The implied convexity of VIX futures

Robert T. Daigler

Brice Dupoyet *

Fernando Patterson

March 13, 2014

Keywords: Implied convexity, VIX futures variance, Volatility

JEL classification: G12; G13; G17

* Correspondence author, Chapman Graduate School of Business, Department of Finance, RB 209A, Miami, Florida 33199. Tel: 305-763-6300, Fax: 305-348-4245, E-mail: dupoyetb@fiu.edu

■ Robert T. Daigler is Knight Ridder Research Professor of Finance, Florida International University, RB 206B, Miami, Florida. Email: daiglerr@fiu.edu
■ Brice Dupoyet is an Associate Professor of Finance, Florida International University, RB 209A, Miami, Florida. Email: dupoyetb@fiu.edu
■ Fernando Patterson is a Doctoral Candidate in Finance at Florida International University, Miami, Florida. Email: fpatt001@fiu.edu

The implied convexity of VIX futures

ABSTRACT

We examine the convexity adjustment needed to value VIX futures prices by extracting the implied convexity component from VIX futures prices and the corresponding spot option market prices. The resultant implied convexity value is the missing adjustment between the observed VIX futures price and the forward-starting variance swap rate. We then study the properties of this convexity adjustment, both as a time series and with respect to various market factors. We find that implied convexity often violates no-arbitrage conditions since it is frequently negative, and that implied convexity is strongly statistically related to both market volatility and VIX futures time to expiration. We also investigate the ability of implied convexity to predict realized VIX futures price *variances*, finding that it underestimates the realized VIX futures variance 71% of the time.

1. Introduction

VIX futures pricing is an active area of futures markets research, both because volatility products are considered a new asset class *and* because no VIX futures pricing model is currently accepted as the definitive model by the financial community. The difficulties in pricing VIX futures arise from the need to accurately estimate a number of complicated factors under different market scenarios, the importance of the stability of these factors, and the size of the mispricing for current models. Consequently, the limited market-based empirical evidence does not support the long-term pricing stability for any of the existing models. As Poon and Granger (2003) point out, "much remains to be done to understand the process and characteristics of volatility."

The lack of sustainable empirical support for standard VIX futures pricing models makes this study both interesting and necessary. In this paper we introduce the concept of *implied* convexity of

VIX futures and then examine its properties.[1] One way to think about the convexity of VIX futures pricing is that it provides the "missing piece" needed to help determine how VIX futures prices behave. Although several researchers previously have modeled convexity (discussed shortly), their convexity component is often an approximation value as well as being model-dependent. Therefore, developing and examining implied convexity from market-based values provides an alternative to understanding a key factor affecting the behavior of VIX futures prices. Such an examination is similar to how traders employ implied volatility and therefore it can provide new insights into how to create a more accurate and realistic VIX futures pricing model.

After describing our approach to the VIX futures convexity issue, we examine the empirical properties of the implied convexity series. We find that implied convexity often violates the no-arbitrage condition related to the forward-starting variance swap rate. These results show that institutional traders (or market makers) can potentially earn arbitrage-free profits. Additionally, regressing implied convexity on S&P 500 volatility and on the spot VIX yields a statistically significant relation. Notably, implied convexity increases during highly volatile periods, such as the Fall of 2008. Alternatively, implied convexity decreases as the VIX futures reaches expiration. Finally, we find that implied convexity is associated with future realized VIX futures variance, although implied convexity does consistently underestimate the future variance of VIX futures. This result is similar to the phenomenon whereby the implied volatility of an option consistently overestimates the future realized volatility of the underlying asset.

This article is organized as follows: section 2 examines current VIX futures models and the associated convexity adjustment; section 3 defines the nature of the problem and provides a

[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.

summary of the theory; section 4 describes our methodology and approach to the problem; section 5 details the data used in this study; section 6 examines the results; and section 7 concludes.

2. Current models of VIX futures pricing

Since the underlying cash asset for VIX futures is not traded as an index, the behavior of the VIX futures price series differs from a typical cost-of-carry futures contract.[2] Moreover, the stochastic process of volatility over time is not well understood, leading to difficulties in designing an effective model for VIX futures pricing. Wang and Daigler (2013) examine part of this issue by undertaking an analysis of the volatility of volatility and the behavior of the implied volatility skew. However, they do not apply these factors to a pricing model. In addition, the VIX futures price is the market's estimate of the future VIX level at the expiration of the contract. Therefore, the futures price is affected by the characteristics of the future VIX, as well as those of the spot VIX, and the S&P 500 underlying the options. In turn, these 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 are Lin (2007), Zhu and Lian (2012), and Mencia and Sentana (2013). However, these complex models contain a large number of parameters. Since the parameters change over time in relation to their estimation from different sampling characteristics, the consistency of these models is questionable. For example, Zhu and Lian (2012) state that "Just as an issue raised in Zhang & Huang (2010), we also believe that searching for a reliable estimation method to determine the model parameters from market data

[2] Tradable out-of-the-money S&P 500 options do determine the underlying cash VIX. However, several factors make this cash "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 cash 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, i.e. 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 cash VIX.

remains a challenge. While our work presented in this paper 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, there are a handful of studies, such as Huskaj and Nossman (2012), that estimate the parameters exogenously with an improved rate of success in capturing the term structure of VIX futures compared to previous studies. Other studies, for example Grunbichler and Longstaff (1996), Zhang and Zhu (2006), and Dupoyet, Daigler, and Chen (2011), provide 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). However, the stability of the results for these models over time needs more empirical support.

Zhang and Zhu (2006) develop a stochastic variance model for the evolution of the VIX over time, showing that their model overprices VIX futures by 16-44%. When using only one year of data, parameter fitting reduces the errors to 2-12%. However, the time period in question they employed was a very low volatility period, unlike any time period that would include 2008. Zhu and Lian (2012) develop 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 find 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. However, the results of their model with different volatility processes typically show errors in excess of 5%, and the data series ends before the volatile period in the Fall of 2008.

