The Implied Convexity of VIX Futures

Robert T. Daigler, Brice Dupoyet, and Fernando M. Patterson

Robert T. Daigler is Knight Ridder Research Professor of Finance at Florida International University in Miami, FL.
daiglerr@fiu.edu

Brice Dupoyet is an associate professor of finance at Florida International University in Miami, FL.
dupoyetb@fiu.edu

Fernando M. Patterson is an assistant professor of finance at The University of Texas Rio Grande Valley in Edinburg, Texas.
fernando.patterson@utrgv.edu

An important component of theoretical CBOE Volatility Index (VIX) futures prices is a term correcting for the negative convexity of the square root function by subtracting from the forward-starting variance swap rate an estimate of the future volatility of VIX futures prices. In the same fashion that an index option's traditional implied volatility can be viewed as an aggregate market consensus of future realized volatility, this convexity value can be viewed as an aggregate market consensus of future volatility of volatility. This article examines the predictive properties and features of this convexity adjustment needed to value VIX futures prices by extracting it from the relationship between observed VIX futures prices and the corresponding spot option market prices used to compute the forward-starting variance swap rate. The authors find that implied convexity levels can indeed be used to forecast the future volatility of VIX futures prices, even though implied convexity consistently underestimates future realized VIX futures variance. They also show that implied convexity can at times violate strict theoretical conditions by being negative, although we are able to rule out arbitrage opportunities. Finally, they examine the properties of this implied convexity adjustment, both as a time series and with respect to various market volatility factors with which they find positive and statistically significant relations.

CBOE Volatility Index (VIX) futures and options pricing are active areas of research, both because volatility products are considered a new asset class and because the financial community has yet to accept a VIX futures or VIX options pricing model as the definitive one. The difficulties in pricing VIX products arise from the need to accurately estimate a number of complex factors under different market scenarios, the importance of the stability of these factors, and the size of the mispricing of current existing models. Furthermore, for VIX options, determining the appropriate value of the volatility of volatility (the implied volatility of VIX options) is challenging. Consequently, the limited market-based empirical evidence does not support the long-term pricing stability of any of the existing models. As Poon and Granger [2003] pointed out, "much remains to be done to understand the process and characteristics of volatility."

In the pricing of VIX futures, the issue of (negative) convexity arises through Jensen's inequality and the presence of a square root linking the VIX futures price to the forward-starting variance swap rate. The expectation of a square root being less than the square root of the expectation, an adjustment must be made downward from the latter in order to match the former. More formally, at a given time t, for any random variable x_T whose value is revealed at time $T > t$, we know that

$$Var_t(x_T) = E_t(x_T^2) - \left[E_t(x_T)\right]^2 \qquad (1)$$

where $Var_t(.)$ represents the variance of a random variable and $E_t(.)$ denotes its expectation.

Substituting the value of the VIX futures price at expiration, F_T, for x_T into (1) and denoting the risk-neutral probability measure by Q, we obtain

$$Var_t^Q(F_T) = E_t^Q(F_T^2) - \left[E_t^Q(F_T)\right]^2 \quad (2)$$

Furthermore, because $E_t^Q(F_T) = E_t^Q(VIX_T) = F_t =$ the fair value of the VIX futures price today, and because $E_t^Q(F_T^2) = E_t^Q(VIX_T^2) =$ the forward-starting variance swap rate, we can express the fair value of the VIX futures price at time t today, expiring at time T as

$$F_t = \sqrt{E_t^Q(VIX_T^2) - Var_t^Q(F_T)} \quad (3)$$

As noted in the CBOE's VIX White Paper [2003], the expression $Var_t^Q(F_T)$ is the variance of the VIX futures prices from the present to the futures contract's expiration, which can be viewed as the daily cumulative variance of VIX futures prices from the current time t to expiration time T. The adjustment for convexity therefore turns out representing the future volatility of volatility. Because this variance is *a priori* unknown, market participants need to estimate it. Its value essentially provides the "missing piece" needed to determine the fair price of a VIX futures contract. Exhibit 1 illustrates graphically how: the larger the VIX_T level, the larger the required convexity adjustment. Because the convexity adjustment consistent with market-observed VIX futures prices is by definition the market's aggregate estimate of the future variance of these VIX futures prices, also referred to as "the volatility of volatility," we develop a methodology to extract it from the data and introduce the concept of *implied* convexity of VIX futures prices and examine its forecasting and statistical properties.[1] Conceptually, this is analogous to extracting an implied volatility from option prices, because both implied volatility and implied convexity levels become the inputs needed for a formula to match market-observed option prices and VIX futures prices, respectively.

We do find that the empirically estimated implied convexity can indeed fairly accurately predict the future realized variance of VIX futures prices, making the implied convexity concept a forward-looking measure of VIX futures future variance. Implied convexity also consistently underestimates the future variance of VIX futures, in a way that resembles—albeit in reverse—the well-known pattern whereby the implied volatility of

Exhibit 1
Concavity/Convexity Adjustment for Small and Large Variances

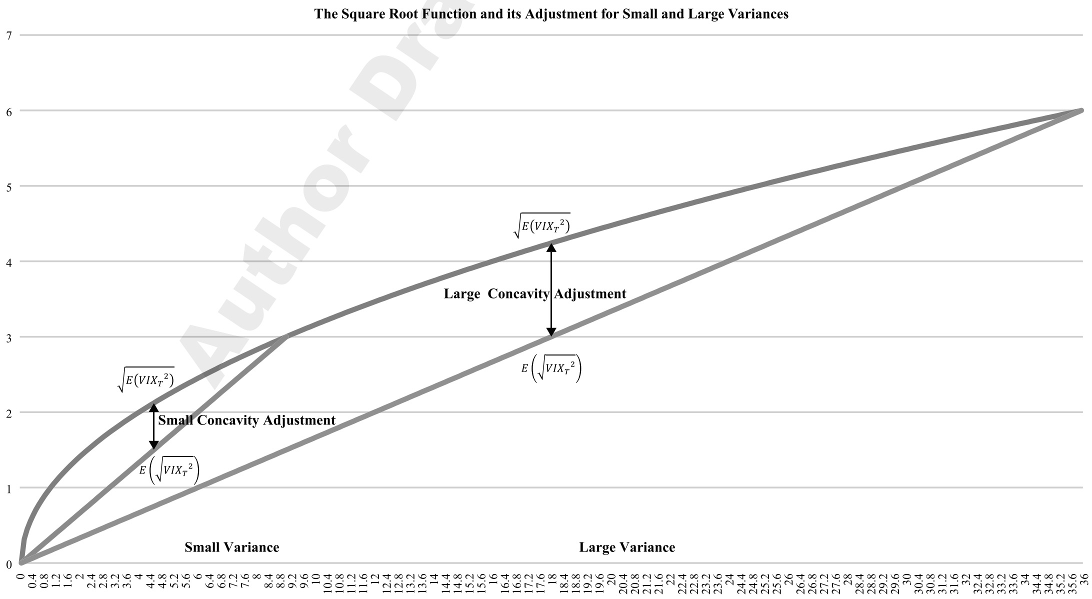

Note: This exhibit illustrates the adjustment required to reduce the square root of the forward-starting variance swap rate to the VIX futures price (by definition equal to the risk-neutral expectation of the VIX level at maturity, VIX_T).

an option consistently overestimates the future realized volatility of the underlying asset. Being a forecast of future volatility, implied convexity also naturally decreases as the VIX futures contract gets closer to expiration, because future potential volatility decreases with the remaining time to expiration—a phenomenon similar to what corporate bonds experience when getting closer to maturity. These forecasted volatility levels can be used to better price derivative products whose values depend on future levels of volatility. We also find that implied convexity levels sometimes violate the theoretical condition related to the forward-starting variance swap rate. However, we show that such results cannot be turned into arbitrage-free profits for a variety of reasons that we will discuss in more detail. Finally, we regress implied convexity levels on a variety of market volatility factors as well as control variables and obtain statistically significant relations in most cases, validating intuitive predictions regarding implied convexity dynamics.

A BRIEF LOOK AT CURRENT MODELS OF VIX FUTURES PRICING

Because the underlying asset in VIX futures contracts is an index of S&P 500 Index options and therefore not directly traded, the behavior of the VIX futures price series differs from that of a typical cost-of-carry futures contract.[2] Moreover, the volatility stochastic process is not completely known, which leads to difficulties in designing an effective model for VIX futures pricing. Finally, such instruments are affected by mean-reversion, jumps, the dynamic nature of the VIX and VIX futures volatility over time, and the shifting term structure of volatility.

