO between January 3, 1986 and December 31, 2014. Visual inspection shows two extreme values for the VIX coinciding with the market crash in October 1987 and the sub-prime mortgage crisis in 2008. The VIX exhibits intermittent periods of small and large variations at irregular intervals. Sudden upward spikes are followed by relatively slower reversion to the mean. Peak values of the VIX appear to roughly coincide with relative market bottoms suggesting a negative correlation with the S&P 500. We now proceed to verify this initial insight through rigorous empirical analysis.

FIGURE 1
DAILY CLOSING PRICES OF S&P 500 VERSUS VIX AND VXO

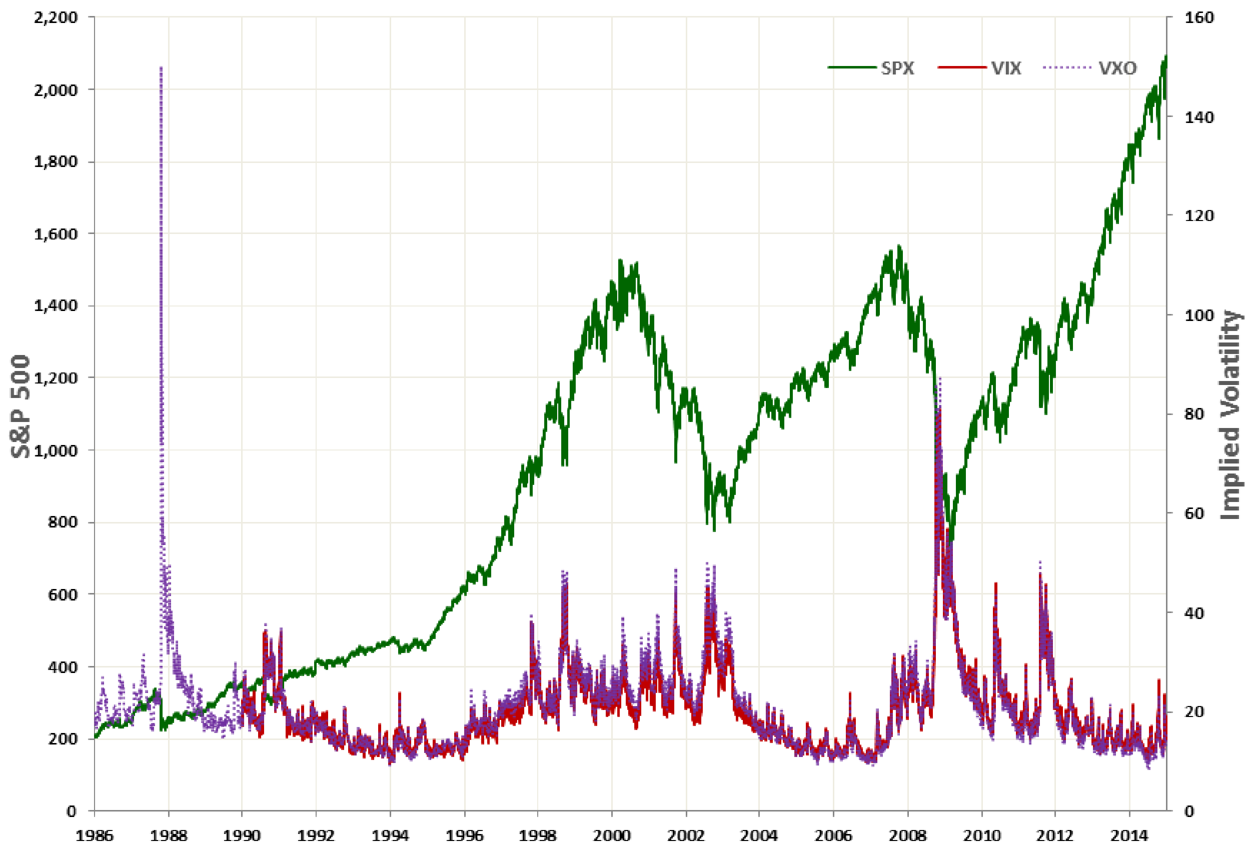

Our studies focus on the period of January 1990 until December 2014 and we have gathered daily closing prices for our variables of interest. *Figure 2* shows a histogram of daily VIX levels which resemble a Levy or power-law distribution rather than a normal distribution. Similarly, the distributions of VIX log returns (*Figure 3*) do not appear to be normal. We observe heavier tails than those expected from Gaussian normality. Power laws are frequently found in financial time series including currency, commodity, and stock market returns (Gabaix 2009) but they are also apparent in time series of VIX levels and VIX returns as our data suggest.

FIGURE 2
FREQUENCY DISTRIBUTION OF VIX LEVELS VS. NORMAL DISTRIBUTION

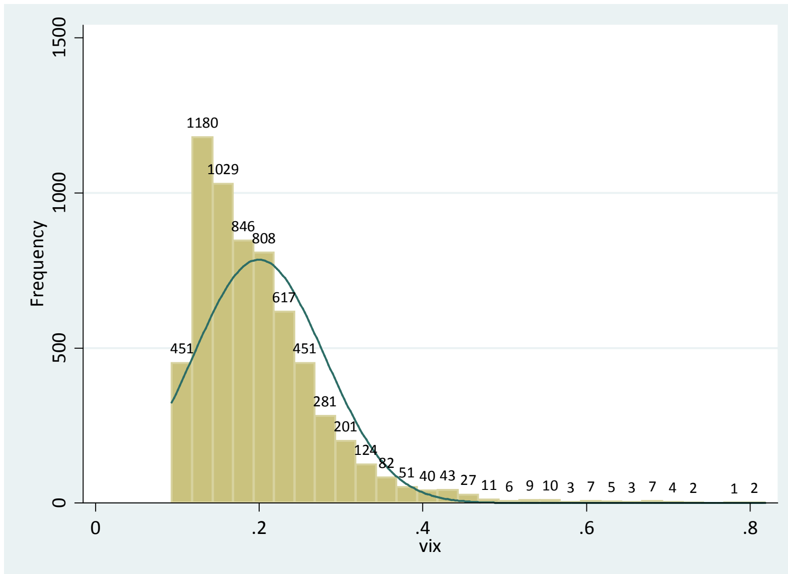

FIGURE 3
HISTOGRAM OF 1-DAY LOG RETURNS OF THE VIX VS. NORMAL DISTRIBUTION

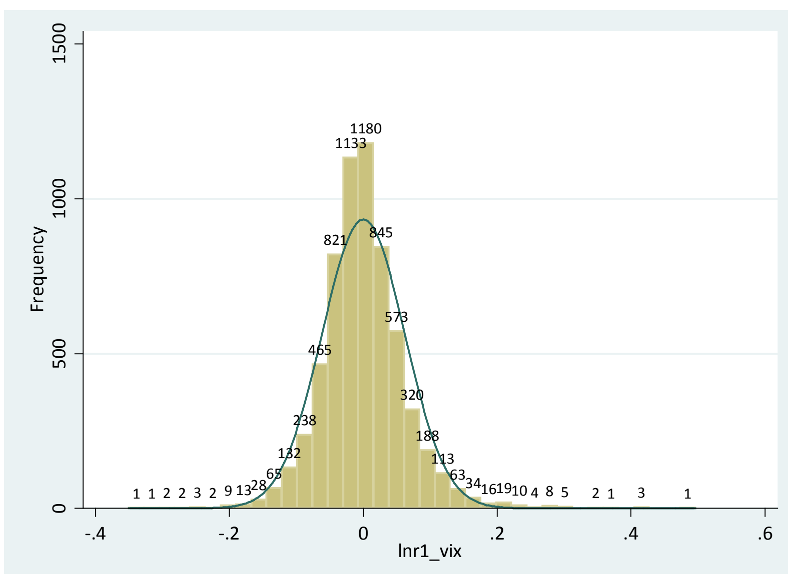

Table 1 shows the summary statistics of the first four moments along with the percentile ranges of VIX distributions between 1990 and 2014. The long-term average is 0.1995 with an interquartile range[12] of 0.0926 indicating the VIX closed between 0.1417 and 0.2343 about 50% of the time. During periods of crisis however, the average VIX often remains outside its long-term interquartile range. The widest interquartile range occurred in 2008 when the VIX closed between 0.2158 and 0.4 fifty percent of the time with an annual mean of 0.3269. The second largest interquartile range occurred in 2009 when the VIX closed between 0.2428 and 0.3943 50% of the time (a range of 0.1515).

Higher moments of the VIX distribution show similar variations. Skewness fluctuates between 0.1783 and 2.3033. It is positive and significant for all of the years except 1992, 1996 and 2000, suggesting that the underlying distribution is asymmetrical.