Another factor affecting VIX futures pricing is the convexity adjustment. A few authors do incorporate the convexity adjustment into their models, using an approximation of the Taylor series and a selected volatility process. For example, Lin (2007) employs a second-order correction approximation, Lu and Zhu (2010) use Kalman filter and maximum likelihood to estimate their model, and Zhang et al. (2010) use a third order approximation based on the Heston stochastic volatility model. Zhu and Lian (2012) compare their "exact" formula to the other approximation methods, finding that these approximations can create large percentage errors. Through the use of simulations, Zhu and Lian also show that the errors increase as volatility increases. Rather than proposing yet another approximation method for convexity, we back out convexity from the actual VIX futures prices, testing this implied convexity for several periods of varying market volatility, including the Fall of 2008. This implied convexity allows us to examine the behavior of the actual convexity adjustment needed to justify the current VIX futures price.

3. Modeling VIX futures pricing and the convexity issue

The pricing of physical commodity futures contracts is based on the cost-of-carry model. Violations of the cost-of-carry model beyond the transaction bounds create arbitrage opportunities. Although the cash VIX is determined from S&P 500 option prices, the conversion from the implied future variance captured by a portfolio of out-of-the-money S&P 500 options to the implied volatility represented by the cash VIX results in a convexity issue. As such, the need arises to estimate the VIX futures price variance (or volatility) from the present time to the contract's expiration. Estimating the futures price variance poses a challenge for pricing purposes, since the future realized variance of VIX futures is unknown by definition and must therefore be estimated. One approach to dealing with this convexity issue is to develop a model to forecast the VIX futures *price* volatility from the present time until the contract's expiration. However, such a model entails assumptions concerning the path of future VIX futures prices and therefore includes model risk. A second approach to examining the convexity of VIX futures prices is to employ actual futures and cash VIX prices and solve for the *implied* convexity of VIX futures. This paper introduces the concept of implied convexity and then examines the associated empirical properties of this new measure.

In this section we first provide a brief summary of the main issues and associated equations related to VIX futures pricing, and then elaborate on our implied convexity approach to VIX futures pricing.[3] The first step in understanding the VIX futures convexity is to define the associated variance. For today at time t, for any random variable x_T whose value is revealed at time $T > t$, we know that

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

Substituting the value of the VIX futures price at expiration, F_T, for x_T into (1) we obtain

$$Var_t(F_T) = E_t(F_T^2) - \left[E_t(F_T)\right]^2 \tag{2}$$

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

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

As noted in the CBOE VIX white paper (2003), the expression $Var_t(F_T)$ is the variance of the VIX futures prices from the present to the futures expiration, which can be viewed as the daily cumulative variance of the VIX futures prices from the current time t to expiration time T. Since $F_t = E_t(VIX_T) = E_t(\sqrt{VIX_T^2}) \leq \sqrt{E_t(VIX_T^2)}$ by reason of the convexity of the square root function, and since $Var_t(F_T)$ is the element that compensates exactly for the difference found in the convexity

[3] In our development of the models we simplify the wide variety of symbols found in this strand of the literature.

inequality, the variance of the VIX futures prices is often referred to as the "concavity/convexity adjustment." However, this adjustment is not directly observable. In order to be estimated, the convexity adjustment requires the specification of both a model and a corresponding set of assumptions: therefore, this correcting factor is model dependent. The use of different VIX dynamic specifications can consequently lead to radically different VIX futures variance estimates for such a model. The next section shows how we examine this issue by using an approach similar to the concept of implied volatility for options written on traditional securities.

4. Methodology

One approach to the convexity adjustment for VIX futures is to embark on a quest for the "perfect model." However, this approach requires some arbitrariness in its model specification, complex calibrations of the model itself, and ambiguous exercises in forecasting for the future variance related to the convexity adjustment. Instead, we compute the *implied* convexity using *actual* VIX futures prices and the S&P 500 near- and far-term option-implied volatility values. We then analyze the properties of the convexity adjustment factor over time, as well as with respect to several market dynamics. This new approach of examining the convexity adjustment provides new information related to VIX futures pricing and potential arbitrage opportunities with VIX futures.

VIX futures contracts started trading in March of 2004 and therefore provide observable prices and volatilities since that date. However, the forward-starting variance swap rates needed for the cash comparison are not traded on exchanges. Alternatively, forward-starting variance swap rates can be determined from the value of a portfolio of S&P 500 options. Therefore, we take advantage of publicly reported volatility indices that reflect the prices of the S&P 500 options portfolios. Forward-starting variance swap rates can then be computed in a way that is similar to inferring the forward curve from a yield curve. As described below, we calculate these values by taking advantage of the additive nature of variances and of the aforementioned index volatility measures tracked and reported by the CBOE.

Here we detail the procedure used to extract the implied convexity value. First, on the date the VIX futures contract expires, the 30-day implied volatility forecast by the VIX futures is simply the implied volatility determined from the S&P 500 options series that expires 30 days after the VIX futures expiration. However, before the expiration of the VIX futures, the 30-day implied volatility underlying the futures contract is a "forward-starting implied volatility," i.e. a 30-day volatility starting at the futures expiration. Consequently, we must calculate the forward implied volatility (variance) as of any specific day/time before VIX futures expiration.

In order to introduce the concept of implied convexity we first examine the calculation of variance in general. 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 relation among the variances associated with the various corresponding periods needed to determine a total variance over time period 1 to 3 is as follows:

$$(t_3 - t_1)\sigma^2_{t_1 \to t_3} = (t_2 - t_1)\sigma^2_{t_1 \to t_2} + (t_3 - t_2)\sigma^2_{t_2 \to t_3} \tag{4}$$

where $\sigma^2_{t_i \to t_j}$ represents the annualized 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 (5) providing the total variance between the present time t and 30 days after the VIX futures expires:

$$(T + 30 - t)\sigma^2_{t \to T+30} = (T - t)\sigma^2_{t \to T} + 30\sigma^2_{T \to T+30} \tag{5}$$

[4] Note that the variances do not technically need to be annualized, provided that they 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 (4) by the proper scaling factor.

where $\sigma^2_{t\to T+30}$ represents the variance between the present and 30 days after expiration of the 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 isolate this rate appearing in equation (5) as follows:

$$\sigma^2_{T\to T+30} = \left(\frac{T+30-t}{30}\right)\sigma^2_{t\to T+30} - \left(\frac{T-t}{30}\right)\sigma^2_{t\to T} \tag{6}$$

