LIQUIDITY CONSIDERATIONS IN ESTIMATING IMPLIED VOLATILITY

ROHINI GROVER
SUSAN THOMAS*

Option markets have significant variation in liquidity across different option series. Illiquidity reduces the informativeness of the price. Price information for illiquid options is more noisy, and thus the implied volatilities (IVs) based on these prices are more noisy. In this study, we propose weighting schemes to estimate IV, which reduce the importance attached to illiquid options. The two indexes using liquidity weights are SVIX, which is a spread-adjusted volatility index, and TVVIX, which is a traded volume weighted VIX. We find SVIX outperforms TVVIX, the conventional schemes such as the traditional VXO, or vega weights, and volatility elasticity weights. © 2012 Wiley Periodicals, Inc. Jrl Fut Mark 32:714–741, 2012

1. INTRODUCTION

When options markets became established and liquid, market prices of options were used to directly calculate the market forecast of volatility called the

We are grateful to Abhisar Srivastava, Tirthankar Patnaik, and the National Stock Exchange of India Ltd. for the data used in this study. We thank Nidhi Aggarwal, Ajay Shah, Rajat Tayal, and the participants of the IGIDR Finance Research Seminars for useful discussions. We also thank our discussant Prof. Moon, Seong Ju, and participants of the 7th Annual APAD Conference for useful suggestions. The views expressed in this study belong to the authors and not their employer.

*Correspondence author, Indira Gandhi Institute of Development Research, Goregaon East, Mumbai 400065, India. Tel: (+91) (022)284-165-50, e-mail: susant@igidr.ac.in

Received October 2011; Accepted November 2011

■ *Rohini Grover is a Ph.D. Scholar at Indira Gandhi Institute of Development Research, Mumbai, India.*

■ *Susan Thomas is an Assistant Professor at Indira Gandhi Institute of Development Research, Mumbai, India.*

The Journal of Futures Markets, Vol. 32, No. 8, 714–741 (2012)

Published online February 10, 2012 in Wiley Online Library (wileyonlinelibrary.com).
DOI: 10.1002/fut.21543

"implied volatility" (IV). IV is a direct measure of the *forecast* of volatility made by economic agents. An extensive literature has documented the high quality of volatility forecasting that is embedded in IV.

Different options on the same underlying yield different values for the IV. Analytical methods are thus required to reduce multiple values for IV from different traded options on the same underlying into an efficient point estimate of IV for that underlying.

One endeavor to isolate a volatility forecast from multiple values of IV was based on developing models of option pricing that incorporated the factors that caused IV to change deterministically, such as moneyness of the option, volatility dynamics, or the liquidity of the underlying.

The other strategy adopted was to develop an index of IV. This was discussed extensively in the literature, which led up to the introduction of the Chicago Board Options Exchange (CBOE) VIX in 1993. This implied volatility index was calculated for the S&P 100 index options using the methodology proposed by Whaley (1993) and was disseminated by the CBOE in real time. The index was calculated using only at-the-money (ATM) options with a defined weighting scheme over the IV values calculated. In 2003, the CBOE shifted the VIX calculation methodology to one that used option prices over a wide range of strike values. In the following years, similar computation of an implied volatility index has commenced on numerous options exchanges worldwide. Trading in derivatives on VIX has also commenced.

The CBOE VIX methodology is predicated on all option prices being measured sharply. However, in the real world, there is substantial cross-sectional variation in the liquidity of option series. As an example, in tranquil times (September 2007), the bid–offer spread of options on the S&P 500 index at the CBOE ranged from near 0% to 200%. In turbulent times (September 2008), many more options were afflicted with illiquidity.

At present, a variety of heuristics are being utilized by exchanges worldwide in addressing this problem. In this study, we try to frontally address the problem of illiquid options markets by constructing a weighting scheme for the construction of a VIX that directly incorporates the liquidity of the option. The empirical work of this study is based on one of the most active option markets in the world: options on the NSE-50 (Nifty) index, traded at the National Stock Exchange. We use the bid–offer spread in a weighting scheme that adjusts for illiquidity when calculating the VIX. We call this the *spread-adjusted* VIX (SVIX). We also use traded volumes of options as another liquidity proxy and compute *volume-adjusted* VIX (TVVIX).