TABLE 1
VIX MOMENTS & PERCENTILE RANGES FOR DAILY CLOSING LEVELS

											Normal Range	
Year	**# days**	**Mean**	**Std. Dev.**	**Skewness**	**Kurtosis**	**5%**	**25%**	**50%**	**75%**	**95%**	**50%**	**90%**
1990	253	0.2306	0.0474	0.3872	2.3018	0.1641	0.1936	0.2257	0.2716	0.3079	0.0780	0.1438
1991	253	0.1838	0.0368	2.3033	9.6425	0.1490	0.1602	0.1744	0.1974	0.2539	0.0372	0.1049
1992	254	0.1545	0.0212	0.2320	2.3034	0.1229	0.1370	0.1536	0.1685	0.1896	0.0315	0.0667
1993	253	0.1269	0.0133	0.4125	3.0925	0.1087	0.1171	0.1243	0.1356	0.1502	0.0185	0.0415
1994	252	0.1393	0.0207	0.6460	4.0558	0.1115	0.1211	0.1386	0.1550	0.1701	0.0339	0.0586
1995	252	0.1239	0.0097	0.6955	3.3591	0.1112	0.1162	0.1230	0.1297	0.1425	0.0135	0.0313
1996	254	0.1644	0.0194	0.3031	3.1157	0.1334	0.1520	0.1625	0.1746	0.2011	0.0226	0.0677
1997	253	0.2238	0.0414	1.7877	5.9881	0.1856	0.1970	0.2095	0.2383	0.3209	0.0413	0.1353
1998	252	0.2560	0.0686	1.1409	3.4340	0.1826	0.2052	0.2315	0.2873	0.4047	0.0821	0.2221
1999	252	0.2437	0.0288	0.4042	2.8624	0.1998	0.2229	0.2411	0.2621	0.2967	0.0392	0.0969
2000	252	0.2332	0.0341	0.1783	2.4582	0.1788	0.2071	0.2324	0.2594	0.2893	0.0523	0.1105
2001	248	0.2575	0.0478	1.1132	4.0698	0.2009	0.2211	0.2426	0.2862	0.3495	0.0651	0.1486
2002	252	0.2729	0.0691	0.5100	2.1953	0.1843	0.2115	0.2639	0.3250	0.3969	0.1136	0.2126
2003	252	0.2198	0.0524	1.0086	2.6659	0.1647	0.1829	0.1986	0.2496	0.3262	0.0667	0.1615
2004	252	0.1548	0.0192	0.4697	3.2903	0.1267	0.1430	0.1533	0.1655	0.1893	0.0226	0.0626
2005	252	0.1281	0.0147	0.6640	3.1451	0.1077	0.1167	0.1252	0.1365	0.1563	0.0198	0.0486
2006	251	0.1281	0.0225	1.6439	6.0938	0.1053	0.1135	0.1200	0.1365	0.1774	0.0230	0.0721
2007	251	0.1754	0.0536	0.4918	2.1695	0.1034	0.1311	0.1643	0.2168	0.2649	0.0857	0.1615
2008	253	0.3269	0.1638	1.2846	3.2906	0.1859	0.2158	0.2510	0.4000	0.6780	0.1842	0.4921
2009	252	0.3148	0.0908	0.6825	2.1979	0.2112	0.2428	0.2857	0.3943	0.4756	0.1515	0.2644
2010	252	0.2255	0.0527	1.2748	4.9425	0.1647	0.1833	0.2172	0.2527	0.3373	0.0694	0.1726
2011	252	0.2420	0.0814	0.7640	2.3819	0.1581	0.1740	0.2072	0.3157	0.3902	0.1417	0.2321
2012	250	0.1780	0.0254	0.8530	3.5755	0.1438	0.1573	0.1752	0.1905	0.2227	0.0332	0.0789
2013	252	0.1423	0.0174	1.2310	4.3892	0.1231	0.1299	0.1375	0.1498	0.1750	0.0200	0.0519
2014	251	0.1415	0.0260	1.7274	6.8825	0.1133	0.1232	0.1367	0.1511	0.1944	0.0279	0.0811
Overall	6301	0.1995	0.0801	2.0610	10.3932	0.1152	0.1417	0.1816	0.2343	0.3411	0.0926	0.2259
Bull Mkt	5250	0.1836	0.0628	1.3198	4.8792	0.1140	0.1351	0.1673	0.2149	0.3122	0.0798	0.1982
Bear Mkt	1051	0.2790	0.1055	2.1439	8.0352	0.1848	0.2134	0.2444	0.2955	0.5205	0.0821	0.3357

Kurtosis varies between 2.1695 and 9.6425 among the years. 16 out of the 25 years of observations showed fat tails, i.e., kurtosis greater than three. The three years with highest kurtosis levels are 1991 (9.64), 2014 (6.88) and 2006 (6.09). Due to some statistical properties of volatility, kurtosis for the entire data set is greater than the values in each given year. For longer time periods, the VIX approaches a mean of about 20%. In each given year, however, the annual mean value could be significantly different from its long-term average. Since kurtosis measures how much of the variance arises from extreme values compared to its mean, longer time periods are likely to have more extreme values while the VIX over longer time periods is mean reverting. Several studies (e.g. Hafner 2003; Bali et al 2006; Wong and Lo 2008; Fouque et al 2008) have identified a mean-reversion property of the VIX. We find support for this notion within our data set, since most of the sudden upticks in volatility are rather short-lived.

Long before the inception of the VIX, researchers have identified *volatility clustering*, the notion that "large changes tend to be followed by large changes - of either sign - and small changes tend to be followed by small changes" Mandelbrot (1963). In addition to clustering, we are interested in the duration of elevated VIX levels as this directly affects risk-management strategies. In our data set, there was only a 5% chance that the VIX would be above 0.3411. The clustering property of volatility then suggests that an unusually high VIX level is likely to be followed by another unusually high day. Since regulatory risk-management procedures for banks require a volatility forecast using a 10-day value-at-risk at the 99 percent confidence level (BIS 2009), we wanted to examine how often elevated VIX levels appear longer than 10 consecutive days. There were 7 periods of ten or more consecutive trading days with VIX levels above 0.3411. In particular, during the period of September 26, 2008 until April 16, 2009, the VIX was above this elevated level, a period of 139 days. Clearly, the notion of a mean-reverting fear index was put to the test during this crisis period when the VIX remained at elevated levels much longer than typical market conditions would suggest. *Table 2* shows the time periods of these high VIX levels.

TABLE 2
PERIODS OF CONSECUTIVE TRADING DAYS OF THE VIX ABOVE 0.3411

From	To	# days
26-Sep-08	16-Apr-09	139
3-Sep-02	14-Oct-02	30
27-Aug-98	22-Sep-98	18
24-Sep-98	14-Oct-98	15
21-Sep-11	7-Oct-11	13
20-Apr-09	4-May-09	11
15-Jul-02	26-Jul-02	10

Some of the statistical properties of the VIX, in particular mean reversion and clustering, are unlike many other financial instruments. Given the non-Gaussian distribution of volatility, it would suggest that statistical inferences drawn from traditional mean and standard deviation estimates are probably not useful for risk management. We submit that traditional linear regressions are not helpful in determining the efficiency of the VIX as an estimator of future risk either. Instead, our examinations will focus on a pure empirical examination of the VIX using basic comparative studies to determine how reliable its estimates of expected short-term volatility are.

THE VIX AS A PREDICTOR OF ACTUAL VOLATILITY

For our studies, we define the main variables as follows: spx, vix, and r_vol representing the daily closing prices of the S&P 500 Index, the VIX, and realized (historical) volatility, respectively. d1_vix, d5_vix, d1_vix, d21_vix, d60_vix are the 1-, 5-, 10-, 21-, and 60-day forward looking VIX changes or VIX differences calculated as vix_{t+1} - vix_t. Log returns are calculated as ln(VIXt+1 / VIXt) for 1-day returns and ln(VIXt+21 / VIXt) for 21-day returns. lnr1_vix, lnr5_vix, lnr10_vix, lnr21_vix, lnr60_vix are the 1-, 5-, 10-, 21-, and 60-day forward looking log returns of the VIX. Log returns and differences for other variables are calculated with the same methodology. Lagged values of our variables are defined as L.#_of_days_(variable) so that a 21-day lagged value of the VIX is denoted as L.21(vix).

Although we don't run traditional regressions, for completeness, we tested for the presence of auto-correlation, non-stationarity and heteroskedasticity. Our tests confirm numerous earlier studies that suggest daily prices of stocks and stock indices such as the S&P 500 are auto-correlated and exhibit a strong non-stationary stochastic trend. However, the VIX, realized volatility, changes of the VIX and changes in realized volatility are not auto-correlated and do not display a stochastic trend *(Figure 4)*. Heteroskedasticity was present for all our variables of interest. We also note that the VIX has a correlation of 89% with realized volatility *(See Table 3)*.