The CBOE tracks two indices relating to the cash VIX that also employ a portfolio of S&P 500 options and are calculated using the methodology developed for the VIX, namely the VIN (the implied volatility for the near-term S&P 500 option expiration) and the VIF (the implied volatility for the far-term S&P 500 option expiration). These series are the individual components that determine the cash VIX. However, instead of representing a constant *n*-day-ahead volatility as the VIX does, the VIN and the VIF determine the volatility between now and the first-month option expiration (VIN) and the volatility between now 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 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 (5) and (6). Consequently, the forward-starting variance swap rate in that situation is expressed as:

$$E_t(VIX^2_T) = \sigma^2_{T\to T+30} = \left(\frac{T+30-t}{30}\right)VIF^2 - \left(\frac{T-t}{30}\right)VIN^2 \tag{7}$$

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 VIN and VIF is either 28 or 35 days, not 30 days.[5] For equation (7) 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 indices. 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 indices by computing:

$$Adjusted\_VIN = \left(\frac{30+n_2}{30+n_2-n_1} \right) VIN + \left(\frac{-n_1}{30+n_2-n_1} \right) VIF \tag{8}$$

$$Adjusted\_VIF = \left(\frac{n_2}{30+n_2-n_1} \right) VIN + \left(\frac{30-n_1}{30+n_2-n_1} \right) VIF \tag{9}$$

We can then apply equation (7) to the new horizon-adjusted volatility indices, yielding an equation for the forward-starting variance swap rate when time adjustments are needed to match the VIX futures expiration:

$$E_t(VIX_T^2) = \sigma^2_{T \to T+30} = \left(\frac{T+30-t}{30} \right) Adjusted\_VIF^2 - \left(\frac{T-t}{30} \right) Adjusted\_VIN^2 \tag{10}$$

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:

[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. On the other hand, 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.

$$IC_t = Var_t(F_T) = E_t(VIX_T^2) - F_t^2 \tag{11}$$

where IC_t is the implied convexity, $Var_t(F_T)$ is the expected variance of future VIX futures prices between time *t* and time *T*, $E_t(VIX_T^2)$ is the forward-starting variance swap rate computed in (10), and F_t^2 is the square of the current price of the VIX futures contract expiring at time *T*.

5. Data

We obtain daily data from TradeStation for August 25, 2008 through the end of 2011 for the VIN and VIF series described in the previous section.[6] When necessary, the series are adjusted according to equations (8) and (9), 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&P500 options. Thus, any unusual behavior in the options in the last week (as part of the VIN) does not influence the indices. 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 (10).

VIX futures data are obtained from the CQG Data Factory for all contracts expiring between August 2008 and December 2011. The VIX futures series employed here is composed of only the nearby contracts, which are rolled over to the next nearby contract on the Monday before the third Friday of every month. 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

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

daily closing VIX futures mid-price and the contemporaneous forward-starting variance swap rate using equation (11).

6. Implied convexity properties and results

6.1. Implied convexity: the "normal" range

Given the lack of prior studies on the behavior of implied convexity, a reasonable starting point is to examine the time series history of this measure. Figure 1 displays implied convexity levels in conjunction with VIX index levels from August 2008 through December 2011. This figure 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 faster than for the VIX. Furthermore, the implied convexity is fairly stable as the VIX reverses to lower levels.

The important characteristics observed in Figure 1 are logical for two reasons. First, implied convexity is the market's forecast of the future realized variance of VIX futures prices between now and the futures contract's expiration. Thus, the implied convexity represents the variance of futures prices, not futures 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, since a high spot VIX level produces high VIX futures prices, this will generate high implied convexity values, since 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 therefore lower implied convexity values. In conclusion, the size of the convexity adjustment is strongly positively correlated to the level of the VIX, as reflected by Figure 1.

The second reason why the results in Figure 1 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 towards a lower value and therefore the future realized futures variance associated with that lower VIX level should 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* should also decrease. Since the implied convexity is an estimate of the futures *price* variance, and since the variance of a price is a function of both the variance of its returns and of the level of that price,[7] the combination of these two effects will lower the convexity adjustment, thus accelerating the mean reversion of the implied convexity compared to that of the underlying VIX.

Figure 1 also shows that the implied convexity dips below zero in many instances (41.21% of the time), a counterintuitive phenomenon for a no-arbitrage situation. This argument is based on the fact that the variance of the VIX futures prices can never be negative. Consequently, by definition, implied convexity should always be positive, since it is a forecast of the *variance* of VIX futures prices from time t until contract expiration, According to equation (11), 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 (10).[8] The fact that negative implied convexity values are a common occurrence in reality shows that VIX futures often trade above their theoretical upper bounds, an observation consistent with possible mispricing of VIX futures prices.

The results from Figure 1 raise the issue of what typical implied convexity values should be,

[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.

at least in terms of how the market currently values VIX futures that include negative implied convexities. Table 1 displays the summary statistics for the implied convexities by year in Panel A and by implied convexity quintile in Panel B. Panel A shows that implied convexity exhibited its highest value of 0.0660 and its lowest value of -0.0727 during the last four months of 2008. Since 2008 the volatility in implied convexity has steadily declined. Overall, implied convexity exhibited positive skewness and a relatively large kurtosis. However, during the last quarter of 2008 implied convexity exhibited negative skewness. The results in Panel B also show that high levels of positive implied convexity (quintile 5) possess a very large standard deviation of values. Figure 2 additionally reports the implied convexity's frequency distribution over the sample period. In particular, Figure 2 shows 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 and a negatively skewed distribution of outliers.

Table 2 describes the historical level percentiles of implied convexity overall, as well as by years. Over the sample time period, the median daily closing implied convexity level is 0.0006. The 50th, 75th, and 95th interquartile ranges show that implied convexity oscillates within a narrow range the majority of the time. However, as market volatility periodically changes, so does implied convexity. The results in Table 2 clearly show that substantial year-to-year variations in implied convexity levels exist. Such results are to be expected, since the VIX itself fluctuates considerably during the sample period, causing an increase in the volatility of the VIX and the 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 remained slightly above zero. In fact, the 5% and 95% cutoff values for implied convexity in Table 2 have steadily trended towards zero during this period. Moreover, generally the entire implied convexity range declined with time since 2008, which is consistent with an improvement in pricing efficiency.[9] Overall, these results show that the distribution of convexity values has evolved over time.