Some of the more complete endeavors to price VIX futures include research by Lin [2007], Zhang and Huang [2010], Zhu and Lian [2012], and Mencia and Sentana [2013]. Some papers incorporate the convexity adjustment into their models using an approximation of the Taylor series and a selected volatility process. For example, Lin [2007] employed a second-order correction approximation; Lu and Zhu [2010] used a Kalman filter and a maximum likelihood approach to estimate their model; and Zhang, Shu, and Brenner [2010] used a third-order approximation based on a Heston stochastic volatility framework. Zhu and Lian [2012] compared their "exact" formula to other approximation methods, finding that these approximations can create large percentage errors and, through the use of simulations, show that the errors increase as volatility increases. However, these complex models contain a large number of parameters whose estimation can be both challenging and unreliable. Additionally, because the parameters have the potential to change over time, the consistency of these models is questionable. For example, Zhu and Lian [2012] stated that "Just as an issue raised in Zhang and Huang [2010], we also believe that searching for a reliable estimation method to determine the model parameters from market data remains a challenge. While our work presented in this article has demonstrated another alternative, complicated stochastic models probably will not gain popularity among market practitioners, until a convincing approach can be accepted and agreed upon by a majority of researchers."

As such, a handful of studies such as Huskaj and Nossman [2013] estimated the parameters exogenously with some degree of success in capturing the term structure of VIX futures compared with previous studies. Other studies, for example, Grunbichler and Longstaff [1996], Zhang and Zhu [2006], and Dupoyet, Daigler, and Chen [2011], provided simpler, more tractable models involving some combination of the constant elasticity of variance (CEV) and Poisson jumps (generally based on the Cox–Ingersoll–Ross [1985] approach). Still, the stability of the results for these models over time needs more empirical support. In fact, empirical studies of such models typically show large percentage errors.[3] Similarly, studying the volatility of volatility and the implied volatility skew (most important for VIX options), as in Wang and Daigler [2013], only seems to scratch the surface of this complex issue.

Rather than proposing yet another stochastic process for the VIX and a convexity term approximation, we back out implied convexity values from observed VIX futures prices and the strip of underlying S&P 500 options that determine the spot VIX (adjusted for time to expiration) and test their predictive and statistical properties for several periods of varying market conditions, including the Fall 2008 period. Similar in concept to the implied volatility of an index option, implied convexity provides a form of aggregate market forecast of the future VIX futures volatility. An obvious promising use for practitioners as well as academics is the ability of implied convexity to provide volatility forecasts. Moreover, implied convexity can be compared with the implied volatility of VIX options as an alternative method to examine the volatility of volatility (see Wang and Daigler [2013] for research on this area). Of course, as with implied volatility, changes in the market will affect future VIX futures volatility. That being said, a model that is a reasonable forecast of future VIX futures volatility will nevertheless be a useful benchmark. We initiate this investigation in this article.

METHODOLOGY

In this section, we detail the procedure used to extract the implied convexity values. First, on the day the VIX futures contract expires, the 30-day implied volatility forecast by the VIX futures is the implied volatility determined from the S&P 500 options series that expires 30 days after the VIX futures contract's expiration. Before the expiration of the VIX futures contract, however, the 30-day implied volatility underlying the futures contract is a "forward-starting implied volatility"; that is, a 30-day volatility starting at expiration of the futures contract. Consequently, one must calculate the forward implied volatility (variance) from any specific day before the contract's expiration.

In order to introduce the concept of implied convexity, we first examine the calculation of variances in general. Over a period of N days, treating the sample mean as equal to zero, the return variance is calculated as

$$\sigma^2 = \frac{1}{N}\sum_{i=1}^{N} r_i^2 \quad \text{where } r_i \text{ is the return on day } i \tag{4}$$

Splitting the time series into two periods of lengths n and p, where $n + p = N$, we have

$$\sigma^2 = \frac{1}{N}\left[\sum_{i=1}^{n} r_i^2 + \sum_{i=n+1}^{N} r_i^2\right] = \frac{1}{N}\sum_{i=1}^{n} r_i^2 + \frac{1}{N}\sum_{i=n+1}^{N} r_i^2$$

$$\sigma^2 = \frac{n}{N}\left(\frac{1}{n}\sum_{i=1}^{n} r_i^2\right) + \frac{p}{N}\left(\frac{1}{p}\sum_{i=n+1}^{N} r_i^2\right) \tag{5}$$

Therefore, the return variance for the full period can be expressed as a weighted average of the variances of the two subperiods, with the weights being the proportion of time occupied by each variance regime:

$$\sigma^2 = \frac{n}{N}\sigma_n^2 + \frac{p}{N}\sigma_p^2 \tag{6}$$

where σ_n^2 and σ_p^2 are the return variances during the periods of length n and p, respectively.

It is important to note that Equation (6) is related to, yet different from, the usual statement that the variance of a return over a horizon or frequency of length n will be n/N of the variance of a return over a horizon or frequency of length N. This last statement requires both independence of the returns (so that variances are additive) and a constant variance over time. Alternatively, Equation (6) does not require such an assumption, and in fact, it is specifically designed to deal with possible changes in variances. Therefore, in the first period or regime of length n, the variance σ_n^2 could be high, and in the subsequent period or regime of length p, the variance σ_p^2 could be low, perhaps as a result of mean reversion, for instance. Consequently, Equation (6) simply provides a tool to compute the overall variance over the full period, given that two different volatility regimes are experienced. However, in the special case where the variance is constant across subperiods, the formula collapses back to the usual statement that the variance of returns over a horizon or frequency of length n will be n/N of the variance of returns over a horizon or frequency of length N. Additionally, one could split the full sample into as many subperiods as needed, and Equation (6) would simply expand to more terms. Finally, Equation (6) also can be expressed as

$$N\sigma^2 = n\sigma_n^2 + p\sigma_p^2 \tag{7}$$

This means that more generally, for three ordered time points or markers identified as t_1, t_2, and t_3, such that $t_1<t_2<t_3$, the relationship among the variances associated with the various corresponding periods needed to determine a total variance over time period 1 to 3 is as follows:

$$(t_3 - t_1)\sigma^2_{t_1\to t_3} = (t_2 - t_1)\sigma^2_{t_1\to t_2} + (t_3 - t_2)\sigma^2_{t_2\to t_3} \tag{8}$$

where $\sigma^2_{t_i\to t_j}$ represents the variance for the period of time between markers i and j.[4] However, in the context of a VIX futures contract priced at time t, expiring at time T, and "delivering" the *future* implied market volatility (VIX) based on the S&P 500 options expiring in 30 days from VIX futures expiration—and using calendar days in this description—we develop Equation (9) providing the total variance between the present time t and 30 days after the VIX futures contract expires:

$$(T + 30 - t)\sigma^2_{t\to T+30} = (T - t)\sigma^2_{t\to T} + 30\sigma^2_{T\to T+30} \tag{9}$$

where $\sigma^2_{t \to T+30}$ represents the variance between the present and 30 days after expiration of the VIX futures contract, $\sigma^2_{t \to T}$ is the variance between the present time and the VIX futures expiration, and $\sigma^2_{T \to T+30}$ is the variance between the VIX futures expiration and 30 days after expiration. The latter is also the forward-starting variance swap rate representing the implied variance embedded in VIX futures prices, and we can therefore solve for it as follows:

$$\sigma^2_{T \to T+30} = \left(\frac{T+30-t}{30}\right)\sigma^2_{t \to T+30} - \left(\frac{T-t}{30}\right)\sigma^2_{t \to T} \quad (10)$$

The CBOE also tracks two indexes relating to the spot VIX that similarly employ a portfolio of S&P 500 options and are calculated using the methodology developed for the VIX: the VIN (the implied volatility for the near-term S&P 500 options expiration) and the VIF (the implied volatility for the far-term S&P 500 options expiration). These series are the individual components that determine the spot VIX. Instead of representing a constant *n*-day-ahead volatility as the VIX does, however, the VIN and the VIF determine the volatility between the present and the first-month option expiration (VIN) and the volatility between the present and the next-month option expiration (VIF). Therefore, the VIX can be computed as a weighted average between the VIN and the VIF.

We can employ the VIN and VIF directly to compute the forward rate implied variance, if the maturity of the VIN corresponds exactly to that of the expiration of the VIX futures contract and the maturity of the VIF corresponds exactly to 30 days after the VIX expiration. In that case, $\sigma^2_{t \to T+30}$ represents the square of the far-term VIX (VIF) and $\sigma^2_{t \to T}$ represents the square of the near-term VIX (VIN) in Equations (9) and (10). Consequently, the forward-starting variance swap rate in that situation is expressed as

$$E_t^Q(VIX_T^2) = \sigma^2_{T \to T+30} = \left(\frac{T+30-t}{30}\right)VIF^2 - \left(\frac{T-t}{30}\right)VIN^2 \quad (11)$$

In reality, the horizons of the VIN and the VIF do not perfectly match the expiration dates of the VIX futures contract for the VIN and the corresponding 30 days later for the VIF; in fact, the difference in the number of days between the VIN and the VIF is either 28 or 35 days, not 30 days.[5] For Equation (11) to be applicable, we need an adjusted horizon for the VIN to match that of the expiration of the VIX futures contract, and an adjusted horizon for the VIF to match 30 days after the expiration of the VIX futures contract, which would then provide the 30-day forward-starting variance for the VIX futures. This can be accomplished by judiciously weighting the VIN and the VIF indexes. This procedure is analogous to a linear interpolation of yield curve rates (term structure), albeit applied to the volatility term structure in this case. It is precisely because the volatility is not the same for different horizons (as a result of expectations of mean reversion, for instance) that volatility levels for various terms or horizons will not be equal. The weighted average performed is simply a linear interpolation of two volatility levels on the volatility term structure curve.