We compare the performance of SVIX and TVVIX against three alternative weighting schemes: the 1993 CBOE index (called VXO), the vega-weighted index (VVIX), and a volatility elasticity weighted index (EVIX). The performance is measured as the forecasting success of each VIX candidate against the realized volatility (RV) of the market index. The testing procedure employed is the Model Confidence Set (MCS) test (Hansen, Lunde, & Nason, 2003). We find that the new SVIX is a better predictor of future RV. We also run univariate regressions of RV on each VIX and find that although all VIXs contain information about future volatility, they are biased forecasts.

Option IV is an important component of the information set of the financial system. The world over, options markets are being used to create implied volatility indexes using ideas similar to that of the CBOE VIX. Because all options markets have substantial cross-sectional variation in option liquidity, the ideas of this study may potentially yield improved measurement of VIXs.

The study is presented as follows: In Section 2, we present the issues surrounding the creation of a VIX and also present the evaluation framework used to compare the performance of alternative VIXs. In Section 3, we review alternative schemes to construct VIXs. In Section 4, we describe the data used for the analysis. In Section 5, we discuss the liquidity-adjusted weighting scheme, after which we present our analysis in Section 6. In Section 7, we conclude.

2. ISSUES IN CONSTRUCTING IV INDEXES

Gastineau (1977) proposed the use of an index to resolve the problem of multiple values of IVs from different options on the same underlying. An implied volatility index, calculated as a weighted average of the IVs from different option prices, would be the summary measure of underlying future volatility.

The first weighting schemes were suggested by Trippi (1977) and Schmalensee and Trippi (1978), which placed equal weights on all the IVs used in calculating the index. However, because the literature showed that the Black–Scholes model priced some options more accurately than others, schemes where the weights varied according to different factors were proposed. In following years, several researchers made significant progress in developing these concepts further (Galai, 1989; Cox & Rubinstein, 1985; Brenner & Galai, 1993; Whaley, 1993).

The maturation of knowledge in this field was signaled with the launch of an information product in 1993: an implied volatility index based on trading in options on the S&P 100 index. This was called the CBOE VIX. A research literature rapidly demonstrated that VIX was useful in volatility forecasting, over and beyond the state-of-the-art volatility models, because option prices harnessed the superior information set of traders.[1]

[1]Christensen and Prabhala (1998); Christensen, Hansen, and Prabhala (2001) identified and corrected some of the data and methodological problems present in the early studies on this question. They conclude that

Given the importance of VIXs such as VIX in the global financial system, it is useful to explore the methodological issues in the construction of these indexes. The process of creating an optimal methodology for a VIX involves the following two parts:

1. Identifying alternative weighting schemes based on available data about factors that directly influence the shape of the IV smile.
2. Choosing an optimal weighting scheme.

2.1. Factors Influencing IV Values

If the Black–Scholes model held exactly, all options should have the same IV. However, an extensive literature has demonstrated that IV varies with moneyness, maturity, vega, and liquidity. We discuss each of these in turn.

Moneyness/strike

The first documented variation in IV was as a function of strikes or the moneyness. IV was consistently lower for lower values of the moneyness of the option. This variation came to be known as the *volatility smile*. Rubinstein (1994), Jackwerth and Rubinstein (1996), Dumas, Fleming, and Whaley (1998) showed that the pattern of the IV of the S&P 500 index options changed from a smile to a sneer after the 1987 crash.

Maturity

Prices of near-month options show lower IV than far months. Heynen, Kemna, and Vorst (1994), Xu and Taylor (1994), and Campa and Chang (1995) show that IVs are a function of time to expiration and thus exhibit a term structure.

Vega

The derivative of the Black–Scholes price with respect to volatility is called vega. The vega can be shown to be consistently different for different values of

IV is a more efficient forecast for future volatility than volatility calculated from historical returns. Latane and Rendleman (1976), Chiras and Manaster (1978), and Beckers (1981) find that IV performs better in capturing future volatility than standard deviations obtained from historical returns. Blair, Poon, and Taylor (2001) find that volatility forecasts provided by the early CBOE VIX are unbiased, and they outperform forecasts augmented with GARCH effects and high-frequency observations. Similar results were reported early on by Jorion (1995) for foreign exchange options.

Corrado and Miller (2005) examine the forecasting quality of three implied VIXs based on S&P 100, S&P 500, and NASDAQ 100 (National Association of Securities Dealers Automated Quotations). They find that the forecasting quality of the VIX based on the S&P 100 and S&P 500 has improved since 1995, and that those based on the NASDAQ 100 provides better forecasts of future volatility.

the strike, as well as the maturity of the contract.[2] Thus, the vega of an option naturally lent itself as an input to differentiating the IV of different options when calculating an implied volatility index (Latane & Rendleman, 1976). Chiras and Manaster (1978) suggested weighting by volatility elasticity instead of vega.