TABLE 3
CORRELATION OF CONCURRENT LEVELS OF VIX AND R_VOL

```
. correlate vix r_vol, means
(obs=6301)
```

Variable	Mean	Std. Dev.	Min	Max
vix	.1994918	.0800668	.0931	.8086
r_vol	.156297	.0938717	.0482	.8536

	vix	r_vol
vix	1.0000	
r_vol	0.8900	1.0000

FIGURE 4
AUTO-CORRELATION OF VIX DIFFERENCES UP TO 40 LAGS

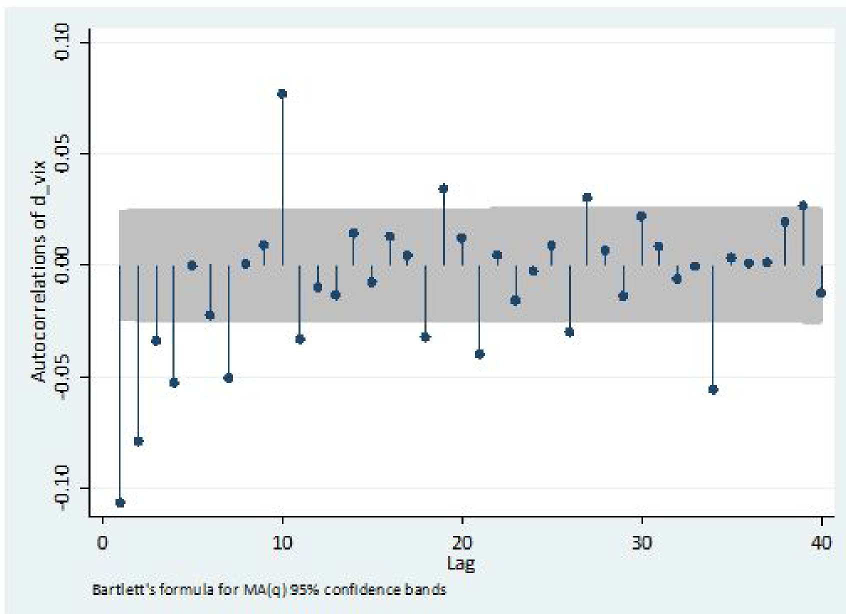

Next, we examine to what extent implied volatility exceeds realized volatility and identify periods when implied volatility is greater than realized volatility and when it is smaller. For this study, we begin with the same data set using daily observations from January 1990 until December 2014. Calculating the daily log returns of the S&P 500 as well as its 21-day standard deviation of returns we derive an annualized value of realized volatility (r_vol) that could be compared with the VIX. Finally, a lag of 21 trading-days is imposed on the VIX. This way, we can directly compare to what extent the implied volatility estimate via the VIX matched actual volatility 30 calendar days ahead.

Referencing *Tables 1, 4 and 5* we compare the data between the lagged VIX and realized volatility. For the entire time period, the lagged VIX exceeds realized volatility by about 430 basis points on average. The VIX also exceeds realized volatility in every given year except for 2008. During the crisis year of 2008, the VIX understates realized volatility by about 180 basis points on average. Variability and interquartile ranges are fairly similar with only small differences throughout the years. During bear market periods[13], the standard deviation of the VIX understates that of r_vol by almost 350 basis points and the correlation between VIX and r_vol reads just over 70% (*Table 6*). This is in line with earlier results which suggest that the forward looking quality of the VIX is reduced during periods of market declines.

More importantly, it would be of great interest to see how the differences between the VIX and realized volatility behave. Hence, we test the forecasting efficiency of the VIX with one simplifying assumption: *If the VIX was an efficient estimator of realized volatility, we would expect that a histogram of the differences between the 21-day lagged VIX and realized volatility would approximate a normal distribution with a mean of zero.*

TABLE 4
MOMENTS & PERCENTILE RANGES FOR DAILY LEVELS OF REALIZED VOLATILITY

											Normal Range	
Year	**# days**	**Mean**	**Std. Dev.**	**Skewness**	**Kurtosis**	**5%**	**25%**	**50%**	**75%**	**95%**	**50%**	**90%**
1990	253	0.1538	0.0433	0.7916	2.4668	0.1053	0.1187	0.1334	0.1784	0.2387	0.0597	0.1334
1991	253	0.1366	0.0008	-0.2281	2.2538	0.0842	0.1153	0.1370	0.1591	0.1828	0.0438	0.0986
1992	254	0.0994	0.0208	0.3262	2.5604	0.0667	0.0817	0.0993	0.1100	0.1415	0.0283	0.0748
1993	253	0.0851	0.0213	0.5040	2.6212	0.0546	0.0676	0.0834	0.1009	0.1221	0.0333	0.0675
1994	252	0.0962	0.0222	0.1096	2.4680	0.0620	0.0770	0.0992	0.1100	0.1358	0.0331	0.0738
1995	252	0.0757	0.0177	0.4695	2.1612	0.0528	0.0597	0.0746	0.0883	0.1075	0.0286	0.0547
1996	254	0.1121	0.0286	0.5923	2.5651	0.0735	0.0872	0.1087	0.1289	0.1678	0.0417	0.0943
1997	253	0.1724	0.0597	2.2358	7.4472	0.1192	0.1397	0.1551	0.1793	0.3542	0.0396	0.2350
1998	252	0.1849	0.0857	1.4419	4.3614	0.1029	0.1210	0.1580	0.2097	0.4116	0.0887	0.3087
1999	252	0.1820	0.0285	-0.3719	2.6362	0.1281	0.1660	0.1862	0.2004	0.2242	0.0344	0.0961
2000	252	0.2117	0.0605	-0.0770	2.4175	0.1047	0.1555	0.2201	0.2512	0.3221	0.0957	0.2174
2001	248	0.2096	0.0636	0.6847	2.0336	0.1394	0.1626	0.1795	0.2627	0.3210	0.1001	0.1816
2002	252	0.2436	0.0902	0.7768	2.7626	0.1226	0.1777	0.2237	0.3005	0.4302	0.1229	0.3076
2003	252	0.1673	0.0490	0.5138	2.4992	0.1037	0.1229	0.1609	0.2045	0.2649	0.0816	0.1612
2004	252	0.1103	0.0176	0.4725	2.9082	0.0862	0.0964	0.1091	0.1228	0.1399	0.0264	0.0537
2005	252	0.1012	0.0220	1.0439	3.3715	0.0746	0.0868	0.0963	0.1064	0.1506	0.0196	0.0760
2006	251	0.0977	0.0267	0.9517	2.6644	0.0693	0.0770	0.0875	0.1111	0.1529	0.0341	0.0836
2007	251	0.1473	0.0595	0.3068	1.8890	0.0648	0.0911	0.1365	0.1964	0.2496	0.1053	0.1848
2008	253	0.3446	0.2251	1.1205	2.6210	0.1438	0.1944	0.2306	0.5072	0.8026	0.3128	0.6588
2009	252	0.2596	0.1046	0.5637	2.0954	0.1324	0.1690	0.2155	0.3560	0.4514	0.1871	0.3190
2010	252	0.1696	0.0660	0.6050	2.6511	0.0758	0.1165	0.1628	0.2068	0.3018	0.0904	0.2260
2011	252	0.2048	0.1106	0.8057	2.8390	0.0737	0.1120	0.1572	0.2895	0.4488	0.1776	0.3751
2012	250	0.1267	0.0315	0.3764	2.2831	0.0837	0.1026	0.1192	0.1475	0.1838	0.0449	0.1001
2013	252	0.1104	0.0269	0.4386	2.2576	0.0749	0.0902	0.1034	0.1308	0.1595	0.0406	0.0846
2014	251	0.1053	0.0359	0.5171	2.3925	0.0572	0.0789	0.0998	0.1328	0.1722	0.0539	0.1150
Overall	6301	0.1563	0.0939	2.9446	16.6954	0.0676	0.0982	0.1333	0.1838	0.3172	0.0856	0.2496
Bull Mkt	5250	0.1374	0.0671	1.9636	8.3788	0.0651	0.0929	0.1191	0.1624	0.2738	0.0695	0.2087
Bear Mkt	1051	0.2505	0.1403	2.3754	9.0624	0.1274	0.1663	0.2127	0.2785	0.5873	0.1122	0.4599

TABLE 5
DIFFERENCES BETWEEN VIX & REALIZED VOLATILITY

Differences between VIX and Realized Volatility

Year	Mean	Std. Dev.	50% Norm. Range
1990	0.0768	0.0041	0.0183
1991	0.0473	0.0360	-0.0066
1992	0.0551	0.0004	0.0032
1993	0.0417	-0.0081	-0.0148
1994	0.0431	-0.0015	0.0009
1995	0.0482	-0.0080	-0.0151
1996	0.0523	-0.0092	-0.0191
1997	0.0514	-0.0183	0.0017
1998	0.0711	-0.0171	-0.0066
1999	0.0617	0.0003	0.0048
2000	0.0215	-0.0264	-0.0434
2001	0.0479	-0.0158	-0.0350
2002	0.0293	-0.0210	-0.0093
2003	0.0525	0.0034	-0.0149
2004	0.0445	0.0016	-0.0038
2005	0.0269	-0.0072	0.0001
2006	0.0304	-0.0042	-0.0111
2007	0.0281	-0.0059	-0.0196
2008	-0.0177	-0.0613	-0.1286
2009	0.0552	-0.0138	-0.0356
2010	0.0559	-0.0132	-0.0210
2011	0.0372	-0.0292	-0.0359
2012	0.0513	-0.0062	-0.0117
2013	0.0319	-0.0095	-0.0206
2014	0.0362	-0.0098	-0.0260
Overall	0.0432	-0.0138	0.0070
Bull Mkt	0.0461	-0.0043	0.0103
Bear Mkt	0.0285	-0.0348	-0.0301