6.2. Implied convexity in relation to stock market volatility

In this section we investigate the relation between implied convexity and market volatility. Theoretically, the implied convexity component of the VIX futures price is the current expectation of the future realized variance of the VIX futures prices from the present to futures expiration. Therefore, if market volatility is high, the VIX is high, VIX futures are high, and consequently the variance of the VIX futures also should be high. The opposite relation is true when market volatility is low. Thus, a direct link should exist between implied convexity and various market volatility measures.

Table 3 displays the regression results between implied convexity and the *prior* 30-day historical variance of the S&P 500 index.[10] 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 implied convexity levels possess a positive and significant relation in general, with the historical variance of the S&P 500 index explaining approximately 18 percent of the variation in implied convexity for the entire time period. Grouping the data into 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 8%, 24%, and 12% R-squares, respectively. As such, the results

[9] In particular, during 2009 the implied convexity levels were between -0.0094 and 0.0160 (a range of 0.0254) at the 90th percentile, compared to a range of 0.0123 in 2010 and 0.0114 in 2011.

[10] The regressions between implied convexity and different market measures use a constant 1-year horizon implied convexity. To obtain the 1-year constant maturity implied convexity, we divide the implied convexity backed out of equation (11) by the number of calendar days left to the VIX futures expiration and multiply it by 365.

show that the relation between historical market volatility and current implied convexity becomes evident for extreme values of implied convexity; otherwise, any relation is undetectable.

A similar relation is found between implied convexity and implied market volatility. The results in Table 4 show that implied market volatility (the VIX) has a positive and significant positive relation with the implied convexity for quintiles 4 and 5, explaining approximately 11% and 36% of the total variation in implied convexity. However, for quintile 1 the relation is negative and significant, explaining approximately 22% of the total variation. The relation between VIX and implied convexity can be visualized in Figure 3, where the positive relation is evident for the larger values of implied convexity (or for the VIX). In other words, large implied convexity values are almost always positive. Overall, the results show that large implied convexity levels contain information associated with historical and expected market volatility.

6.3. Implied convexity as an estimate of the 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. Figure 4 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 figure 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, Figure 4 shows that implied convexity is positively related to VIX futures time to expiration. This relation 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.[11] Table 5 shows the results of the multiple regression analysis of all three explanatory variables on the level of the implied convexity. The results show that the *level* of the implied convexity is related to the VIX futures' time to expiration in the entire sample. However, the *sign* of the implied convexity (i.e. whether negative or positive) is not related to either VIX futures time to expiration, VIX futures liquidity, or the dummy variable determining the presence of a jump in the VIX series (a one for a jump and zero otherwise).

Figure 4 also shows that implied convexity exhibits a pronounced jagged pattern as well as a dipping below zero with 13 days to expiration.[12] Alternatively, the future *realized* VIX futures variance exhibits a smoother pattern, and by definition remains above zero when the VIX futures contract is close to expiration. Additionally, the VIX futures variance stays relatively constant throughout the 38 calendar days before expiration. The largest difference between the implied convexity and the realized VIX futures variance (0.0042) occurs 13 days from VIX futures expiration, whereas the smallest difference (0.0001) occurs 9 days before expiration. Overall, the implied convexity forecasts the VIX futures variance relatively closely.

6.4. Violations of the VIX futures upper bound

Previously we found that implied convexity values are negative nearly 50% of the time, meaning that the VIX futures contract often trades above their 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

[11] Jumps are examined in detail in the following section.

[12] Thus, in Figure 4 the *average* values for the implied convexities for each day before expiration are positive (except for day 13), whereas Figures 2 and 3 show that almost 50% of the observations in the sample possess negative convexity values.

futures bid-ask spread (94.93% of all violations), except for a few cases when the violation is very small.

Figure 5 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 often it is much larger. Given the significant size of the adjustment needed to correct for the VIX futures upper bound violation, we explore its relation to implied convexity. Table 6 shows the results of the 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 (93.14%), meaning that implied convexity is strongly related to the upper bound adjustment. However, when the data is separated into quintiles based on the size of the implied convexity, the results show that only the lowest (quintile 1) and highest (quintile 5) implied convexity levels explain a large amount of the variation in the upper bound violations, whereas the middle quintiles explain much less. Statistically, the results show that implied convexity and the size of the VIX futures upper bound violation are strongly related. However, as implied convexity values tend towards zero, it becomes a less important explanatory variable for the size of the VIX futures upper bound violations.

7. 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 is based on several such assumptions. In conjunction with Jensen's inequality, the lower bound of the VIX futures is the forward-starting volatility swap rate, and the upper bound is the square root of the forward-starting variance swap rate. The convexity adjustment is subtracted from the price of the forward variance to determine the fair value of the VIX futures.

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 paper, is constrained by article length, and is left for future research if needed.

7.1. Jumps in the volatility process

A consideration with our implied convexity analysis stems from 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 plus the quadratic variation of the jump component of the process. In fact, Broadie and Jain (2008) note 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 nonparametric jump tests of Barndorff-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:

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

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

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

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

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. Since 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 1% significance level. We determine that only five likely instances of jumps occur. Since the sample period contains 847 trading days, this is the equivalent of less than one-and-a-half jumps per year. For comparison, Lee (2012) investigates 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. We therefore 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 bias due to jumps.

7.2. 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 discrete strikes for out-of-the-money options with non-zero bids until two adjacent non-zero bids occur. Jiang and Tian (2007) examine 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, since the simulation study employed by Jiang and Tian to achieve the results employs an unrealistic large range of strike prices relative to real-life strikes. Moreover, only the relative difference between VIN and VIF would affect the term structure volatility. Consequently, using the actual strike series employed by the markets would have a minimal bias on the outcomes found in this study.

7.3. 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 utilized to calculate the VIX, VIN, and VIF indices. Pavlova and Daigler (2008) do determine that large differences existed through mid-May 2007 for certain expirations between the cash VIX and the VIX futures settlement value. However, only six of 33 months had biases where the VIX futures was larger than the cash VIX at futures expiration, i.e. situations consistent with a negative implied convexity. Thus, any biases due to settlement in past expirations did not promote the negative implied convexities consistent with arbitrage opportunities. 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.