Among other influential papers, Beckers (1981) and Whaley (1982) suggested minimizing $\sum_i w_i[C_i - BS_i(\hat{\sigma})]^2$, where C_i refers to market price and BS_i refers to Black–Scholes price of option i and w_i could either be vega or equal weights.[3]

Liquidity

A more recent literature has explored the impact of option liquidity on estimated IV. Brenner, Eldor, and Hauser (2001) show that there is a significant illiquidity premium between two sets of currency options, when one set is traded and the other is not. Bollen and Whaley (2004) documented an empirical link between the shape of the IV smile and the depth of the market on the buy and the sell side of options with different moneyness. They show that *net buying pressure* affects the shape of the IV smile in both the index as well as the single stock options markets. Further, they show that the shape of the IV smile is driven by different market forces for index options compared to single stock options.

Models of asymmetric information have been used to provide theoretical underpinnings for the link between liquidity and option price. Nandi (2000) sets up a model of asymmetric information linking the level and the shape of IV function to net order flow of options. The model shows that an increase in net options order flow increases the mispricing by the Black–Scholes model. Garleanu et al. (2009) formalize the findings in Bollen and Whaley (2004) by incorporating end-user demand in a model for options prices. Here they exploit the feature that end-users tend to hold long index options and short equity options to explain the relative expensiveness of index options. Another model to explicitly incorporate liquidity in the price of stock options was Cetin et al. (2006), who show market liquidity premium[4] as a significant part of the option price.

The empirical evidence has also linked IV to option liquidity. Etling and Miller (2000) explore the relationship between bid–ask spread as a liquidity

[2]Vega is higher for options that are further away from the money because they have a lower extrinsic value and are less likely to change with changes in IV. It is also higher for options with longer expiration in order to compensate for additional risk taken by the seller.
[3]This method thus allows the call prices to provide an implicit weighting scheme that yields an estimate of standard deviation that has least prediction error.
[4]The paper models the liquidity using a generic supply function where option price monotonically increases with size of order.

proxy with moneyness of options and find that ATM options have the highest liquidity. Chou et al. (2009) explore how the IV function varies as a function of liquidity in both the spot and options market. They find that *order-based* measures of liquidity (such as the bid–ask spread) better explain the variation in IV than *trade-based* measures (such as traded volume). They also find that both spot and options markets liquidity matter for the variation in IV.

This evidence, about the various factors that influence IV, has led to many alternative approaches to constructing an implied volatility index. The different weighting schemes are further discussed in Section 3. What is the *efficient* weighting scheme rests upon the performance of the forecast from each scheme against some benchmark volatility measure. The framework to carry out such a performance evaluation of different weighting schemes is now examined in the next section.

2.2. Performance Evaluation

One of the reasons that there is no consensus on one best weighting scheme for a VIX is the lack of an observable volatility. The time-series econometrics literature has extensive work on a framework to evaluate the performance of a volatility forecast even though volatility is not observed. For example, these ideas have been used in testing the forecasts of volatility models such as Generalized Autoregressive Conditional Heteroskedasticity (GARCH), Exponentially Weighted Moving Average (EWMA), etc. This framework has two broad approaches: one that delivers a relative measure of performance among a set of candidate models, and the other that delivers a measure of performance of each of the candidate models against a single benchmark.

These questions were revisited when intraday data revealed a superior volatility proxy: RV. Once RV was observed, it became possible to measure how well IV forecasts the RV of the underlying asset over the life of an option. Most studies use a predictive regression of the IV estimate on future volatility where the goodness of prediction is measured through the coefficients of predictive regressions. The early studies by Day and Lewis (1988), Lamoureux and Lastrapes (1993), and Canina and Figlewski (1993) showed that IV is *not* a good predictor for future return volatilities.

The framework of *encompassing regressions* was then used to assess the predictability of IV estimates against other forecast variables. This framework addresses the relative importance of competing volatility forecasts and whether one volatility forecast subsumes all information contained in other volatility forecast(s). Within this approach, Poteshman (2000); Jiang and Tian (2005b); Corrado and Miller (2005) have found that IV estimates are biased, but efficient and informative relative to forecasts from other volatility estimates.

A recent study by Becker, Clements, and White (2007) used an approach that differs from the traditional forecast encompassing approach used in earlier studies and finds that the S&P 500 implied volatility index does not contain any such incremental information relevant for forecasting volatility. Becker, Clements, and White (2008) compare the index against a combination of forecasts of S&P 500 volatility by using the MCS methodology and finds that a combination of forecasts outperforms individual model-based forecasts and IV.