TABLE 6
CORRELATION OF 21-DAY LAGGED VIX AND R_VOL (BEAR-MARKET PERIODS)

```
. correlate L.21(vix) r_vol, means
(obs=988)

    Variable |        Mean    Std. Dev.          Min          Max
-------------+----------------------------------------------------
     _delete |    .2724219     .1036373        .1563        .8086
       r_vol |    .2545801     .1425266        .0911        .8536

             |     L21.
             |      vix    r_vol
-------------+------------------
         vix |
        L21. |   1.0000
       r_vol |   0.7072   1.0000
```

Table 7 shows the summary statistics of the differences between the lagged VIX and realized volatility. Here, we observe a mean of 0.043 with a standard deviation of 0.06. Further, the distribution is negatively skewed with heavy tails as the kurtosis value of over 19 reveals. Visual observation of the histogram showing the differences between the lagged VIX and r_vol may be even more illuminating (*Figure* 5). The distribution is asymmetric with an obvious skew to the left. At the same time, we can see that the majority of the observations are well above zero confirming several earlier results that suggest that the VIX typically overstates realized volatility. We also notice that the right tails, albeit fatter than suggested by a normal distribution, are not nearly as frequent as those observed on the left tail. It is quite apparent that there are a lot more occurrences on the left tail compared to the right tail. Clearly, this distribution is far from being normal. We run the same analysis for bull-market periods. It reveals narrower data ranges, more symmetry and a slightly higher over-estimation of r_vol of about 485 basis points (*Table 8*).

TABLE 7
SUMMARY STATISTICS OF DIFFERENCES BETWEEN VIX AND R_VOL (1990-2014)

```
. sum d_21vix_rvol, detail
```

d_21vix_rvol

	Percentiles	Smallest		
1%	-.1678	-.4958		
5%	-.0375	-.4876		
10%	-.01075	-.4744	Obs	6280
25%	.0218	-.4657	Sum of Wgt.	6280
50%	.0473		Mean	.043292
		Largest	Std. Dev.	.0598602
75%	.07325	.2715		
90%	.0998	.2861	Variance	.0035832
95%	.11795	.2928	Skewness	-2.423001
99%	.1674	.3064	Kurtosis	19.25783

FIGURE 5
HISTOGRAM OF 21-DAY DIFFERENCES BETWEEN VIX AND R_VOL

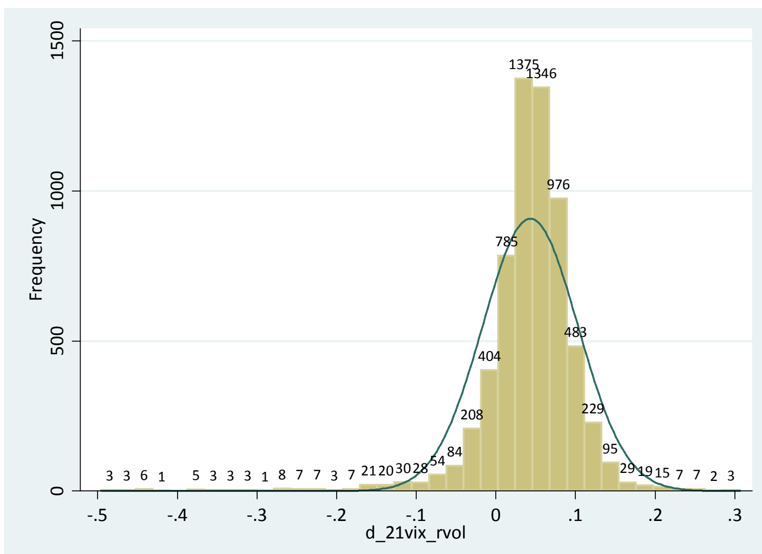

TABLE 8
SUMMARY STATISTICS: DIFFERENCES OF VIX AND R_VOL (BULL-MARKETS 1990 - 2014)

```
. sum d_21vix_rvol if bear==0, detail
```

d_21vix_rvol

	Percentiles	Smallest		
1%	-.1176	-.2786		
5%	-.0186	-.2747		
10%	.001	-.2736	Obs	5229
25%	.0266	-.2734	Sum of Wgt.	5229
50%	.0493		Mean	.0484649
		Largest	Std. Dev.	.0466806
75%	.0742	.2317		
90%	.0988	.2505	Variance	.0021791
95%	.1164	.2524	Skewness	-1.126385
99%	.1521	.2861	Kurtosis	10.22493

During bear-market periods, the mean is closer to zero (0.018) but the data indicate a much larger standard deviation of about 10% (*Table 9*). In other words, the estimation errors appear greater during bear-market periods. The distribution is still negatively skewed but there are essentially no outliers in the right tail. The majority of data points are narrowly dispersed around the mean. Once more, the left tail is longer and heavier than those expected from a normal distribution. The large amounts of absolute differences between VIX and r_vol. displayed in the left tail (*Figure 6*), however, are of much greater concern. Since volatility forecasting is a compulsory risk-management exercise for banks as part of the Basel Accords, (Granger and Poon 2003; BIS 2009), an accurate volatility forecast is critical for the survival of those institutions with large position risks. If the VIX is used as an input to estimate future market risk, this would again suggest a problem for many risk management applications. At the same time, we might ask if the VIX played a part, however small, in causing some of the fallout of the recent crisis.

TABLE 9
SUMMARY STATISTICS: DIFFERENCES OF VIX AND R_VOL (BEAR MARKETS 1990 - 2014)

```
. sum d_21vix_rvol if bear==1, detail

                        d_21vix_rvol
-------------------------------------------------------------
      Percentiles      Smallest
 1%        -.436         -.4958
 5%       -.1435         -.4876
10%       -.0657         -.4744       Obs                1051
25%       -.0152         -.4657       Sum of Wgt.        1051

50%        .0316                      Mean           .0175551
                        Largest       Std. Dev.       .098903
75%        .0677          .2657
90%        .1061          .2715       Variance       .0097818
95%        .1339          .2928       Skewness       -2.00001
99%        .2333          .3064       Kurtosis       10.54842
```

FIGURE 6
21-DAY DIFFERENCES: VIX AND R_VOL (BEAR MARKET PHASES 1990-2014)

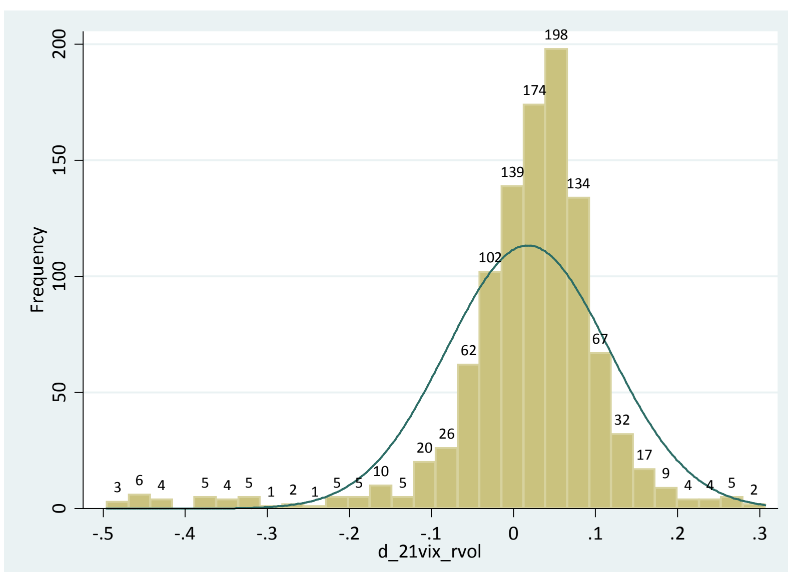

To investigate further, we propose a departure from traditional regressions and analyze range-bound percentage differences between implied and realized volatility instead. *Table 10* shows the percentage range of the differences of over- or under-prediction along with a frequency of how often the estimates occurred within that in that range. These range-bound percentage differences reveal that over 98% of the time, implied volatility over- or understates realized volatility by +/-1%. Over 86% of the time the VIX has an error margin of +/-10%.