8. Conclusions

The VIX futures price theoretically includes a convexity adjustment component to compensate for the difference between the forward start swap rate and the squared VIX futures price, with that difference equal to the future variance of the VIX futures price. One approach to dealing with this issue is to model the future VIX futures variance based on some assumed stochastic behavior of the VIX. We approach the problem differently by solving for the convexity implied by determining 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 then proceed to study the properties of this implied convexity.

We examine the characteristics of the implied convexity adjustment using daily data. The results show that the implied convexity daily values are often negative and therefore outside the no-arbitrage range. Moreover, the implied convexity is related to the S&P 500 volatility. Finally, the implied convexity *underestimates* realized VIX futures variance approximately 71% of the time, which is the opposite of the relation between implied volatility and realized volatility.

Further research is needed to understand why implied convexity possesses the time series properties and relational characteristics discovered in the paper. In particular, these results suggest that arbitrage profits exist for a very liquid futures contract. Such further research should provide insights to VIX futures pricing in general, and therefore to a superior VIX futures pricing model.

References

Ait-Sahalia, Y. (2002). Telling from discrete data whether the underlying continuous-time model is a diffusion. Journal of Finance, 57, 2075-2112.

Becker, R., Clements, A., and White, S. (2007). Does implied volatility provide any information beyond that captured in model-based volatility forecasts? Journal of Banking and Finance, 31, 2535-2549.

Barndorff-Nielsen, O. E., and Shephard, N. (2006). Econometrics of testing for jumps in financial economics using bipower variation. Journal of Financial Econometrics, 4, 1-30.

Broadie, M., and Jain, A. (2008). The effect of jumps and discrete sampling on volatility and variance swaps. International Journal of Theoretical and Applied Finance, 11(8), 761-797.

Carr, P., and Wu, L. (2003). What type of process underlies options? A simple robust test. Journal of Finance, 58, 2581-2610.

Carr, P., and Wu, L. (2006). A tale of two indices. Journal of Derivatives, 13, 13-29.

Exchange, C. B. O. (2003). The VIX white paper. *URL: http://www. cboe. com/micro/vix/vixwhite. pdf*.

Dupoyet, B., Daigler, R., and Chen, Z. (2011). A simplified pricing model for volatility futures. The Journal of Futures Markets, 31, 307-339.

Grunbichler, A., and Longstaff, F. A. (1996). Valuing futures and options on volatility. Journal of Banking and Finance, 20, 985-1001.

Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. The Review of Financial Studies, 6, 327–343.

Huskaj, B., and Nossman, M. (2013). A term structure model for VIX futures. The Journal of Futures Markets, 33, 421-442.

Jiang, G., and Tian, Y. (2007). Extracting model-free volatility from option prices: An examination of the VIX index. Journal of Derivatives, 14, 1-26.

Jiang, G. J., and Oomen, R. C. (2008). Testing for jumps when asset prices are observed with noise - a swap variance approach. Journal of Econometrics, 144, 352-370.

Lee, S. (2012). Jumps and information flow in financial markets. Review of Financial Studies, 25, 439-479.

Lee, S., and Mykland, P. (2008). Jumps in financial markets: a new nonparametric test and jump dynamics. Review of Financial Studies, 21, 2535-2563.

Lin, Y. N. (2007). Pricing VIX futures: Evidence from integrated physical and risk-neutral probability measures. The Journal of Futures Markets, 27, 1175-1217.

Lu, . J., and Zhu, Y. Z. (2010). Volatility components: The term structure of VIX futures. The Journal of Futures Markets, 30, 230-256.

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

Pavlova, I., and Daigler, R. T. (2008). The non-convergence of the VIX futures at expiration. Review of Futures Markets, 17, 201-223.

Poon, S. H., and Clive W. J. G. (2003). Forecasting volatility in financial markets: A review. Journal of Economic Literature, 41, 478-539.

Wang, Z., and Daigler, R. (2013). The option SKEW index and the volatility of volatility. Financial Management Meetings, Chicago.

Zhang, J. E., and Huang, Y. (2010). The CBOE S&P 500 three-month variance futures. Journal of Futures Markets, 30, 48–70.

Zhang, J. E., Shu, J. and Brenner M. (2010). The new market for volatility trading. The Journal of Futures Markets, 30, 809-833.

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

Zhu, S. P., and Lian, G. H. (2012). An analytical formula for VIX futures and its applications. The Journal of Futures Markets, 32, 166-190.

Figure 1

Daily closing levels of the VIX and implied convexity (IC) from August 2008 to December 2011.