In this study, we use the following two steps to compare the performance of our VIXs:

1. Forecasting regressions following Christensen and Prabhala (1998) to test the information content of the volatility measures. Instrumental variable regressions are also run to correct for potential errors-in-variable problems in IV estimates as discussed by previous studies.[5]
2. The MCS methodology of Hansen et al. (2003). This addresses the problem of choosing the *best* forecasting model. It contains the best model with a given level of confidence. It may contain a number of models, which indicates they are of equal predictive ability. It has several advantages over other methods such as superior predictive ability (SPA) test and the reality check (RC) test.[6] The construction of the MCS test is an iterative procedure in that it requires a sequence of tests for equal predictive ability. The set of candidate models is trimmed by deleting models that are found to be inferior. The final surviving set of models in the MCS contains the optimal model with a given level of confidence and these models are not significantly different in terms of their forecast performance.[7]

The critical question that remains in this is still the choice of the benchmark measure for volatility, which we discuss in Section 4.2.

3. CHOICES OF IV INDEXES

In this section, we describe different methods we use in order to calculate implied volatility indexes. We start with a description of the two most widely computed VIXs by several exchanges across the world, namely, VXO and VIX.

3.1. VXO

This VIX is calculated using prices of options on the S&P 100 index. The IVs are calculated using the Black–Scholes model, and the VXO is an average of

[5]Christensen and Prabhala (1998), Jiang and Tian (2005b), Corrado and Miller (2005).
[6]See Hansen et al. (2003).
[7]Hansen et al. (2003).

the IVs on eight near-the-money options, including options at the two nearest maturities.[8]

In 2003, VXO was criticized for using an option pricing model and being biased due to the trading day conversion. In addition, there were two structural changes[9] in the U.S. economy that reduced the usefulness of VXO as a measure of future volatility. These were as follows:

1. S&P 500 options became the most actively traded index options.
2. Earlier index calls and puts were equally important in investor-trading strategies but in later years the market became dominated by portfolio insurers who bought out-of-the-money and ATM index puts for insurance purposes.

Such criticisms of VXO along with changes in the structure of the U.S. options market led to a new approach to calculating the VIX based on the prices of options trading on the S&P 500.

3.2. VIX

In contrast to VXO, VIX has been derived from the concept of fair value of a volatility swap (Demeterfi et al., 1999). Here, even though the variance is derived from market observable option prices and interest rates, the theoretical underpinning is rooted in the broader context of model-free implied variance of Dupire (1993) and Neuberger (1994). This concept was further developed by Carr and Madan (1998), Demeterfi et al. (1999), and Britten-Jones and Neuberger (2000). Jiang and Tian (2005a) establish that the variance measure under this framework is theoretically equivalent to the model-free implied variance formulated by Britten-Jones and Neuberger (2000).

The CBOE calculates and publishes a real-time value of VIX,[10] which has been accepted as the market measure of volatility. In this study, we do not directly analyze the VIX methodology. However, to the extent that the main argument of this study is appropriate—that price information for illiquid option series is less informative—it should impact upon the VIX methodology also.

3.3. Volatility-Linked Weights

The early literature (Latane & Rendleman, 1976; Chiras & Manaster, 1978) suggests two different weighting schemes based on vega and volatility elasticity weighting scheme to calculate the market implied volatility index.

[8]See Whaley (1993) for details on construction of VXO.
[9]See Whaley (2009).
[10]See www.cboe.com/micro/vix/vixwhite.pdf for details on construction of VIX.

1. Vega weights are calculated as

$$\sigma_{tj} = \frac{\sum_i w_{it,j}\sigma_{it}}{\sum_i w_{it,j}},$$

where $w_{it,j}$ is the Black–Scholes vega for the option contract at time t; $j = 1,2$ denotes the two nearest maturities.

2. Volatility elasticity weights are calculated as

$$\sigma_{tj} = \frac{\sum_i w_{it,j}\frac{\sigma_{it,j}}{C_{it,j}}\sigma_{it,j}}{\sum_i w_{it,j}\frac{\sigma_{it,j}}{C_{it,j}}},$$

where $w_{it,j}$ is the Black–Scholes vega and $C_{it,j}$ is the price for the ith option contract at time t; $j = 1,2$ denotes the two nearest maturities.