TABLE 10
PERCENTAGE DIFFERENCES AND NUMBER OF OCCURRENCES OF VIX OVER/UNDERESTIMATING R_VOL

% Differences	Days inside Range	% Trading Days
+/- 0.10%	8	0.13%
+/- 0.25%	15	0.24%
+/- 0.50%	40	0.64%
+/- 0.75%	61	0.97%
+/- 1.00%	80	1.27%
+/- 2.50%	206	3.28%
+/- 5.00%	408	6.50%
+/- 10.00%	851	13.55%

We also wanted to find out if there were specific time periods when these error margins were greatest. To do so, we examine the top 40 and bottom 40 days of these instances. 23 of the top 40 days of over-estimating realized volatility (*Table 11*) occurred after the great recession. More remarkably, 14 of the top 40 days occurred during December 2010. When the Fed started its QE2 program in November 2010, options markets may have priced in concerns about a double-dip recession. While an additional risk-premium for uncertainty is understandable in the aftermath of the credit crisis, it remains puzzling to see the vast percentage differences between implied and actual volatility, again calling into question the effectiveness of the VIX as a forward looking indicator of risk.

Conversely, we investigate the bottom 40 observations of 21-day VIX differences with actual volatility *Table 12*. 26 of the bottom 40 days occur in the midst of the financial crisis in 2008 while 12 of these instances occurred between September 15th and October 1st, 2008. On September 15, 2008, Lehman Brothers filed for bankruptcy protection, which is considered to be the tipping point by many market participants. Despite this massive signal from the Lehman Brothers bankruptcy, the VIX was not able to appropriately capture market risk. Instead, it showed some of its worst periods of underestimating future risk during that same period. It is not just the amount of missing the target but the timing that makes forecast errors all the more devastating. While an overestimation of risk by the VIX may imply that investors are risk averse and would rather err on the side of caution, the level of risk aversion appears to be highest after the worst days of the credit crisis are over. A higher level of risk aversion may be costly but it may not nearly be as devastating as the level of under-estimation of market risk implied by the VIX during the recent crisis. *Figure 7* gives a graphical depiction of the difference between the VIX and realized volatility during the height of the financial crisis of 2008/09. Between October and December 2008, the VIX estimates show the largest discrepancies with realized volatility.

TABLE 11
TOP 40 - % DIFFERENCES: VIX & R_VOL

Date	vix	21-day % Δ	21-day Δ
12/1/2010	0.2136	261%	0.1544
12/2/2010	0.1939	246%	0.1379
3/16/2011	0.2940	222%	0.2028
3/17/2011	0.2637	221%	0.1817
12/6/2010	0.1802	216%	0.1232
12/3/2010	0.1801	216%	0.1231
1/3/2012	0.2297	216%	0.1570
11/22/1993	0.1590	215%	0.1085
12/8/1994	0.1815	210%	0.1230
12/7/2010	0.1799	207%	0.1213
12/10/2010	0.1761	203%	0.1179
10/16/1998	0.3482	199%	0.2317
12/8/2010	0.1774	198%	0.1178
11/4/2014	0.1489	194%	0.0982
11/5/2014	0.1417	194%	0.0935
11/23/1993	0.1427	190%	0.0936
12/9/2010	0.1725	188%	0.1126
10/15/1998	0.3334	187%	0.2172
12/15/2010	0.1794	185%	0.1165
10/20/1998	0.3311	185%	0.2150
10/19/1998	0.3313	185%	0.2151
10/21/1998	0.3321	185%	0.2155
12/16/2010	0.1739	184%	0.1127
8/2/1995	0.1366	184%	0.0884
11/3/2014	0.1473	183%	0.0953
8/1/1995	0.1356	181%	0.0874
8/3/1995	0.1378	179%	0.0884
11/16/1993	0.1511	177%	0.0966
10/29/2014	0.1515	176%	0.0965
12/13/2010	0.1755	175%	0.1118
8/7/2014	0.1666	175%	0.1060
8/14/1995	0.1324	174%	0.0841
12/9/1994	0.1604	174%	0.1019
2/23/2010	0.2137	174%	0.1357
3/1/2010	0.1926	173%	0.1221
10/31/2014	0.1403	173%	0.0888
12/12/1994	0.1546	172%	0.0978
11/29/1993	0.1412	172%	0.0893
12/14/2010	0.1761	172%	0.1113
8/8/2014	0.1577	171%	0.0995

TABLE 12
BOTTOM 40 - % DIFFERENCES: VIX & R_VOL

Date	vix	21-day % Δ	21-day Δ
9/12/2008	0.2566	-65.9%	-0.4958
8/29/2008	0.2065	-62.6%	-0.3451
9/3/2008	0.2143	-62.1%	-0.3515
8/28/2008	0.1943	-62.0%	-0.3169
9/16/2008	0.3030	-61.7%	-0.4876
9/8/2008	0.2264	-61.4%	-0.3609
7/22/2011	0.1752	-61.0%	-0.2736
7/21/2011	0.1756	-60.9%	-0.2734
9/10/2008	0.2452	-60.7%	-0.3790
9/11/2008	0.2439	-60.4%	-0.3720
9/2/2008	0.2199	-60.1%	-0.3317
9/5/2008	0.2306	-59.6%	-0.3397
9/19/2008	0.3207	-59.2%	-0.4657
9/25/2008	0.3282	-59.1%	-0.4744
7/25/2011	0.1935	-59.0%	-0.2786
7/20/2011	0.1909	-58.0%	-0.2636
7/26/2011	0.2023	-57.6%	-0.2747
9/18/2008	0.3310	-57.6%	-0.4491
9/15/2008	0.3170	-57.2%	-0.4228
9/4/2008	0.2403	-57.1%	-0.3195
9/22/2008	0.3385	-56.8%	-0.4450
9/26/2008	0.3474	-56.7%	-0.4543
9/24/2008	0.3519	-56.4%	-0.4560
9/9/2008	0.2547	-56.2%	-0.3264
7/19/2011	0.1921	-55.6%	-0.2407
9/23/2008	0.3572	-55.6%	-0.4465
7/15/2011	0.1953	-55.4%	-0.2431
9/17/2008	0.3622	-54.9%	-0.4403
8/22/2008	0.1881	-54.3%	-0.2235
7/13/2011	0.1991	-53.5%	-0.2293
10/1/2008	0.3981	-52.3%	-0.4360
9/30/2008	0.3939	-52.2%	-0.4308
7/18/2011	0.2095	-52.2%	-0.2291
8/27/2008	0.1976	-52.2%	-0.2158
8/1/2011	0.2366	-52.1%	-0.2575
8/3/2011	0.2338	-52.1%	-0.2544
8/21/2008	0.1982	-52.0%	-0.2143
7/27/2011	0.2298	-51.7%	-0.2456
7/14/2011	0.2080	-51.6%	-0.2219
7/28/2011	0.2374	-50.6%	-0.2433

FIGURE 7
LAGGED VIX AND REALIZED VOLATILITY DURING FINANCIAL CRISIS 2008/09

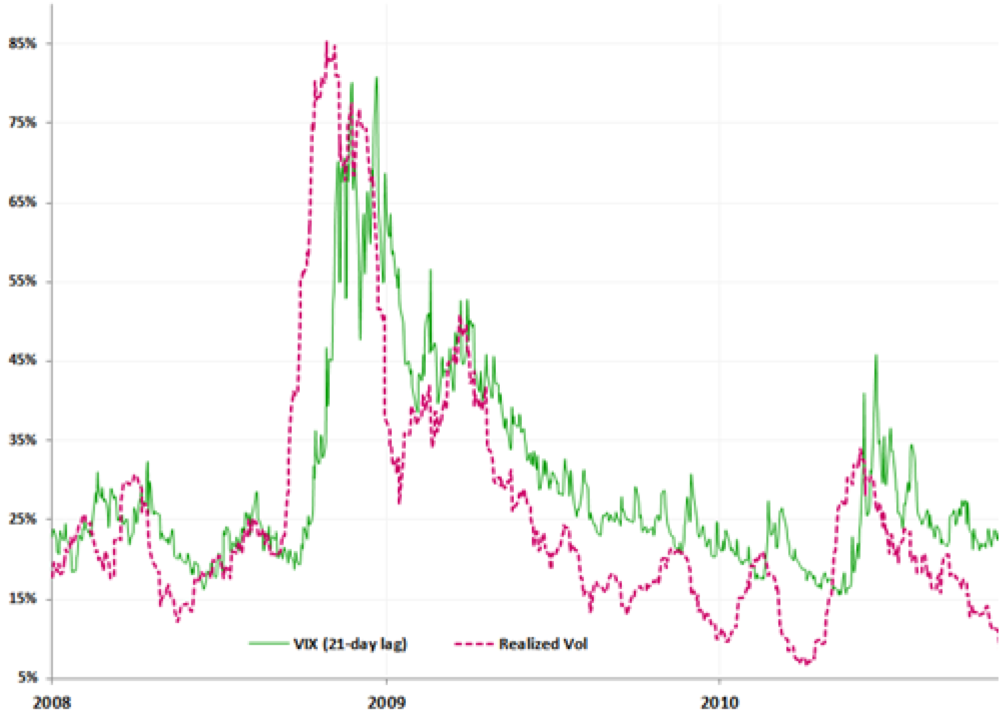

As a final study, we assess whether the VIX forecast of future risk might have better results at shorter time periods. We replicate histograms of the differences between the VIX and realized volatility and examine the 10-day and 1-day time periods to see if they would approximate a normal distribution *Figures 8&9*. Both histograms indicate heavy left-hand tails as well as a negative skew in the distribution. The majority of observations are positive suggesting a general over-estimation of actual volatility at both time periods. Here too, our concerns focus on the left-hand side where too many data points fall far outside the normal curve.