If the horizon of the VIN is n_1 calendar days after the expiration of the VIX futures contract and the maturity of the VIF is n_2 days after $T+30$, we can create the needed adjusted values for both indexes by computing

$$Adjusted\_VIN^2 = \left(\frac{30+n_2}{30+n_2-n_1}\right)VIN^2 + \left(\frac{-n_1}{30+n_2-n_1}\right)VIF^2 \quad (12)$$

$$Adjusted\_VIF^2 = \left(\frac{n_2}{30+n_2-n_1}\right)VIN^2 + \left(\frac{30-n_1}{30+n_2-n_1}\right)VIF^2 \quad (13)$$

We can then apply Equation (11) to the new horizon-adjusted volatility indexes, yielding an equation for the forward-starting variance swap rate when time adjustments are needed to match the VIX futures expiration:

$$E_t^Q(VIX_T^2) = \sigma^2_{T \to T+30} = \left(\frac{T+30-t}{30}\right)Adjusted\_VIF^2 - \left(\frac{T-t}{30}\right)Adjusted\_VIN^2 \quad (14)$$

Finally, using the obtained forward-starting variance swap rate along with the current (time-*t*) observed VIX futures price, the market-implied convexity can be derived as

$$IC_t = Var_t^Q(F_T) = E_t^Q(VIX_T^2) - F_t^2 \quad (15)$$

where IC_t is the implied convexity, $Var_t^Q(F_T)$ is the expected variance of future VIX futures prices between time t and time T, $E_t^Q(VIX_T^2)$ is the forward-starting variance swap rate computed in (14), and F_t^2 is the square of the current price of the VIX futures contract expiring at time T.

DATA

VIX futures contracts started trading in March of 2004. Therefore, VIX futures prices and volatility are observable since that date. The forward-starting variance swap rates needed for the spot VIX comparison are not traded on exchanges, but they can be determined from the value of a portfolio of S&P 500 options. Therefore, we take advantage of publicly reported volatility indexes that reflect the prices of S&P 500 option portfolios to generate forward-starting variance swap rates according to the methodology described earlier.

We obtain daily data for the VIN and VIF series defined in the previous section from TradeStation from August 25, 2008, to July 29, 2015.[6] When necessary, the series are adjusted according to Equations (12) and (13), respectively. The VIN and VIF roll over exactly one week prior to the third Friday of each month, with the third Friday being the expiration of the S&P 500 options. Therefore, any unusual behavior in the options in the last week (as part of the VIN) does not influence the indexes. The VIN horizon (n_1) is calculated as the number of minutes remaining between time t and one week prior to the third Friday of each month. The VIF horizon (n_2) is calculated as the number of minutes remaining between time t and one week prior to the third Friday of the following month. The adjusted VIN and VIF are used to estimate the forward-starting variance swap rate as presented in Equation (14).

VIX futures data are obtained from the CQG Data Factory for all contracts expiring between August 2008 and July 2015. The VIX futures series employed here is composed of only the nearby contracts. We roll over to the next nearby contract on the Monday before the third Friday of every month in order to avoid any pricing issues related to the settlement of the futures contract. The last bid and ask of the VIX futures prices comprise the daily closing values for the analysis. Finally, the daily closing level of implied convexity is determined from the daily closing VIX futures mid-price and the contemporaneous forward-starting variance swap rate using Equation (15).

IMPLIED CONVEXITY PROPERTIES AND RESULTS

Given the lack of prior studies on the behavior of implied convexity, a reasonable starting point is to examine the time series history of the measure. Exhibit 2 displays VIX index levels (Panel A) and implied convexity levels (Panel B) from August 2008 through July 2015. This exhibit shows that the highest levels of implied convexity occur during periods when the VIX also peaks. Visually, the implied convexity exhibits large positive spikes; moreover, the speed of mean reversals in the implied convexity series is higher than that of the VIX. Furthermore, the implied convexity is fairly stable as the VIX reverses to lower levels.

The important characteristics observed in Exhibit 2 are logical for two reasons. First, implied convexity is the market's aggregate forecast of the future realized variance of VIX futures prices between now and the futures contract's expiration. Thus, implied convexity represents the variance of futures prices, not returns. This is important because for a given level of *return* variance, a higher price will display a higher level of variance, simply due to the scaling effect. Additionally, because a high spot VIX level produces high VIX futures prices, this will generate high implied convexity values, because the implied convexity is an estimate of the future realized variance of VIX futures prices. Conversely, lower VIX levels will generate lower VIX futures prices and commensurably lower implied convexity values. In conclusion, the size of the convexity adjustment is strongly positively correlated to the level of the VIX, as shown in Exhibit 2.

The second reason for which the results in Exhibit 2 are sensible is the fact that mean-reversion occurs faster for the convexity adjustment than it does for the VIX itself. If the market anticipates the VIX to gradually continue on a downward trend after an initial decline, then the convexity adjustment should be affected in two ways: First, the VIX will trend toward a lower value and thus the future realized futures variance associated with that lower VIX level would decrease (as explained in the previous paragraph). Second, to the extent that this downward trend is not expected to be unusually erratic, the variance of the VIX futures price *returns* would also

Exhibit 2

Daily Closing Levels of the VIX and Implied Convexity (IC) from August 2008 to July 2015

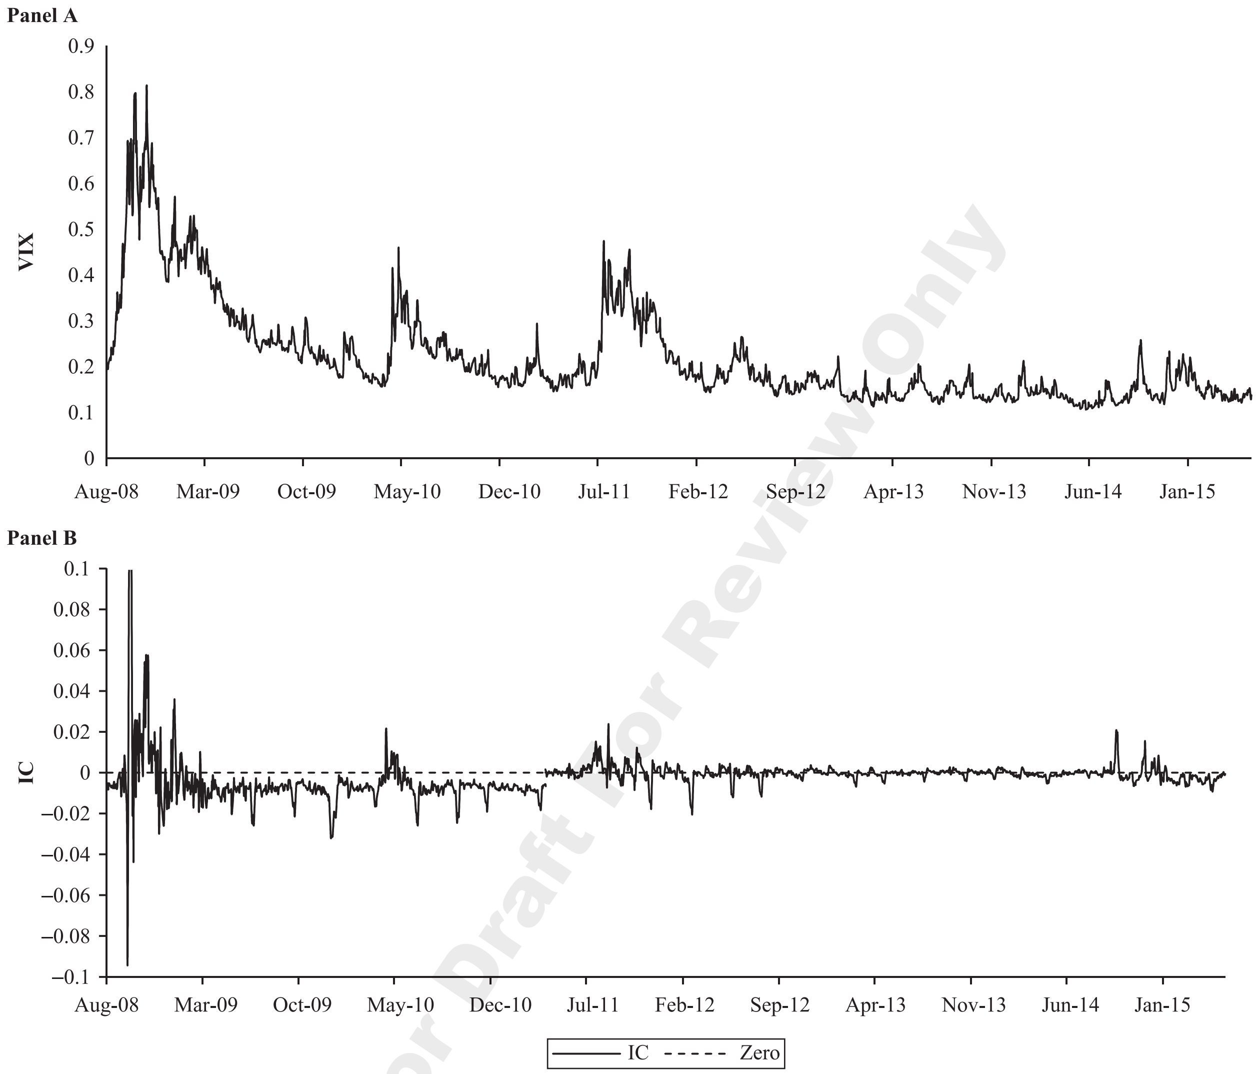

Notes: This exhibit shows the historical level of the VIX (Panel A) and of the implied convexity (Panel B) from August 2008 to July 2015. The "Zero" gradient line corresponds to the lower bound of IC. For illustration purposes, IC levels are capped at 0.1.