The scheme that uses volatility elasticities puts more weight on out-of-the-money options (with low prices $C_{it,j}$) than the vega weights model.

3.4. Adjustment for Rollover

The IV estimates obtained for the two nearest maturity are linearly interpolated to obtain a 30-day estimate. Rollover to the next expiration occurs eight calendar days prior to the expiry of the nearby option. The interpolation scheme used is

$$VIX = 100 \times \left[\sigma_{t1}\left(\frac{N_{c2}-30}{N_{c2}-N_{c1}}\right) + \left(\frac{30-N_{c1}}{N_{c2}-N_{c1}}\right)\right],$$

where σ_{ti} are IVs and N_{ci} is the number of calendar days to expiration. Here, $i = 1,2$ for the near and next month, respectively.

4. MEASUREMENT

Data on the NSE-50 (Nifty) index options at the National Stock Exchange of India Ltd. (NSE) are used. The NSE is an extremely active exchange and is a high-quality source of data on exchange-traded derivatives. NSE is the fifth largest derivative exchange in the world in terms of number of contracts traded (Table I). It is also the third largest exchange in terms of number of contracts traded in equity index (Table II).

TABLE I
Global Exchanges: Number of Contracts Traded

Rank	*Exchanges*	*January to June 2009*	*January to June 2010*	*Change (%)*
1	Korea Exchange	1,464,666,838	1,781,536,153	21.6
2	CME group	1,283,607,627	1,571,345,534	22.4
3	EUREX	1,405,987,678	1,485,540,933	5.7
4	NYSE Euronext	847,659,175	1,210,532,100	42.8
5	National Stock Exchange of India	397,729,690	783,897,711	97.1

Note. CME, Chicago Mercantile Exchange; EUREX, European Exchange; NYSE, New York Stock Exchange.
Source: FIA, http://www.futuresindustry.org/volume-.asp.

TABLE II
Ranked by the Number of Contracts Traded in Equity Index

Rank	*Exchanges*	*January to June 2009*	*January to June 2010*	*Change (%)*
1	KOSPI 200 options, KRX	1,375,065,894	1,671,466,852	21.6
2	Emini S&P 500 futures, CME	308,764,146	299,603,623	–3.0
3	S&P CNX Nifty options, NSE India	146,706,110	221,430,570	50.9
4	SPDR S&P 500 ETF options, CME	181,699,626	219,409,316	20.8
5	DJ Euro Stoxx 50 futures, EUREX	178,923,108	205,280,712	14.7

Note. CME, Chicago Mercantile Exchange; EUREX, European Exchange; KOSPI 200, Korea Composite Stock Price Index 200.
Source: FIA, http://www.futuresindustry.org/volume-.asp.

In Table III, we show the average number of records of intraday data for the Nifty index option contracts from March 2009 to April 2010. The large number of records suggests a highly active market.

The numerical values shown in the three tables (Tables I–III) show very high growth rates at NSE. Thus, although the NSE is an important exchange on the global scale, its importance is likely to go up in the future if these growth rates continue.

4.1. Price Measurement

Even though we have both traded prices as well as mid-quote prices available at high frequency, we choose the mid-quote of the bid–ask orders of the options from the options limit order book as the benchmark input price in the Black–Scholes model to compute IV.

This is because of the relative illiquidity of the options market by way of trade updates compared to order updates. Very frequently, the next-month options market suffers from illiquidity in terms of trades, and there is no traded price that is observable. However, the order book has more liquidity in terms

TABLE III
Records of Intraday Data per Month, with the Nifty Index Options Market

Month	*Average no. of Records*
March 2009	2,891,142
April 2009	3,596,361
May 2009	2,732,869
June 2009	3,409,783
July 2009	4,563,929
August 2009	4,724,429
September 2009	4,855,215
October 2009	4,790,568
November 2009	6,014,413
December 2009	7,768,108
January 2010	7,163,127
February 2010	8,429,391
March 2010	7,527,610
April 2010	8,767,319

Source: NSE.

of order updates. Thus, there is information in the order book data that are not reflected in the traded prices. This makes it meaningful to use the mid-quote prices rather than traded prices to calculate IV for the VIXs because it reduces the effect of missing data.

The use of mid-quote prices has an automatic liquidity/illiquidity impact in the VIX calculation—where the options prices are moving due to changes in the limit order book rather than a realized price from a transaction. This problem is prevalent in other emerging markets, and has implicitly driven an incorporation of liquidity considerations into VIX calculations (Tzang et al., 2010).