FIGURE 8
HISTOGRAM OF 10-DAY DIFFERENCES BETWEEN VIX AND R_VOL

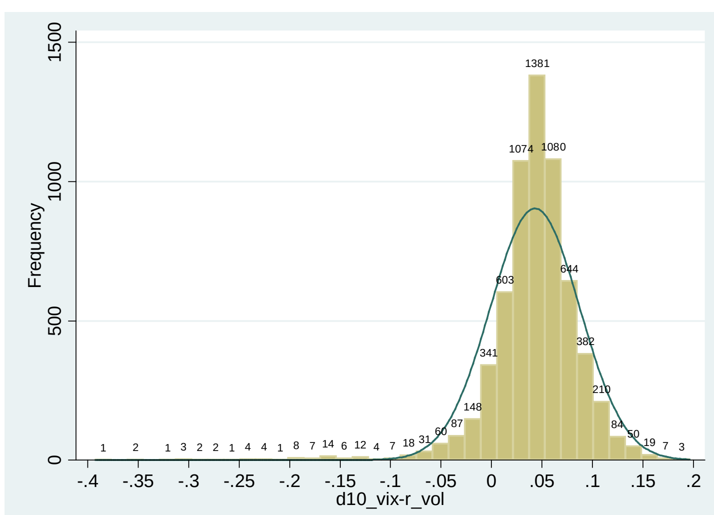

FIGURE 9
HISTOGRAM OF 1-DAY DIFFERENCES BETWEEN VIX AND R_VOL

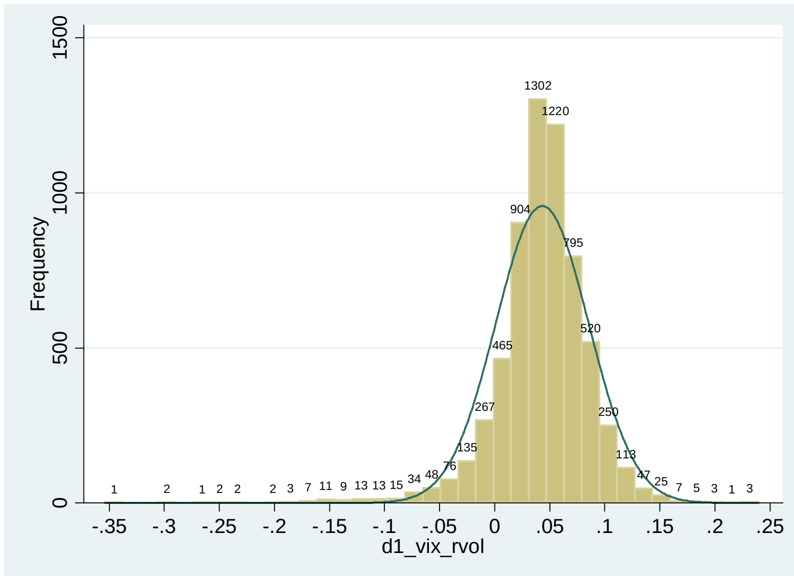

We also reference a more compelling visualization of the differences between VIX and r_vol. *Figure 10* shows the 21-day differences. The vast majority of observations are positive but there are instances, particularly during the crisis periods, when these differences are negative, the most crucial one of which falls into the midst of the financial crisis.

FIGURE 10
21-DAY DIFFERENCES BETWEEN VIX AND R_VOL FROM 1990-2014

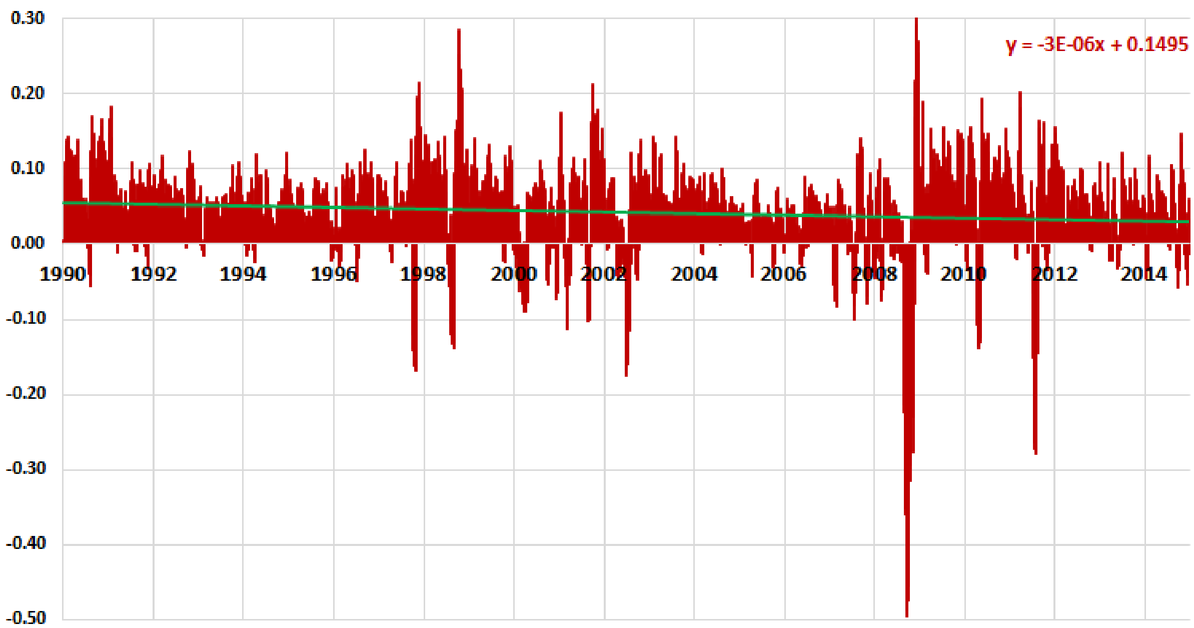

Lastly, we survey the data to see if there is an optimal time period when inter-temporal differences between VIX and r_vol are minimal. We conduct a study of 1-day up to 60-day differences and compute the basic summary statistics (*Table 13*).

TABLE 13
SUMMARY STATISTICS: 1-DAY TO 60-DAY DIFFERENCES BETWEEN VIX AND R_VOL

Days	Avg. Diff.	Std. Dev.	Skew	Kurtosis	Max	Min
0	0.04319	0.04293	-1.2889	7.4081	0.24398	-0.35783
1	0.04319	0.04209	-1.3250	7.6885	0.23960	-0.35372
2	0.04318	0.04153	-1.3727	8.0669	0.23573	-0.37084
3	0.04317	0.04120	-1.3990	7.9694	0.21194	-0.33169
4	0.04317	0.04118	-1.4038	7.7626	0.21058	-0.33188
5	0.04316	0.04121	-1.4726	7.9233	0.20883	-0.32246
6	0.04315	0.04141	-1.5866	8.8571	0.19484	-0.32386
7	0.04314	0.04188	-1.6780	9.5275	0.19355	-0.30297
8	0.04313	0.04247	-1.8109	10.6823	0.18150	-0.35426
9	0.04313	0.04331	-1.9444	11.7493	0.18202	-0.35846
10	0.04313	0.04430	-2.0591	12.6059	0.19647	-0.39247
11	0.04312	0.04551	-2.1984	13.9124	0.19637	-0.40496
12	0.04313	0.04668	-2.2999	14.7517	0.20797	-0.42416
13	0.04313	0.04798	-2.3771	15.4723	0.21327	-0.44317
14	0.04313	0.04927	-2.4617	16.2269	0.22401	-0.46237
15	0.04313	0.05054	-2.5061	16.6061	0.22274	-0.47425
16	0.04314	0.05186	-2.5051	16.6750	0.24838	-0.45194
17	0.04314	0.05331	-2.5224	16.8428	0.25004	-0.45820
18	0.04314	0.05480	-2.5097	16.7554	0.25772	-0.46987
19	0.04314	0.05640	-2.5150	16.8896	0.26598	-0.48175
20	0.04314	0.05802	-2.4750	16.5566	0.27291	-0.47972
21	0.04314	0.05984	-2.4136	16.2107	0.30642	-0.49576
25	0.04314	0.06543	-2.3980	16.5512	0.34024	-0.55855
30	0.04315	0.07074	-2.3591	16.9216	0.46501	-0.57627
40	0.04317	0.07844	-2.4487	16.9506	0.45146	-0.63366
50	0.04321	0.08437	-2.3823	15.8268	0.48138	-0.65054
60	0.04319	0.08841	-2.3043	14.4294	0.44454	-0.63393