decrease. Because the implied convexity is an estimate of the futures *price* variance, and because the variance of prices is a function of both the variance of returns and of price levels,[7] the combination of these two effects will lower the convexity adjustment, thus accelerating the mean reversion of the implied convexity compared with that of the underlying VIX.

Exhibit 2 also shows that the implied convexity can sometimes dip below zero, a counterintuitive phenomenon considering the fact that the theoretical variance of the VIX futures prices can never be negative. By definition, implied convexity should always be positive, because it is a forecast of the *variance* of VIX futures prices from time t until contract expiration, According to Equation (15), a negative implied convexity occurs when the square of the VIX futures prices F_t^2 is larger than the forward-starting variance swap rate estimated in Equation (14).[8] The fact that negative implied convexity values are a common occurrence in practice shows that VIX futures often trade above their theoretical upper bounds.

The resulting violations from the theoretically impossible negative implied convexity values suggest violations of the "strict" theoretical argument. Consequently, investment strategies using VIX futures should consider implications of said "mispricing" for the management of the portfolio. Additionally, a theoretical pricing violation in this context does not have the same consequences as it would in a cost-of-carry setting, where arbitrage strategies quickly bring prices back in equilibrium. A few considerations make these violations difficult to take advantage of directly. First, trading the forward-starting variance swap rate entails buying a continuum of options of two different maturities, a task that would involve large transaction costs. Second, because there is no cost-of-carry relationship, an arbitrage strategy involving the buying of the forward-starting variance swap rate and the selling of the VIX futures contract—when the convexity adjustment would presumably be negative and the square of the VIX futures price "wrongly" above the forward-starting variance swap rate—would still not be able to guarantee a profit, given that the rate at which one would borrow the funds needed to purchase the portfolio of options designed to replicate the forward-starting variance swap rate would impact the final profit. Finally, and most importantly, the strategy would call for the need to sell the square of the futures contract. Although this at first might appear impossible, one could technically sell 15 VIX futures contracts when they trade at a price of $15. As the futures price begins changing, however, dynamic adjustments would have to be made, thereby increasing trading costs further. Additionally, decimals in the price would not be accounted for, turning the strategy into an approximation exercise that could potentially wipe off profits.

Exhibit 3 displays the summary statistics for the implied convexity levels by year in Panel A and by quintile in Panel B. Panel A shows that implied convexity exhibited its highest value of 0.2286 as well as its lowest value of –0.0816 during the last four months of 2008. Since 2008, the volatility in implied convexity has generally declined. Overall, implied convexity exhibits positive skewness and a relatively large kurtosis. During some years (2010, 2012, and 2013), however, implied convexity exhibited negative skewness. The results in Panel B also demonstrate that high levels of positive implied convexity (quintile 5) possess a very large range of values. These results are also evident in Exhibit 4, which shows the implied convexity's frequency distribution over the sample period. In particular, Exhibit 4 points at a large aggregation of data points around zero that appear almost normally distributed, as well as the presence of a large number of positive outliers.

EXHIBIT 3

Summary Statistics for Daily Closing Levels of Implied Convexity

	No. obs.	Mean	Median	Min.	Max.	Std. Dev.	Skewness	Kurtosis
Panel A: Summary Statistics for Entire Sample and by Year								
All	1,693	*0.0007*	–0.0002	–0.0816	0.2286	0.0122	11.32	189.75
2008	87	*0.0185*	0.0059	–0.0816	0.2286	0.0462	2.91	10.67
2009	249	–0.0003	–0.0004	–0.0184	0.0432	0.0069	2.46	13.38
2010	251	–0.0001	0.0008	–0.0248	0.0286	0.0068	–0.57	4.64
2011	251	*0.0006*	0.0003	–0.0176	0.0232	0.0046	0.24	4.75
2012	246	*–0.0007*	–0.0001	–0.0203	0.0051	0.0037	–2.54	8.34
2013	249	*–0.0003*	–0.0001	–0.0068	0.0034	0.0016	–1.22	3.11
2014	248	–0.0001	–0.0006	–0.0066	0.0205	0.0034	3.40	16.58
2015	147	*–0.0025*	–0.0026	–0.0092	0.0084	0.0031	0.97	2.08
Panel B: Summary Statistics by Quintile								
Quintile 1	339	*–0.0065*	–0.0042	–0.0816	–0.0022	0.0066	–5.45	51.47
Quintile 2	338	*–0.0014*	–0.0013	–0.0021	–0.0007	0.0004	–0.14	–1.07
Quintile 3	339	*–0.0002*	–0.0002	–0.0007	0.0004	0.0003	0.18	–1.27
Quintile 4	338	*0.0010*	0.0010	0.0004	0.0019	0.0004	0.34	–1.03
Quintile 5	339	*0.0103*	0.0040	0.0019	0.2286	0.0234	6.90	55.30

Notes: Panel A presents the summary statistics for the daily closing levels of implied convexity for the entire sample and by year. Panel B shows the implied convexity statistics by quintiles, where quintile 1 is the lowest quintile and quintile 5 is the highest quintile. Means that are statistically different from zero ($p = 0.05$) are in italics.

E X H I B I T 4

Frequency Distribution of Implied Convexity

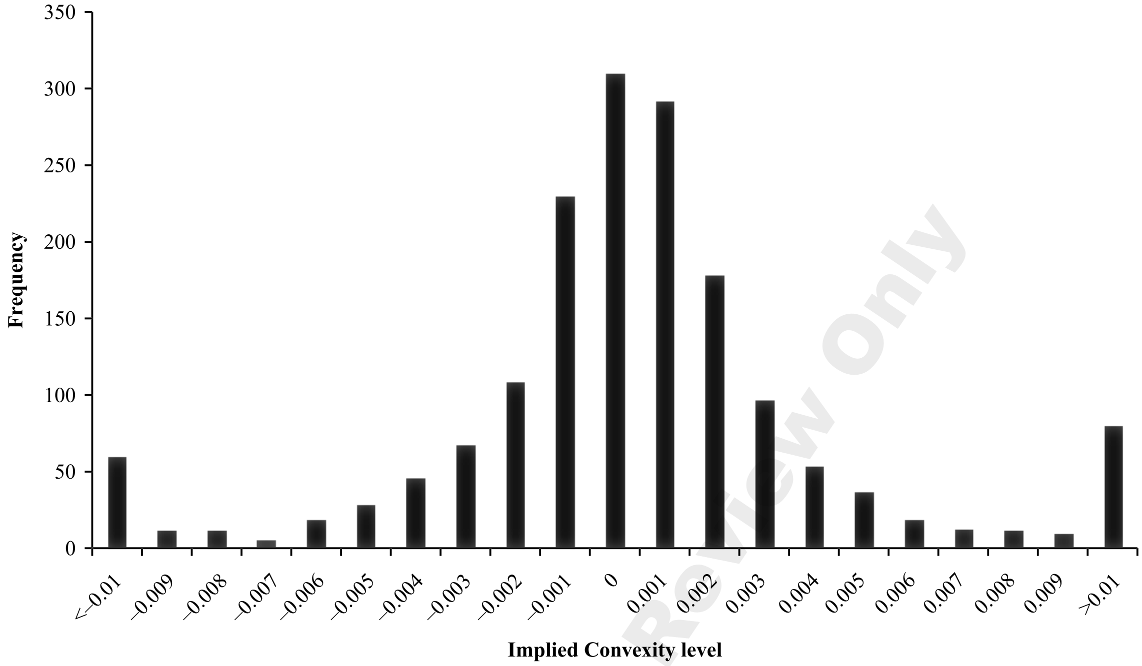

Note: This exhibit shows the distribution of the implied convexity daily closing values for a period going from August 2008 to July 2015.