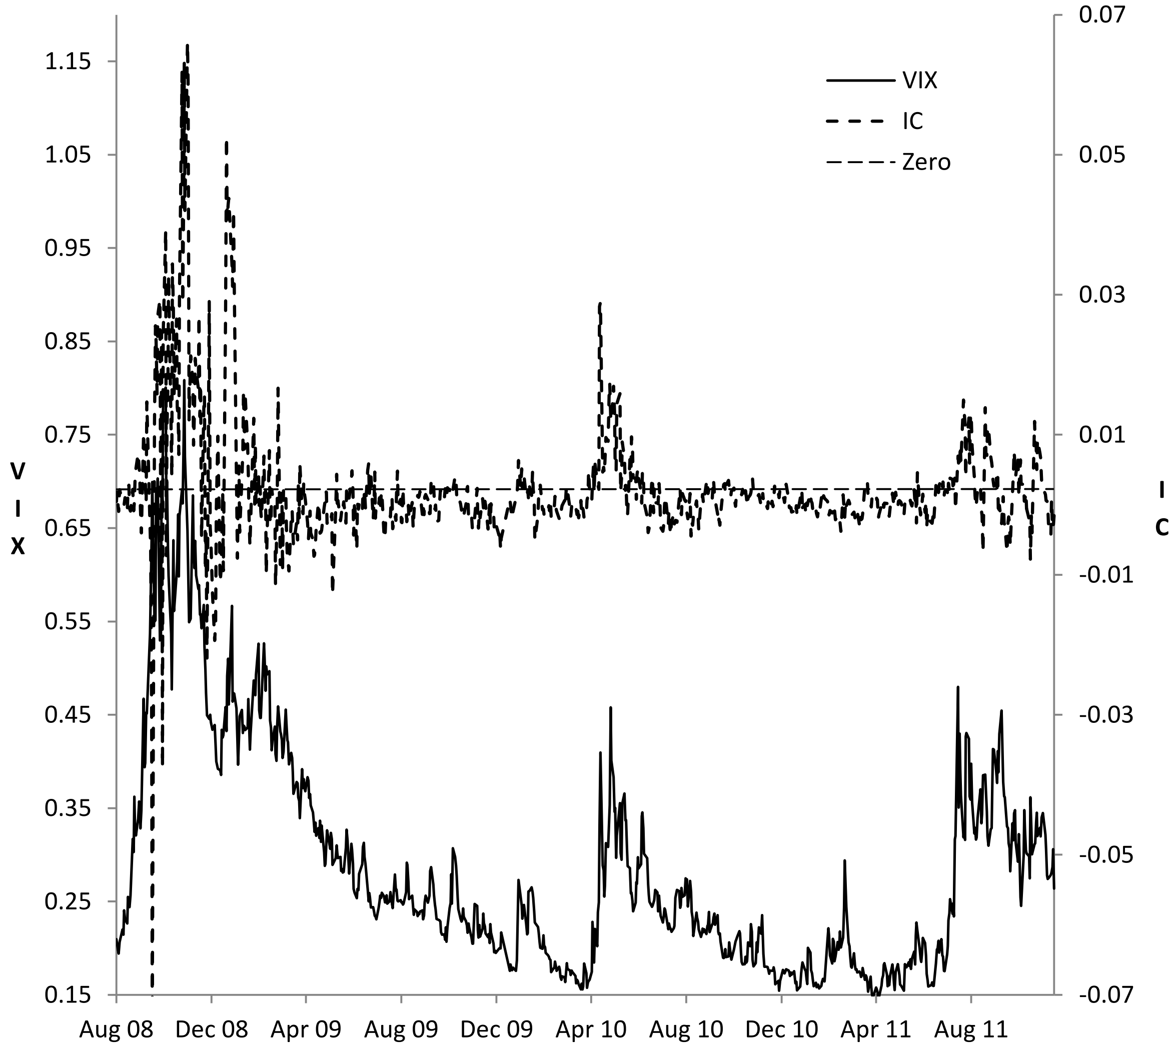

Note: This figure shows the historical level of the VIX (left vertical axis) and of the implied convexity, IC, (right vertical axis) from August 2008 to December 2011. The "Zero" gradient line corresponds to the lower bound of implied convexity.

Figure 2
Frequency distribution of implied convexity

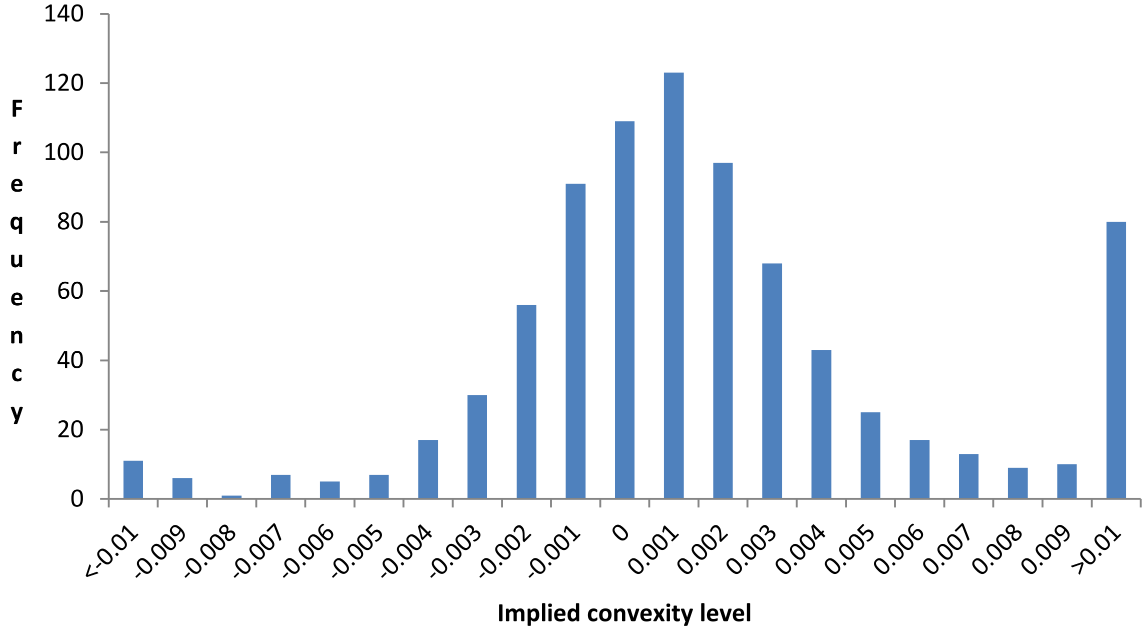

Note: This figure shows the distribution of the implied convexity daily closing values from August 2008 to December 2011.