Our findings are summarized in graphical format (*Figure 11*). Interestingly, the average differences across all time periods are incredibly steady at about 0.043. It clearly shows that the VIX consistently overstates actual volatility, no matter which forecasting period is chosen. It also reconfirms our intuition and gives a clear rationale to option traders who are, on balance, net-sellers of options thereby taking advantage of this overstatement of future risk. Similarly, the Skew is always negative. For periods of over 10 days, the skew remains consistently below -2 with only minor deviations between -2.3 and -2.5 throughout all periods greater than 13 days. The remaining moments are much less consistent. Standard deviation starts at just over 4% and declines slightly until 4-day periods. From 5-day differences onwards, we notice a gradual increase of standard deviation which finally peaks at 60-day differences. Kurtosis starts at just over 7 and stays between 7 and 8 until 5-day differences. From then on, it increases to around 17 and stays at very high levels up to 40-day differences. After that, it gradually decreases but remains above 14 for all periods up to 60-day differences.

FIGURE 11
MOMENTS OF 1-DAY to 60-DAY DIFFERENCES BETWEEN VIX AND R_VOL

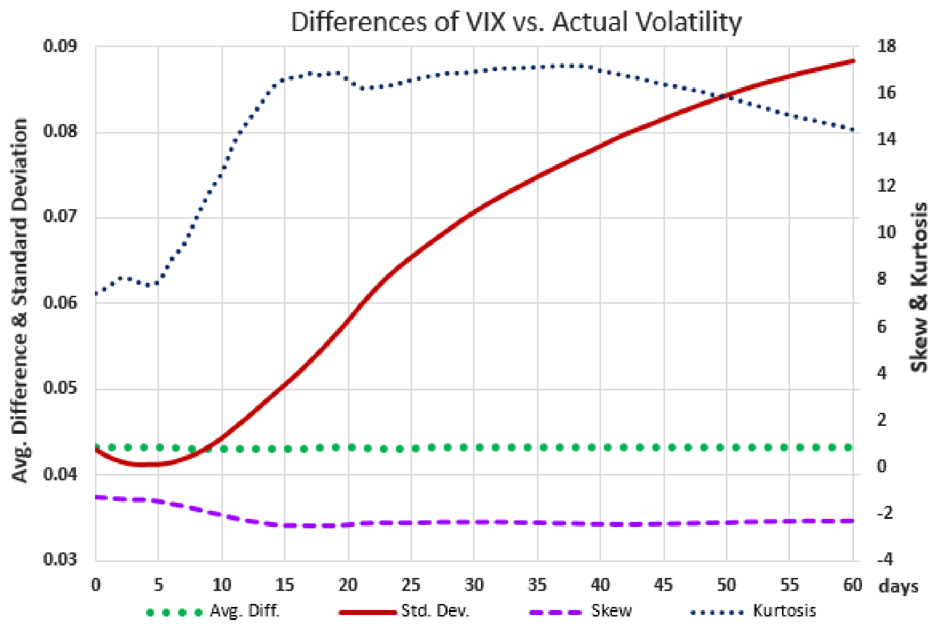

In our view, there are two major takeaways from this final study. The average over-estimation of actual volatility across all time periods is almost suspiciously consistent and we cannot find an optimal time period when the mean differences are substantially smaller than during other periods. Again, it provides an incentive for option traders to short implied volatility, which is typically their go-to strategy. However, the results also suggest that for periods of greater than 5-day differences, the potential for outliers increases rapidly, making it all-the-more difficult to manage risk when selling implied volatility – a big caveat for option traders. Standard deviations increase, suggesting larger forecast errors. Kurtosis doubles from 5-day to 15-day differences and the negative skew suggests that forecast errors are concentrated in the negative territory, in other words, an under-estimation of risk. Our results confirm previous studies by Martins and Zein (2002) and Christoffersen and Diebold (2000), who demonstrate that volatility forecasts have greater explanatory powers at shorter time horizons but may not be of much value for time horizons of more than ten days. To round up our analysis, we reference a 3-dimensional histogram that shows the distributions of inter-temporal differences between the VIX and actual volatility from 1-day to 30-day periods (*Figure 12*). This visualization confirms our previous studies that suggest, despite a relatively high correlation with actual volatility, the VIX may not be a useful predictor of risk.

FIGURE 12
3D HISTOGRAM OF 1-DAY - 30-DAY DIFFERENCES BETWEEN VIX AND R_VOL

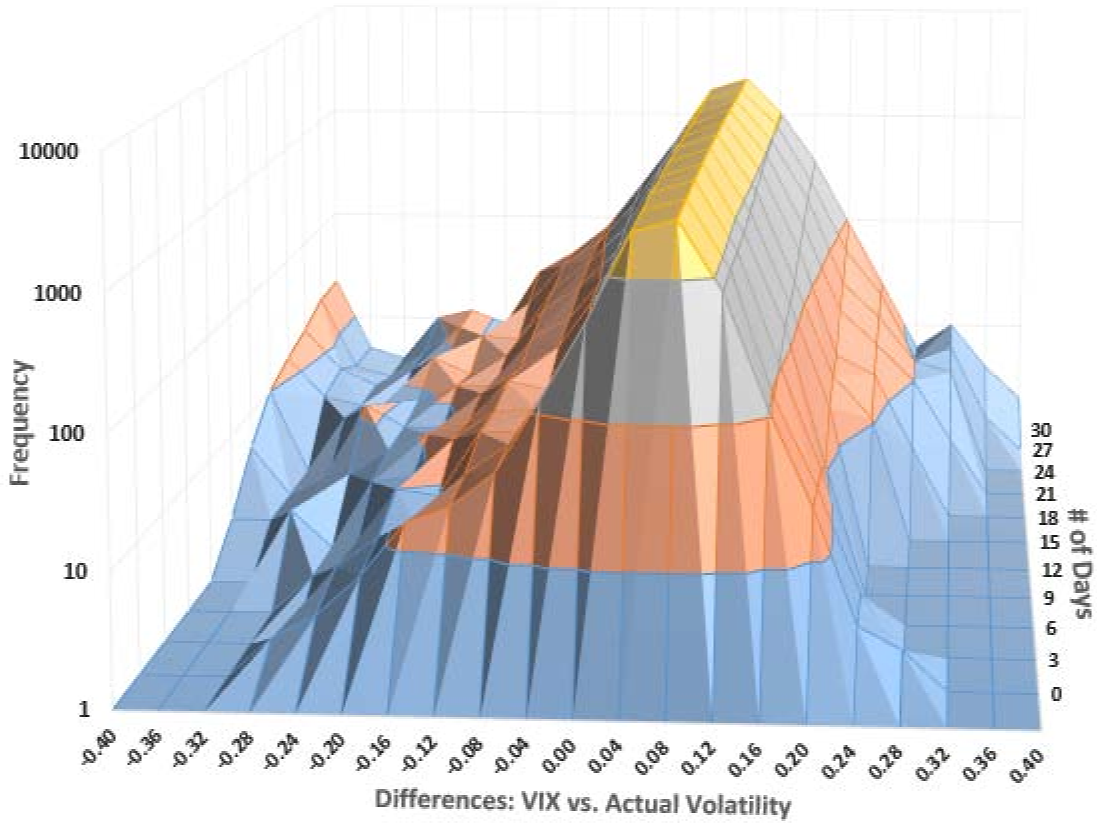

CONCLUSIONS

The forward-looking aspect of the VIX has been praised as a critical feature of investor expectations for a risk metric (Whaley 2009). Yet, our findings suggest that the VIX consistently overstates actual volatility by about 430 basis points across all time periods from 1-day to 60-day differences. Investor expectations for risk, as derived from options prices, may be inflated.

Interestingly, during the most critical time periods such as market crashes or crises, the VIX does not hold up to its promise. In the midst of the financial crisis of 2008, the VIX understates realized volatility by about 180 basis points on average. Timing plays an important role too. 26 of the 40 worst days of underestimating actual volatility occur between September and December 2008. Moreover, for the two-week period right after the Lehman Brothers bankruptcy, the VIX had some of its worst moments in terms of under-estimating actual market risk. Poor timing increases forecast errors.