Exhibit 5 presents the overall historical level percentiles of implied convexity, as well as by year. Over the sample time period, the median daily closing implied convexity level is −0.0002. The 50th, 75th, and 95th percentile ranges show that implied convexity oscillates within a narrow range the majority of the time. As market volatility continuously evolves, however, so does implied convexity. The results in Exhibit 5 clearly show that substantial year-to-year variations in implied convexity levels exist. Such results are to be expected, because the VIX itself fluctuates considerably during the sample period, creating an increase in the volatility of the VIX and related VIX futures prices, which in turn directly affects implied convexity levels. Although negative implied convexity values are found in all of the sampled years, since 2008 the median implied convexity levels have generally remained very close to zero. In fact, the 5% and 95% cutoff values for implied convexity in Exhibit 5 have generally trended toward zero during this period. Overall, these results show that the distribution of convexity values has evolved over time.

Implied Convexity as an Estimate of the Future Realized VIX Futures Variance

The implied convexity component should reflect the market's expectation of the future realized variance of VIX futures prices between the present time and contract expiration. Panel A of Exhibit 6 displays the mean value of the *future* VIX futures price variance and the implied convexity by calendar days left until VIX futures contract expiration. The exhibit shows that both the VIX futures variance and the implied convexity steadily decline as the VIX futures contract becomes shorter in duration, which is expected with less time remaining until the futures contract's expiration. As such, Panel A shows that implied convexity is positively related to VIX futures time to expiration. This relationship brings into question whether implied convexity is related to other characteristics of the VIX futures contract or its underlying asset, such as the volume (liquidity) of the VIX futures and the presence of jumps in the VIX series.[9]

Panel A of Exhibit 6 also shows that implied convexity exhibits a jagged pattern and displays a dip below

EXHIBIT 5
Normal Ranges for daily closing levels of implied convexity over the sample period

Year	No	5%	10%	25%	50%	75%	90%	95%	Range
All	1,693	–0.0079	–0.0042	–0.0017	–0.0002	0.0014	0.0040	0.0087	0.0166
2008	87	–0.0188	–0.0093	–0.0004	0.0059	0.0229	0.0605	0.1132	0.1320
2009	249	–0.0099	–0.0067	–0.0026	–0.0004	0.0015	0.0041	0.0082	0.0181
2010	251	–0.0137	–0.0084	–0.0014	0.0008	0.0022	0.0046	0.0095	0.0232
2011	251	–0.0064	–0.0026	–0.0013	0.0003	0.0020	0.0058	0.0086	0.0150
2012	246	–0.0106	–0.0031	–0.0016	–0.0001	0.0011	0.0025	0.0033	0.0139
2013	249	–0.0038	–0.0016	–0.0009	–0.0001	0.0006	0.0014	0.0021	0.0059
2014	248	–0.0038	–0.0028	–0.0016	–0.0006	0.0006	0.0019	0.0048	0.0086
2015	112	–0.0066	–0.0059	–0.0044	–0.0026	–0.0010	0.0005	0.0033	0.0099

Notes: This exhibit shows the distribution of daily closing levels of implied convexity by percentiles for the entire dataset and for each year of the data, as well as the range between the 95th and 5th percentiles.

EXHIBIT 6
Realized VX Variance and the Implied Convexity

Panel A

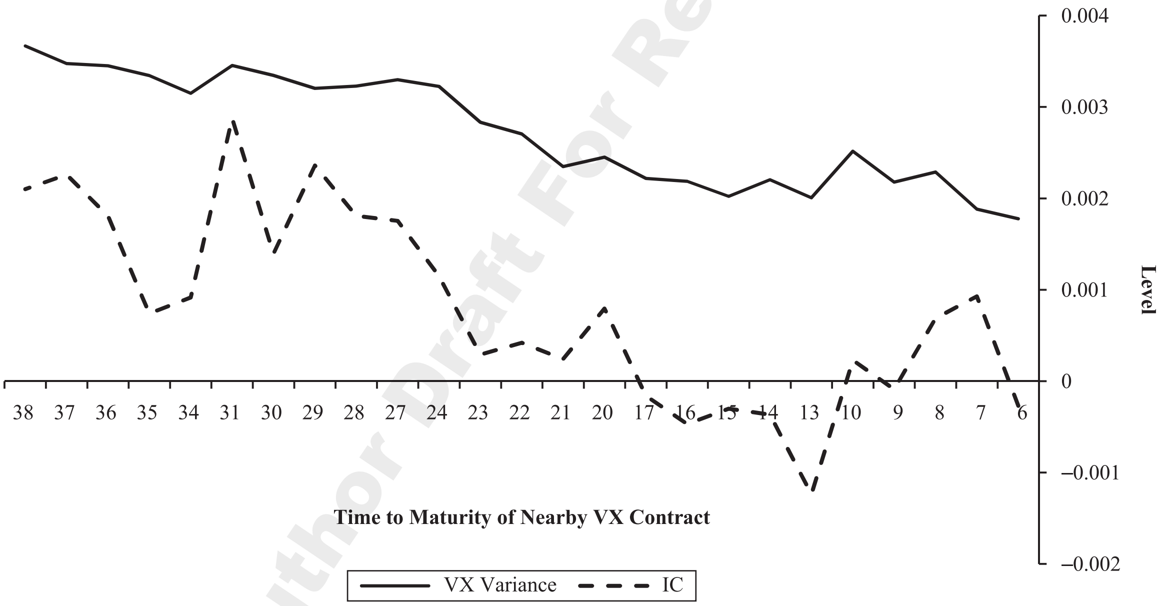

Panel B

	VX Variance
Constant	0.002
	(17.65)
Implied Convexity	0.20
	(3.88)
VX Maturity	0.00004
	(7.76)
Adj. R^2	0.91

EXHIBIT 6 (*Continued*)

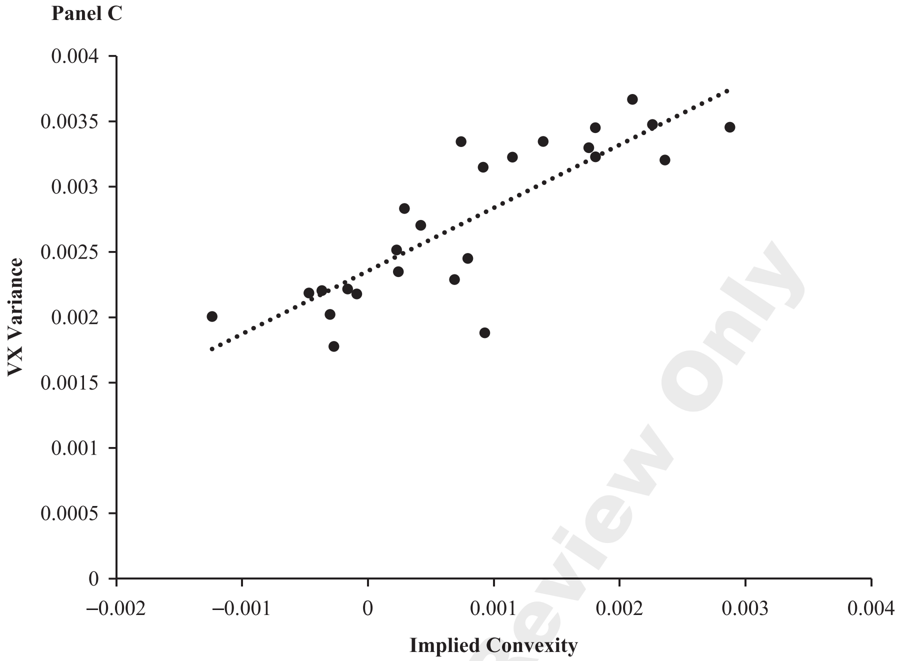

Notes: Panel A of Exhibit 6 shows the average implied convexity daily closing levels (IC) and the realized VIX futures (VX) variance of the nearby contract averaged by days remaining to VX maturity. Panels B and C show the regression results where the VX variance is regressed against the average implied convexity by days remaining to VX maturity.

zero with 13 days to expiration.[10] Alternatively, the future *realized* VIX futures variance exhibits a smoother pattern and by definition remains above zero even when the VIX futures contract is close to expiration. Panel B of Exhibit 6 reports results when the future realized variance is regressed on the implied convexity and the maturity of the contract. All coefficients, including the constant, are highly statistically significant, with an adjusted R-squared of 91%. Overall, the implied convexity forecasts the VIX futures variance very closely, as can also be seen graphically in Panel C. Therefore, implied convexity can on average be used as a reliable and accurate forecast of future volatility of volatility. In this context, its use as a proxy for the volatility of volatility can prove useful in models that need to estimate a forward-looking measure of volatility of volatility, such as VIX options pricing models.

Violations of the VIX Futures Upper Bound

Previously, we established that implied convexity values can sometimes be negative, indicating that the VIX futures contract can trade above its theoretical upper bound. To examine this phenomenon more closely, we determine the VIX futures price that would eliminate this violation (i.e., the value that would make the negative implied convexity values equal to zero) for each day in the sample. We find that this adjustment value is almost always outside of the VIX futures bid–ask spread (92% of all violations), except for a few cases when the violation is very small.