Figure 3
Relation between VIX and implied convexity

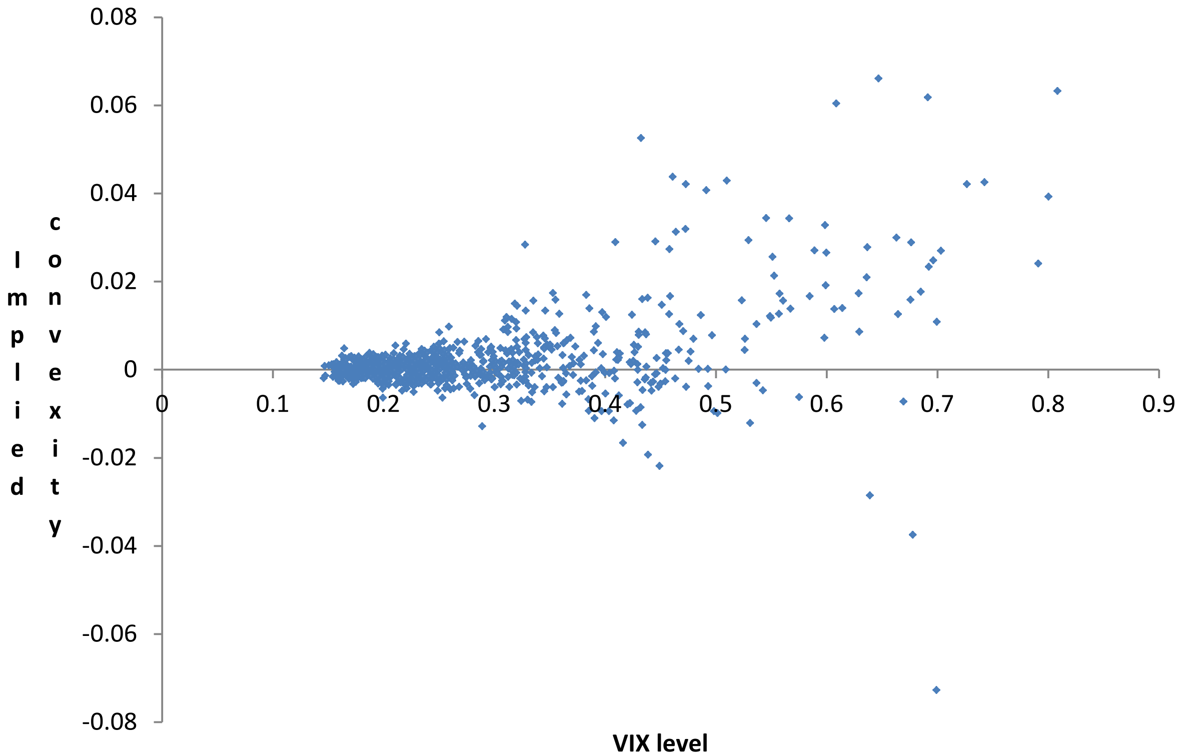

Note: This figure plots the implied convexity's daily closing values (vertical axis) against the VIX daily closing values (horizontal axis). The sample period is August 2008 to December 2011.

Figure 4
Average realized VIX futures variance (Var VIX) and implied convexity (IC) values by days left until VIX futures expiration.

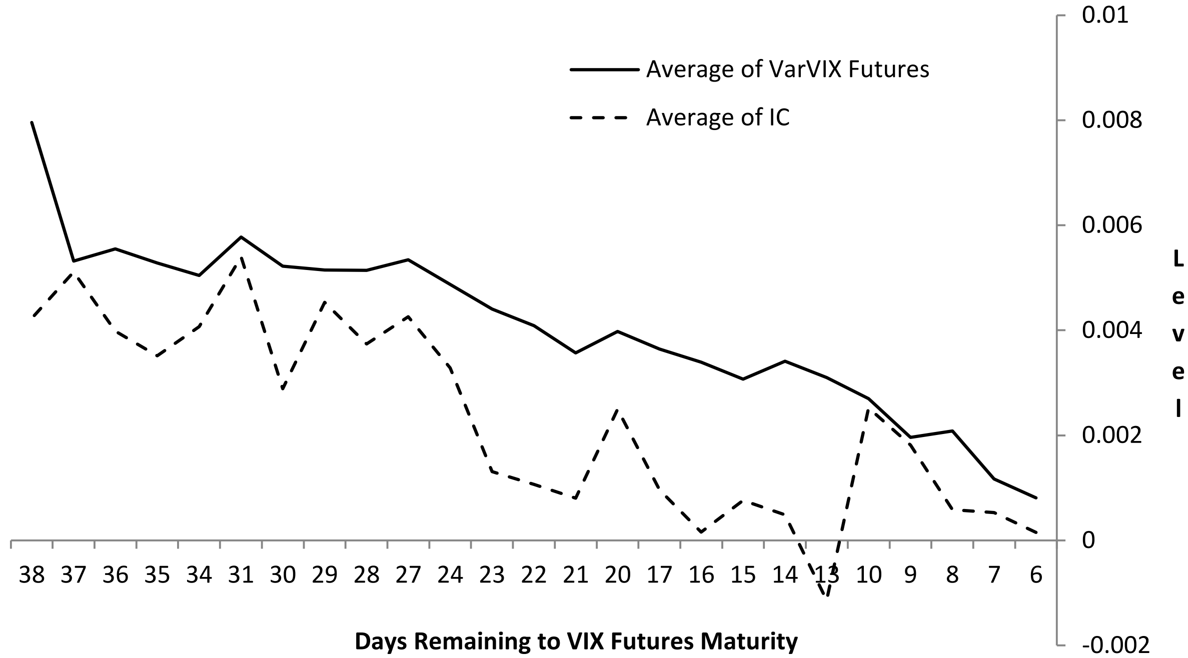

Note: This figure shows the average implied convexity daily closing levels (IC) and the average future realized VIX futures variance by days remaining to VIX futures maturity. At any point in time, the VIX futures contract is the nearby contract.

Figure 5

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)

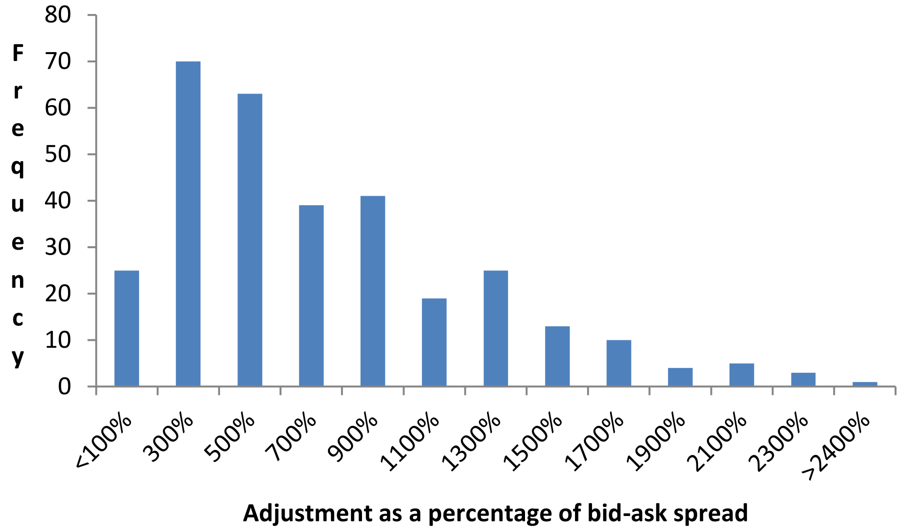

Note: This figure shows the frequency distribution for the VIX futures downward adjustment needed in order to prevent an upper bound violation of the VIX futures, as defined by Carr and Wu (2006). The adjustment is expressed as a percentage of the VIX futures bid-ask spread. Only violations that fall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