Our studies suggest that the VIX is not effective as a forward-looking risk metric, particularly during periods when it matters most i.e. market crashes and crises. The '*investor fear gauge'* (Whaley 2000) may need some revisions. One might argue to what extent and frequency incorrect volatility estimates are tolerable for investors, but we must question the benefits of the VIX as an ideal risk management tool. Despite its negative correlation with the S&P 500, it fails to correctly anticipate market crashes and crises. We must also question why the VIX remains at elevated levels far too long after a market crash took place. In fairness, other studies suggest that implied volatility and the VIX are still better estimators of future volatility than more complex models (Granger and Poon, 2003 and Blair, Poon and Taylor, 2000) and perhaps we expect too much from a single variable in terms of quantifying the very complex nature of risk. However, a reliable risk metric for investors must meet a higher level of accuracy than currently provided by the VIX. At this moment, the VIX spreads fear too often and when fear is actually warranted, its warning signals are not loud and clear enough.

This paper left several potential follow up questions unanswered. The VIX showed an unusual consistency of over-estimating actual volatility by about 430 basis points across all time periods from 1-day to 60-day differences. We propose to investigate this unusual consistency in the hope of finding a better approach to model market risk. In addition, we put ourselves in the shoes of an investor who is typically concerned about down-side risk. But rather than forecasting volatility, a more appropriate question might be how implied volatility and the VIX affect stock market returns. More to the point, we would like to explore if there are specific VIX levels or volatility trends that can help us improve our risk-management or trading strategies. This will be examined in an extension to this paper.

ENDNOTES

1. The CBOE Volatility Index®, VIX®, was introduced in 1993 and is a registered trademark of the Chicago Board Options Exchange® (CBOE®)
2. For a comprehensive review of financial crises spanning eight centuries, see Reinhart and Rogoff (2009). The authors demonstrate that the recent U.S. subprime crisis is hardly unique.
3. See Sornette (2003) and Montier (2002) for discussions on outliers.
4. In the 1987 paper "How we came up with the option formula," Fischer Black describes how he developed a differential equation for valuing a warrant based on Treynor's work. His notes containing the differential equation were dated June 1969.
5. 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)
6. Andrew G. Haldane, The dog and the Frisbee, Speech given at Fed Kansas City, 31 August 2012
7. Hull explains that a VaR calculation is aimed at making a statement in the following form: *"We are X percent certain that we will not lose more than V dollars in the next n days"* whereby *V* is the VaR, *X%* is the confidence level, and *n* days is the time horizon (Hull 2012, p. 517)
8. (cited by 16,470 as of 9/14) http://community.stern.nyu.edu/rengle/research/
9. ARCH specified conditional variance as a linear function of past sample variances only, but 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)
10. The objective of this chapter was to give a coherent overview of the major lines of work in this field. The literature on volatility, risk, and return in different financial markets is so vast, however, that we have unavoidably omitted many important studies.
11. VIX values are usually stated as☐ x 100. However, for our studie s, we converted official VIX and VXO data were to decimals to better facilitate statistical analysis. For instance, a VIX level of 20 is converted to 0.2
12. The *interquartile range* is defined as the difference between the upper quartile and the lower quartile. It is often used as a measure of dispersion for skewed distributions since it is insensitive to outliers.
13. Bear market dates are identified as: 7/16/90 – 10/11/90; 4/3/00 – 10/9/02; 10/9/07 – 3/9/09
    Bull market dates are identified as: 1/1/90 - 7/15/90; 10/12/90 - 4/2/00; 10/10/02 - 10/08/07; 03/10/09 - 12/31/14

REFERENCES

Aboura, Sofiane and Wagner, Niklas (2015). Extreme Asymmetric Volatility: Stress and Aggregate Asset Prices. *Available at SSRN: http://ssrn.com/abstract=1348563*

Andersen, Torben G., Bollerslev Tim, Diebold Francis X and Ebens Heiko (2001). The Distribution Of Realized Stock Return Volatility. *Journal of Financial Economics, 2001, v61 (1,Jul), 43-76.*

Bali, Turan G. and Demirtas, K. Ozgur (2006). Testing Mean Reversion in Stock Market Volatility. Journal of Futures Markets, Vol. 28, No. 1, pp. 1-33, 2008.

Bank for International Settlements (2009). Revisions to the Basel II market risk framework, *BIS (2009) ISBN print: 92-9131-774-8.*

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 (2001). Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns. *Journal of Econometrics 105.*

Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics, 31, 307-327.*

Bollerslev, T. (2007). Glossary to ARCH (GARCH). *Duke University and NBER, October 25, 2007.*

Brenner, Menachem, F and Galai, Dan (1989). New Financial Instruments for Hedging Changes in Volatility. *Financial Analysts Journal, July/August 1989.*

CBOE (2009). THE CBOE VOLATILITY INDEX® - VIX®. *Chicago Board Options Exchange.*

Chen, Kan and Jayaprakash, Ciriyam and Yuan, Baosheng (2005). Conditional Probability as a Measure of Volatility Clustering in Financial Time Series. *SSRN: http://ssrn.com/abstract=688741.*

Christoffersen, Peter F. and Diebold Francis X. (2000). How Relevant is Volatility Forecasting for Financial Risk Management? *The Review of Economics and Statistics, Vol. 82, No. 1, pp. 12-22*

Cont, Rama (2005) Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models. *Springer, Long memory in economics.*

Engle, Robert (2004). Econometric Models and Financial Practice. *The American Economic Review, Vol. 94, No. 3 (Jun., 2004), pp. 405-420.*

Engle, Robert (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica, Vol. 50, No. 4 (July 1982): 987-1007.*

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*

Fieldhouse, Stuart (2012). Trading Volatility as an Asset Class, *The Hedge Fund Journal.*

Fouque Jean-Pierre, Papanicolaou George, Sircar K. Ronnie (2008). Mean-Reverting Stochastic Volatility. *Working Paper, 2008.*

Gabaix, Xavier (2009). Power Laws in Economics and Finance. *Annual Review of Economics, Annual Reviews, vol. 1(1), pages 255-294, 05.*

Granger, Clive W. J. and Poon, Ser-Huang (2003). Forecasting Volatility in Financial Markets: A Review. *Journal of Economic Literature Vol. XLI, pp. 478–539.*

Grant, Maria and Gregory, Krag and Lui, Jason (2007). Volatility as an asset. *Goldman Sachs Options Research, November 15, 2007.*

Hafner Christian (2003). Simple approximations for option pricing under mean reversion and stochastic volatility. *Econometric Institute Research Papers, No EI 2003-20.*

Haug, Espen Gaarder and Taleb, Nassim Nicholas (2009). Option Traders Use (very) Sophisticated Heuristics, Never the Black–Scholes–Merton Formula. *Journal of Economic Behavior and Organization, Vol. 77, No. 2, 2011.*

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*

Lintner, John (1965a). The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets. *Review of Economics and Statistics. 47:1, 13-37.*

Lintner, John. (1965b). Securities Prices, Risk and Maximal Gains from Diversification. *The Journal of Finance 20(4), 587-615.*

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

Markowitz, Harry (1952). Portfolio Selection. *Journal of Finance. 7:1, pp. 77-99.*

Markowitz, Harry (1959). Portfolio Selection: Efficient Diversification of Investments. *Cowles Foundation Monograph No. 16. New York: John Wiley & Sons, Inc.*

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

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

Mixon, Scott (2008). Option Markets and Implied Volatility: Past Versus Present. *Journal of Financial Economics (JFE). Available at SSRN: http://ssrn.com/abstract=889543*

Montier, James (2002), Behavioral Finance, *John Wiley Publishing, Hoboken, New Jersey.*

Mossin, J. (1966). Equilibrium in a Capital Asset Market. *Econometrica 34 (4), 768-783.*

Nobelprize.org (1997). Advanced Information - Additional background material on the Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel 1997. *Nobel Media AB 2014. Web. http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/1997/advanced.html*

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, 2005, pp 269-281.*

Reinhart, C. M. and Rogoff, K.S. (2009). This Time It's Different: Eight Centuries of Financial Folly, *Princeton University Press, Princeton, New Jersey.*

Sharpe, William F. (1964). Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *Journal of Finance. 19:3, pp. 425-442.*

Sornette, Didier (2003). Why Stock Markets Crash: Critical Events in Complex Financial Systems, *Princeton University Press, Princeton, New Jersey.*

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

Taleb, Nassim Nicholas (2007). The Black Swan: The Impact of the Highly Improbable, *Random House*

Tobin, James (1958). Liquidity Preference as Behavior Towards Risk. *Review of Economic Studies, February 1958, 25(2), pp. 65-86.*

Treynor, Jack L. (1962). Toward a Theory of Market Value of Risky Assets, *Unpublished manuscript. A final version was published in 1999, in Asset Pricing and Portfolio Performance: Models, Strategy and Performance Metrics. Robert A. Korajczyk (editor) London: Risk Books, pp. 15-22.*

Treynor, Jack.L. (1961). Market Value, Time, and Risk. *Unpublished manuscript. "Rough Draft"*

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

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

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

Wong, Hoi Ying and Lo, Yu Wai (2009). Option Pricing with Mean Reversion and Stochastic Volatility. *European Journal of Operational Research 197, 179-187, 2009.*