Exhibit 7 shows the frequency distribution of the adjustment relative to the size of the bid–ask spread, showing that the majority of the time the adjustment is in excess of 100% of the bid–ask spread, and that it is often much larger. Given the significant size of the adjustment needed to correct for the VIX futures upper bound violation, we explore its relationship to implied convexity levels. Exhibit 8 shows the results of a regression where the adjustment is the dependent variable and the implied convexity is the independent variable. Overall, the results show a significantly large regression R-squared of 78% (adjusted for VIX futures maturity), meaning that implied convexity is strongly related to the upper bound adjustment. When the data are separated into quintiles based on the size of the implied convexity, the results remain statistically

EXHIBIT 7

Frequency Distribution of the Downward VIX Futures Price Adjustment Needed to Avoid a VIX Futures Upper Bound Violation (as a Percentage of the VIX Futures Bid–ask Spread)

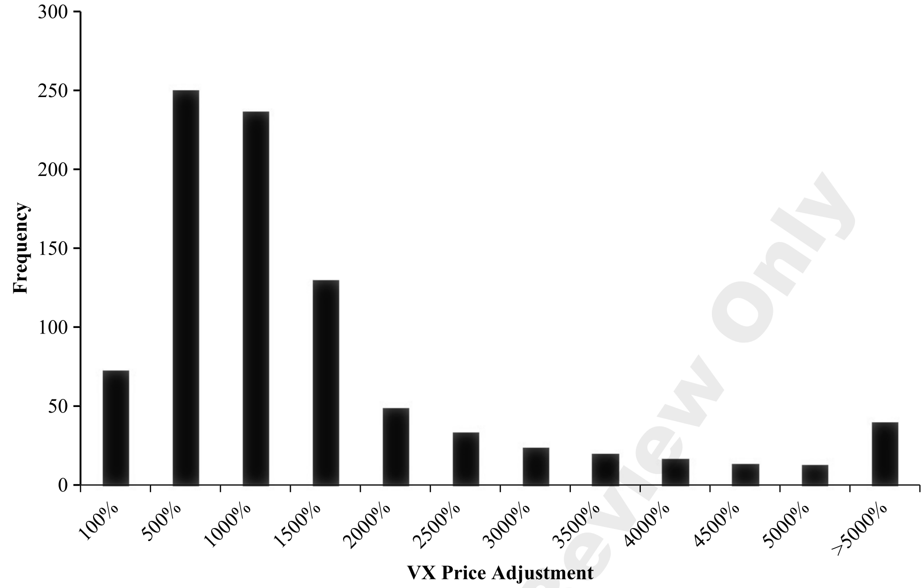

Notes: This exhibit shows the frequency distribution for the VIX futures downward adjustment needed in order to prevent an upper bound violation of the VIX futures, as defined by Carr and Wu [2006]. The adjustment is expressed as a percentage of the VIX futures bid–ask spread. Only violations that fall outside of the bid–ask spread are shown (94.93% of all violations), all of which fall below the VIX futures bid price.

EXHIBIT 8

Regression Results for Adjustment Needed To Correct VIX Futures Upper Bound Violation Regressed on the Implied Convexity

	Quintile 1	Quintile 2	Quintile 3	All
Constant	0.01	0.001	0.0001	0.004
	(10.10)	(3.86)	(1.25)	(10.72)
Implied Convexity	–1.34	–2.48	–2.91	–1.59
	(–26.95)	(–17.99)	(–31.04)	(–55.86)
VX Maturity	–0.0001	–0.00002	–0.00001	–0.0001
	(–3.53)	(–3.34)	(–2.32)	(–4.40)
Regression Statistics				
Adj.R^2	0.70	0.49	0.82	0.78
No. obs.	339	338	224	901

Notes: This exhibit shows the results of the regression where the adjustment in the VIX futures price (needed to correct the VIX futures upper bound violation that occurs when implied convexity is negative) is the dependent variable and the implied convexity closing level and VX contract maturity are the independent variables. T-values are reported in parentheses below the regression coefficient. Note that the total number of observations is based on the number of negative implied convexity values. Therefore, only the first three quintiles have observations, and the number of observations here is lower than the total number of observations in the sample period. Quintile 1 represents the smallest implied convexity values.

significant, and the adjusted R-squared values are high. Therefore, implied convexity is an important explanatory variable for the size of the VIX futures upper bound violations.

Implied Convexity in Relation to Stock Market Volatility

In this section, we perform an additional verification of the theory that market volatility and implied convexity should be related. Theoretically, if market volatility is high, the VIX is high, VIX futures prices are high, and consequently, the variance of the VIX futures would also be high. The opposite relationship is true when market volatility is low. Thus, a direct link should exist between implied convexity and various market volatility measures.

Exhibit 9, Panel A, displays regression results between implied convexity and the *prior* 30-day historical variance of the S&P 500, computed as the average squared return.[11] The purpose of this regression is to measure the extent to which prior volatility in the market is associated with current implied convexity. The results show that recent realized market volatility and current

Exhibit 9

Regression Results for Implied Convexity Regressed on the Prior Month for the S&P 500 Return Variance and on the VIX Daily Closing Levels

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5	All
Panel A						
Constant	–0.43	–0.07	–0.01	0.04	0.59	–0.10
	(–17.38)	(–35.77)	(–11.23)	(24.32)	(3.57)	(–3.41)
S&P 500 Variance	–9.10	–0.29	–0.32	0.64	46.33	39.81
	(–3.41)	(–0.56)	(–0.90)	(2.33)	(6.91)	(15.54)
VX Maturity	0.02	0.002	0.0005	–0.001	–0.03	0.001
	(10.81)	(19.89)	(8.70)	(–14.81)	(–4.06)	(0.54)
Regression statistics						
Adj. R^2	0.26	0.54	0.18	0.41	0.19	0.12
No. obs.	339	338	339	338	339	1,693
Panel B						
Constant	–0.37	–0.07	–0.01	0.03	0.07	–0.32
	(–11.03)	(–24.27)	(–7.81)	(15.28)	(0.32)	(–8.16)
VIX	–0.36	–0.02	0.002	0.02	2.48	1.45
	(–3.44)	(–1.59)	(0.23)	(3.32)	(6.37)	(12.56)
VX Maturity	0.01	0.002	0.0005	–0.001	–0.02	0.001
	(10.58)	(20.03)	(8.69)	(–14.42)	(–3.88)	(0.50)
Regression statistics						
Adj. R^2	0.26	0.54	0.18	0.42	0.17	0.08
No. obs.	339	338	339	338	339	1,693

Notes: Panel A provides the regression where the dependent variable is the daily closing level of implied convexity and the independent variables are the S&P 500 return variance over the prior 30 days (computed as the average squared return) and the time to maturity of the nearby VX contract. Panel B displays the results of the regression where the dependent variable is the daily closing level of implied convexity and the independent variables are the daily closing level of the VIX and the time to maturity of the nearby VX contract. The data are annualized for consistency. Quintile 1 represents the smallest implied convexity values and t-statistics are reported in parentheses below the regression coefficients.

implied convexity levels possess a positive and significant relationship in general, with the historical variance of the S&P 500 index explaining approximately 12% of the variation in implied convexity for the entire time period when controlling for the VIX futures horizon of the nearby contract. Grouping the data by implied convexity quintiles shows that only the very high (quintiles 4 and 5) and very low (quintile 1) levels of current implied convexity are significantly related to recent realized historical market volatility, with 41%, 19%, and 26% adjusted R-squared values, respectively. As such, the results show that the relationship between historical market volatility and current implied convexity becomes evident for extreme values of implied convexity; otherwise, any relationship is undetectable.

A similar relationship is found between implied convexity and implied market volatility. The results in Exhibit 9, Panel B, show that implied market volatility (the VIX) has a positive and significant relationship with the implied convexity when the implied convexity is very high (quintiles 4 and 5) and very low (quintile 1). Adjusting for the nearby VIX futures contract horizon, implied market volatility explains approximately 42%, 17%, and 26% of the total variation in implied convexity for quintiles 4, 5, and 1, respectively. The relationship between VIX and implied convexity can be visualized in Exhibit 10, where the positive relation is evident for very large and very small values of implied convexity (or of the VIX). Additionally, the exhibit shows that large absolute values of implied convexity are almost always positive. Overall, the results show that when implied convexity deviates largely from zero, its levels contain information associated with historical and expected market volatility.