Table 1
Summary statistics for daily closing levels of implied convexity

	No. obs.	Mean	Median	Min.	Max.	Std. Dev.	Skewness	Kurtosis
Panel A: by year								
All	825	*0.0022*	0.0006	-0.0727	0.0660	0.0092	2.11	18.97
2008	77	*0.0118*	0.0103	-0.0727	0.0660	0.0208	-0.28	3.79
2009	250	0.0011	0.0000	-0.0219	0.0526	0.0090	2.88	11.84
2010	246	*0.0014*	0.0008	-0.0064	0.0289	0.0044	2.83	12.69
2011	252	*0.0012*	0.0005	-0.0078	0.0150	0.0036	1.28	2.46
Panel B: by Implied Convexity quintile								
1	165	*-0.0049*	-0.0031	-0.0727	-0.0017	0.0070	-6.76	58.04
2	165	*-0.0009*	-0.0008	-0.0017	-0.0001	0.0005	-0.06	-1.29
3	165	*0.0006*	0.0006	-0.0001	0.0013	0.0004	0.06	-1.03
4	165	*0.0023*	0.0022	0.0013	0.0037	0.0007	0.51	-0.86
5	165	*0.0141*	0.0085	0.0038	0.0660	0.0129	1.98	3.96

Note: 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 are in italics.

Table 2
Normal ranges for daily closing levels of implied convexity over the sample period

Year	No. obs.	5%	10%	25%	50%	75%	90%	95%
All	825	-0.0048	-0.0031	-0.0013	0.0006	0.0028	0.0085	0.0166
2008	77	-0.0122	-0.0040	0.0003	0.0103	0.0233	0.0344	0.0604
2009	250	-0.0094	-0.0050	-0.0023	0.0000	0.0020	0.0057	0.0160
2010	246	-0.0032	-0.0025	-0.0009	0.0008	0.0024	0.0051	0.0091
2011	252	-0.0029	-0.0021	-0.0010	0.0005	0.0026	0.0061	0.0085

Note: This table shows the distribution of daily closing levels of implied convexity by percentiles for the entire data set and for each year of the data.

Table 3
Regression results for implied convexity regressed on the prior month for the S&P 500 return variance

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5	All
Constant	-0.0700***	-0.0170***	0.0099***	0.0330***	0.1151***	-0.0197***
	(0.0168)	(0.0010)	(0.0006)	(0.0012)	(0.0244)	(0.0071)
SP500 variance	-0.4982***	-0.0012	-0.0124	0.0612***	0.6363***	0.5777***
	(0.1038)	(0.0169)	(0.0102)	(0.0159)	(0.0776)	(0.0428)
Regression statistics						
R-squared	0.1238	0.0000	0.0089	0.0829	0.2920	0.1815
No. obs.	165	165	165	165	165	825

Note: This table provides the regression where the daily closing level of implied convexity is the dependent variable and the S&P 500 return variance over the prior 30 days is the independent variable. The data is annualized for consistency. Quintile 1 represents the smallest implied convexity values. Standard errors are reported in parentheses below the regression coefficients.*, **, *** shows significance at the 10%, 5%, and 1% levels, respectively.

Table 4
Regression results for implied convexity regressed on the VIX daily closing levels

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5	All
Constant	0.0977***	-0.0167***	-0.0124***	0.0228***	-0.1865***	-0.1383***
	(0.0343)	(0.0027)	(0.0015)	(0.0032)	(0.0489)	(0.0152)
VIX	-0.7204***	-0.0013	-0.0129**	0.0531***	1.0572***	0.5975***
	(0.1054)	(0.0112)	(0.0064)	(0.0121)	(0.1095)	(0.0488)
Regression statistics						
R-squared	0.2227	0.0001	0.0241	0.1053	0.3638	0.1541
No. obs.	165	165	165	165	165	825

Note: This table gives the results of the regression where the daily closing level of implied convexity is the dependent variable and the daily closing level of the VIX is the independent variable. Quintile 1 represents the smallest implied convexity values. Standard errors are reported in parentheses below the regression coefficient.
*, **, *** shows significance at the 10%, 5%, and 1% levels, respectively.

Table 5
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	-0.0011	-0.0028***	0.0048
	(0.0008)	(0.0006)	(0.0013)
VIX Jumps	-0.0003	0.0012	-0.0023
	(0.0029)	(0.0031)	(0.0037)
VX Volume^	-2.2796	1.6261	-6.9391
	(2.5814)	(1.7526)	(5.4117)
VX Maturity^	0.0111***	-0.0004	0.0038
	(0.0023)	(0.0022)	(0.0034)
Regression Statistics			
R-squared	0.0277	0.0031	0.0068
No. obs.	830	345	485

Note: This table gives 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+ denotes negative and positive values of implied convexity, respectively. Standard errors are reported in parentheses below the regression coefficient.
*, **, *** shows significance at the 10%, 5%, and 1% levels, respectively.
^Regression coefficients and standard errors are multiplied by 100,000 for readability.

Table 6

Regression results for adjustment needed to correct VIX futures upper bound violation regressed on the implied convexity

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5	All
Constant	0.0031***	0.0006	0.0011	0.0001	-0.0000	0.0017***
	(0.0003)	(0.0011)	(0.0010)	(0.0003)	(0.0000)	(0.0000)
Implied Convexity	-0.9163***	-1.6900***	-1.5466***	-2.0846***	-2.3394***	-1.0019***
	(0.0232)	(0.4058)	(0.5657)	(0.2812)	(0.1020)	(0.0149)
Regression statistics						
R-squared	0.9599	0.2106	0.1031	0.4582	0.8900	0.9314
No. obs.	67	67	67	67	67	335

Note: This table gives 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 is the independent variable. Standard errors 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, the number of observations here is lower than the total number of observations in the sample period. Quintile 1 represents the smallest values.
*, **, *** shows significance at the 10%, 5%, and 1% levels, respectively.