Exhibit 10
Relationship between VIX and Implied Convexity

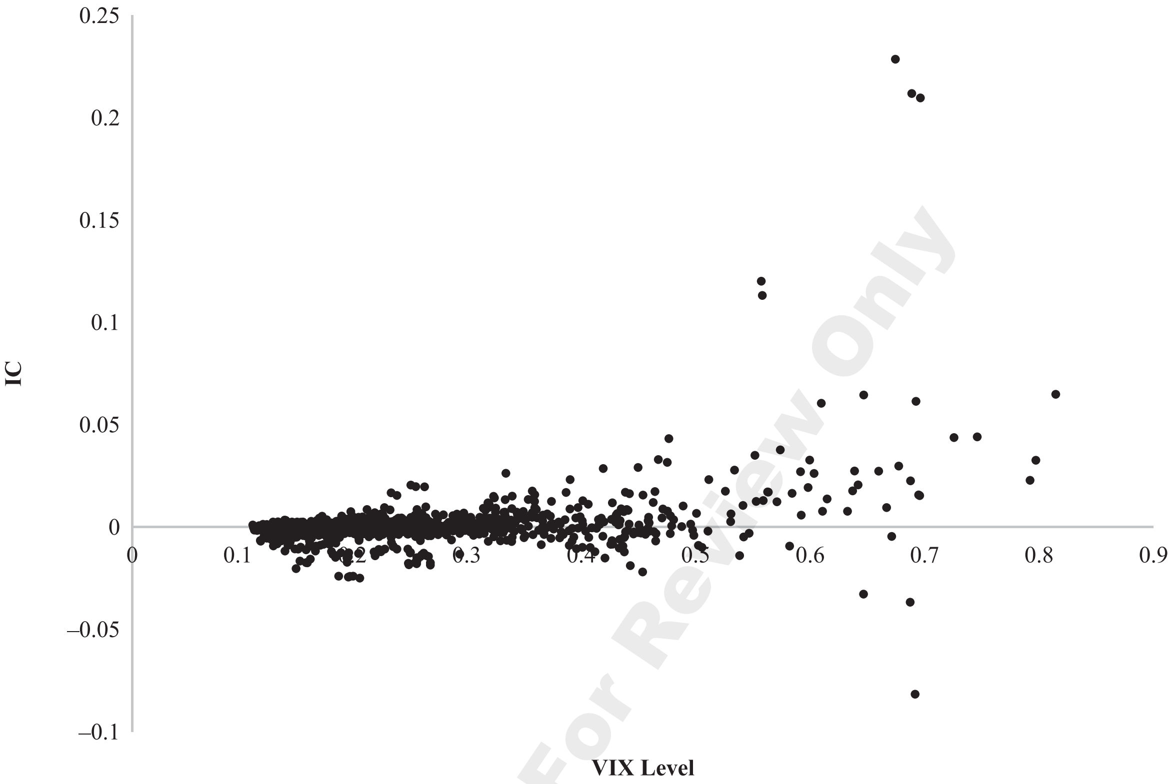

Notes: This exhibit plots the implied convexity's daily closing values (vertical axis) against the VIX daily closing values (horizontal axis). The sample period extends from August 2008 to July 2015.

POTENTIAL VIOLATIONS OF THE MODEL

All models are based on assumptions designed to allow for tractable closed-form solutions. The Carr and Wu [2006] approach of presenting a lower and upper bound for the price of the VIX futures contract is based on several such assumptions. In conjunction with Jensen's inequality, the lower bound for the VIX futures price is the forward-starting volatility swap rate, and the upper bound is the square root of the forward-starting variance swap rate. In the following sections, we discuss the Carr and Wu [2006] simplifying assumptions. Although each assumption can be empirically analyzed in detail, we argue how possible assumption violations should be minimal; therefore, rare violations should not materially affect the results. Moreover, an extensive examination of each assumption is beyond the scope of this article, is constrained by article length, and is left for future research if needed.

Jumps in the Volatility Process

One element to consider is the fact that the proof in Carr and Wu [2006] assumes that the volatility price process is continuous. If instead one allows for jumps, the future realized variance is then a combination of the variances arising from the quadratic variation of the diffusion component of the process and the quadratic variation of the jump component of the process. In fact, Broadie and Jain [2008] noted that the effect of ignoring jumps in computing the fair variance swap rate from the portfolio of S&P 500 options used in the VIX definition could be significant. Therefore, we proceed to test for the presence of jumps in the (spot) VIX as a means to verify whether the continuous volatility assumption in Carr and Wu [2006] could be violated.

Various non-parametric methods designed to identify the occurrence of jumps in stochastic processes have emerged in the last few years (for example, see Ait-Sahalia

[2002], Carr and Wu [2003], Barndorff-Nielsen and Shephard [2006], and Jiang and Oomen [2008], to name a few). We choose to implement the straightforward test found in Lee and Mykland [2008], both for its intuitive appeal as well as for its ability to outperform the non-parametric jump tests of Barndorf–Nielsen and Shephard [2006] and Jiang and Oomen [2008].

The statistic *L(i)* tests at time t_i whether there is a jump in the process *S(t)* from t_{i-1} to t_i and is defined as follows:

$$L(i) = \frac{\ln[S(t_i) / S(t_{i-1})]}{\widehat{\sigma}(t_i)} \tag{12}$$

where

$$\widehat{\sigma}(t_i) = \frac{1}{K-2} \sum_{j=1-(K-2)}^{i-1} \times \left|\ln[S(t_j) / S(t_{j-1})]\right| \left|\ln[S(t_{j-1}) / S(t_{j-2})]\right| \tag{13}$$

Using daily frequencies, we adopt the Lee and Mykland [2008] recommended optimal window size K of 16 days. For n, the number of observations, and $c = \sqrt{2}/\sqrt{\pi}$, let

$$C_n = \frac{[2\ln(n)]^{1/2}}{c} - \frac{\ln(\pi) + \ln[\ln(n)]}{2c[2\ln(n)]^{1/2}} \text{ and}$$
$$S_n = \frac{1}{c[2\ln(n)]^{1/2}} \tag{14}$$

The actual test is then whether the ratio $\frac{|L(i)| - C_n}{S_n}$ is above the critical value corresponding to the chosen significance level. Because under the null hypothesis of no jumps, the ratio has a cumulative distribution function of the form $\exp(-e^{-x})$, the critical value corresponding to the 1% upper tail will therefore be $\beta^* = -\ln(-\ln(0.99))$.

We test for the presence of jumps in the VIX over our entire sample period at the 5% significance level. We determine that a total of only 17 likely instances of jumps occur. Because the sample period contains 1,735 observations (trading days), this is the equivalent of less than 2.5 days per year. By comparison, Lee [2012] investigated the predictability of jump arrivals in U.S. stock markets, identifying the number of jumps experienced by individual equities from the DJIA and the S&P 500 Index from January 4, 1993, to December 21, 2008. The associated average number of jumps per security is 21.79 per year. Therefore, we conclude that the assumption of a continuous process by Carr and Wu [2006] is a reasonable one for the VIX, at least in the context of our study. Accordingly, our computed variance swap rate is consistent with the conclusion that it is essentially free of any jump-related bias.

Finally, for completeness, Exhibit 11 shows the results of a multivariate regression where implied convexity is regressed against the presence of jumps in the VIX (dummy variable, 1 for jump and 0 otherwise) and the volume of the nearby VIX futures contract. Like before, we control for the number of days left to maturity of the nearby VIX futures contract. The results show that implied convexity is significantly related to the VIX futures' volume, whether implied convexity is negative or positive. However, implied convexity is not related to the presence of a jump in the VIX series.

Approximation Error

The Carr and Wu [2006] lower and upper bound proof is based on a continuous range of option strikes, whereas the actual VIX (and VIN and VIF) employ

EXHIBIT 11

Relation of implied convexity (IC) to the presence of jumps in the VIX, VIX futures volume, and the VIX futures time to expiration

	IC	IC–	IC+
Constant	–0.0003	–0.01	0.01
	(–0.44)	(–13.71)	(7.16)
VIX Jumps	0.001	0.001	0.001
	(0.27)	(0.91)	(0.16)
VX Volume	–1.75E–8	2.88E–8	–4.28E–8
	(–2.24)	(6.02)	(–3.21)
VX Maturity	0.0001	0.0001	–0.0002
	(2.56)	(3.58)	(–3.69)
Regression Statistics			
Adj. R^2	0.005	0.06	0.03
No. obs.	1,693	889	794

Notes: This exhibit displays the results of the regression where the implied convexity is the dependent variable and the presence of jumps in the VIX (dummy variable), VIX futures volume (contract liquidity), and VIX futures time to expiration are the independent variables. IC– and IC+ denote negative and positive values of implied convexity, respectively, while t-statistics are reported in parentheses below the regression coefficients. Regression coefficients for VX Volume are shown in scientific notation for readability.

discrete strikes for out-of-the-money options with non-zero bids until two adjacent non-zero bids occur. Jiang and Tian [2007] examined biases associated with the discrete and truncation calculation procedure of the CBOE, leading to a misspecification of true volatility, especially in regards to the term structure of volatilities. This bias translates to the VIN being relatively more downward biased than the VIF at high volatility levels, causing the upper bound to be downward biased. However, the importance of this bias is suspect, because the simulation study employed by Jiang and Tian [2007] to achieve the results employs an unrealistically large range of strike prices relative to real-life strikes. Moreover, only the relative difference between VIN and VIF would affect the term structure of volatility. Consequently, using the actual strike series employed by the markets would have a minimal bias on the outcomes found in this study.

Settlement Procedure

As examined in Pavlova and Daigler [2008], a settlement bias exists due to the procedure employed to determine the individual option prices used to calculate the VIX futures settlement price. In particular, any option employed to calculate the VIX at the open on the Wednesday settlement day uses the trade price at 8:31 a.m. Central Time if a trade exists, rather than the average of the bid–ask price used to calculate the VIX, VIN, and VIF indexes. Pavlova and Daigler [2008] determined that large differences existed through mid-May 2007 for certain expirations between the spot VIX and the VIX futures settlement value. Only 6 of 33 months, however, had biases where the VIX futures price was larger than the spot VIX at futures expiration (i.e., situations consistent with a negative implied convexity). Thus, biases due to settlement in past expirations do not generally promote negative implied convexities. Moreover, market participants report that the settlement bias in recent years is near zero. Consequently, previous literature does not support the settlement bias causing the negative implied convexities.

CONCLUSION

The most interesting aspect of VIX futures prices is a convexity adjustment whose value is an estimate of their future realized variance between now and the expiration date of the contract, bridging the gap between the forward-starting variance swap rate and the squared VIX futures price. Rather than positing a stochastic process for the VIX, we approach the problem in a novel way by solving for the convexity implied by the difference between the forward-starting swap rate (determined by the adjusted value of the next two expirations of the S&P 500 options series) and the square of the current VIX futures price and proceed to evaluate its predictive and statistical properties.

Our main result is that the implied convexity adjustment can indeed be a useful forecast of the future volatility of VIX futures levels. The most important implication of our study is that implied convexity is a generally reliable and accurate estimate of future VIX volatility, which can be used as an input to any derivative model involving the volatility of volatility as in the case of VIX options pricing models. The implied convexity consistently, albeit at a small scale, underestimates realized VIX futures variance—a phenomenon that is the opposite of the well-known relationship between implied volatility and realized volatility.

Examining the characteristics of the implied convexity adjustment using daily data, our results also show that the implied convexity daily values can sometimes be negative, falling outside the theoretical range; however, we are able to rule out arbitrage opportunities. We also confirm the positive relationship between implied convexity and various market volatility measures and rule out a series of possible model assumption violations, such as jumps, approximation error, and settlement procedure.

ENDNOTES

[1]Whereas this convexity adjustment might not be "implied" in the pure conventional model-dependent sense, it is nevertheless implied by the difference between the forward-starting variance swap rate and the square of the VIX futures price.

[2]Tradable out-of-the-money S&P 500 options do determine the underlying spot VIX. However, several factors make this spot "asset" difficult to trade on a risk-free basis for arbitrage purposes. One major factor is that the portfolio of options would need to be dynamically traded to make it always equivalent to the underlying spot index used for VIX futures settlement. Such trading would be costly; moreover, many of the needed options have limited liquidity. Another issue includes the uncertainty that the settlement procedure creates; that is, actual option prices traded at the opening of settlement day are employed to calculate the settlement price rather than the bid–ask average used to calculate the spot VIX.

[3]An example of empirically testing a VIX futures model is by Zhang and Zhu [2006], who developed and tested a stochastic variance model for the evolution of the VIX over time. However, their model overprices VIX futures by 16%–44%. When using only one year of data, parameter fitting reduces the errors to 2%–12%. Additionally, the time period they employed was a very low volatility period, unlike any time period that would include 2008. Zhu and Lian [2012] developed a VIX futures pricing model displaying stochastic volatility and simultaneous jumps for both the underlying asset and the corresponding volatility, with the Heston [1993] stochastic volatility model as the underlying volatility process. They found that jumps in the underlying asset improve the pricing of VIX futures, whereas jumps in volatility do not. Their empirical comparisons of the various models and their individual factors are interesting. The results of their model with different volatility processes typically show errors in excess of 5%, however, and the data series ends before the volatile period in Fall 2008.

[4]Note that the variances may or may not be annualized, but they should reflect the same frequency of returns. For instance, the variances could reflect daily, weekly, or monthly returns, as long as they are all of the same frequency type. Thus, one can easily annualize all the variances by simply multiplying both sides of Equation (8) by the proper scaling factor.

[5]The VIN and the VIF are based on the S&P 500 index option contracts that expire on the third Friday of the expiration month. Alternatively, VIX futures contracts expire on the Wednesday that is 30 days prior to the third Friday of the calendar month immediately following the month in which the futures contract expires. Additionally, there are also some unusual cases where the VIN and the VIF do not “straddle” the VIX (in terms of expiration dates) due to the timing of the rollover of the S&P 500 option contracts occurring the week before the expiration of the VIX futures contract, creating even larger timing differences between the expirations of the futures contract, the VIN, and the VIF.

[6]Data are not available for the nearby and far-term VIX series prior to August 25, 2008.

[7]Given a certain return volatility, a higher price will imply fluctuation levels of larger magnitude, translating into a higher price variance.

[8]The upper bound of the VIX futures price is equal to the square root of the forward-starting variance swap rate, as shown in Carr and Wu [2006]. This is because the price of the VIX futures is the expected value of the square root of future expected variance, which is always smaller than or equal to the square root of the expected value of future expected variance. As such, implied convexity is negative if and only if the VIX futures price is above its theoretical upper bound and positive if and only if the VIX futures price is below its theoretical upper bound.

[9]Jumps are examined in detail in another section.

[10]Thus, in Exhibit 6, Panel A, the *average* values for the implied convexities for each day before expiration are positive (except for day 13), whereas some extreme observations in the sample can possess negative values.

[11]The regressions between implied convexity and different market measures use a constant one-year horizon implied convexity. To obtain the one-year constant maturity implied convexity, we divide the implied convexity backed out of Equation (15) by the number of calendar days left to the VIX futures expiration and multiply it by 365.

REFERENCES

Ait-Sahalia, Y. “Telling from Discrete Data Whether the Underlying Continuous-Time Model Is a Diffusion.” *Journal of Finance*, 57 (2002), pp. 2075-2112.

Barndorff-Nielsen, O.E., and N. Shephard. “Econometrics of Testing for Jumps in Financial Economics using Bipower Variation.” *Journal of Financial Econometrics*, 4 (2006), pp. 1-30.

Broadie, M., and A. Jain. “The Effect of Jumps and Discrete Sampling on Volatility and Variance Swaps.” *International Journal of Theoretical and Applied Finance*, 11 (2008), pp. 761-797.

Carr, P., and L. Wu. “What Type of Process Underlies Options? A Simple Robust Test.” *Journal of Finance*, 58 (2003), pp. 2581-2610.

——. “A Tale of Two Indices.” *The Journal of Derivatives*, 13 (2006), pp. 13-29.

CBOE. “The VIX White Paper.” CBOE, 2003. Available at www.cboe.com/micro/vix/vixwhite.pdf.

Dupoyet, B., R. Daigler, and Z. Chen. “A Simplified Pricing Model for Volatility Futures.” *Journal of Futures Markets*, 31 (2011), pp. 307-339.

Grunbichler, A., and F.A. Longstaff. “Valuing Futures and Options on Volatility.” *Journal of Banking and Finance*, 20 (1996), pp. 985-1001.

Heston, S. “A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options.” *Review of Financial Studies*, 6 (1993), pp. 327-343.

Huskaj, B., and M. Nossman. “A Term Structure Model for VIX Futures.” *Journal of Futures Markets*, 33 (2013), pp. 421-442.

Jiang, G., and Y. Tian. "Extracting Model-Free Volatility from Option Prices: An Examination of the VIX Index." *The Journal of Derivatives*, 14 (2007), pp. 1-26.

Jiang, G.J., and R.C. Oomen. "Testing for Jumps when Asset Prices Are Observed with Noise—Swap Variance Approach." *Journal of Econometrics*, 144 (2008), pp. 352-370.

Lee, S. "Jumps and Information Flow in Financial Markets." *Review of Financial Studies*, 25 (2012), pp. 439-479.

Lee, S., and P. Mykland. "Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics." *Review of Financial Studies*, 21 (2008), pp. 2535-2563.

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

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

Mencia, J., and E. Sentana. "Valuation of VIX Derivatives." *Journal of Financial Economics*, 108 (2013), pp. 367-391.

Pavlova, I., and R.T. Daigler. "The Non-Convergence of the VIX Futures at Expiration." *Review of Futures Markets*, 17 (2008), pp. 201-223.

Poon, S.H., and C.W.J. Granger. "Forecasting Volatility in Financial Markets: A Review." *Journal of Economic Literature*, 41 (2003), pp. 478-539.

Wang, Z., and R. Daigler. "The Option SKEW Index and the Volatility of Volatility." Presentation, Financial Management Association Annual Conference, Chicago, IL, 2013.

Zhang, J.E., and Y. Huang. "The CBOE S&P 500 Three-Month Variance Futures." *Journal of Futures Markets*, 30 (2010), pp. 48-70.

Zhang, J.E., J. Shu, and M. Brenner. "The New Market for Volatility Trading." *Journal of Futures Markets*, 30 (2010), pp. 809-833.

Zhang, J.E., and Y. Zhu. "VIX Futures." *Journal of Futures Markets*, 26 (2006), pp. 521-531.

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