4.2. Measurement of Realized Volatility

The dynamic behavior of different VIXs is compared with RV as the benchmark measure of market volatility.

The theory of quadratic variation suggests that, under suitable conditions, RV is an unbiased and highly efficient estimator of volatility of returns (Andersen et al., 2001). *RV* is computed as sum of intraday squared returns. *RV* over $[0,T]$ is defined as

$$RV_T = \sum_{i=1}^{n} r^2_{\frac{it}{n}}.$$

Here $r_{\frac{it}{n}}$ refers to index returns from time $(i-1)\frac{T}{n}$ to $i\frac{T}{n}$.

For the calculation of *RV*, data on Nifty index price that is available at within-one-second intervals from the trades and orders dataset, are used. As a first step, these data are discretized at 10-minutes. These discretized data are then used to calculate daily market index volatility.

Earlier studies such as Canina and Figlewski (1993) use overlapping samples to evaluate the performance of IV estimates, although other studies such as Christensen and Prabhala (1998), Jiang and Tian (2005b), Corrado and Miller (2005) use nonoverlapping samples by using data at a lower frequency (monthly) in evaluating the performance of IV estimates.

For our analysis, all VIXs are reduced to daily values (at the end of the trading day), by dividing them by the square root of the number of calendar days, 365. Because VIXs are ex-ante measures of the volatility, each day's VIX is adjusted to the next period.

5. A VOLATILITY INDEX THAT EXPLICITLY UTILIZES LIQUIDITY IN WEIGHTS

The linkages between liquidity and IVs presented in Section 2.1 appear to lead to a calculated *VIX* value, which may be biased due to illiquidity and noncontinuous strike prices. The literature has documented that across different underlyings, options on less liquid underlyings have a larger premium compared to those on more liquid underlyings. An extreme version of the difficulties caused by illiquidity is documented in Jiang and Tian (2005a), who found that the VIX constructed by the CBOE is flawed due to truncation errors that arise from the unavailability of option data for very low and very high strikes in practice.

5.1. Spread-Adjusted VIX

Two elements of a strategy are proposed for confronting the problem of illiquidity. First, the mid-quote prices rather than traded prices are utilized. This reduces noise. Second, option IV is explicitly weighted by option liquidity, which is measured as the bid–ask spread available at that point of time in the limit order book for that option. These weights are calculated as follows:

$$\sigma_{it} = \frac{\sum_i w_{it,j}\sigma_{it}}{\sum_i w_{it,j}},$$

where $w_{it,j} = 1/s_{it,j}$ and $s_{it,j}$ refers to the percentage spread defined as (ask–bid)/mid-price of option i at t, and $j = 1, 2$ stands for the two nearest maturities.

This strategy attaches greater weight to liquid products, where observed prices or quotes have reduced noise. The lack of availability of options prices traded at a wide range of strikes is known to magnify the truncation error of the CBOE VIX calculation methodology, and increase the bias of the VIX measure. Our method automatically adjusts for the lack of data by incorporating it in the value of the spread. If there are data missing on either side of the book, the spread would take a value of infinity, and the weight attributed to that option would be zero.

5.2. Volume-Adjusted VIX

Another liquidity-adjusted weighting scheme is considered, where options are weighted using traded volume. Mayhew and Stivers (2003), Dennis, Mayhew, and Stivers (2006), Brous, Ufuk, and Ivilina (2010) have found that stocks with higher traded volume result in IV estimates that outperform historical volatility forecasts. These weights are computed as follows:

$$\sigma_{it} = \frac{\sum_i w_{it,j}\sigma_{it}}{\sum_i w_{it,j}},$$

where $w_{it,j}/\sum_i w_{it,j}$ refers to the fraction of volume traded for option i at the end of day t, and $j = 1{,}2$ stands for the two nearest maturities. Options with a higher traded volume have a greater impact on the IV estimate.

5.3. Stylized Facts on the Cross-Sectional Variation of Option Liquidity

The crucial issue that affects this research is the cross-sectional variation of option liquidity. Our empirical work is based on Indian data. This raises the concern that the results are an artifact of this emerging markets setting—perhaps one where liquidity is spotty, arbitrage is weak, or one where liquidity risk is large.

In order to evaluate this question, we plot bid–offer spreads on put options in the United States (Figure 1) and India (Figure 3). We also plot bid–offer spreads on call options in the United States (Figure 2) and India (Figure 4).

In both countries, we see high cross-sectional variation of option liquidity. If anything, option illiquidity is a smaller problem in India. Thus, our empirical

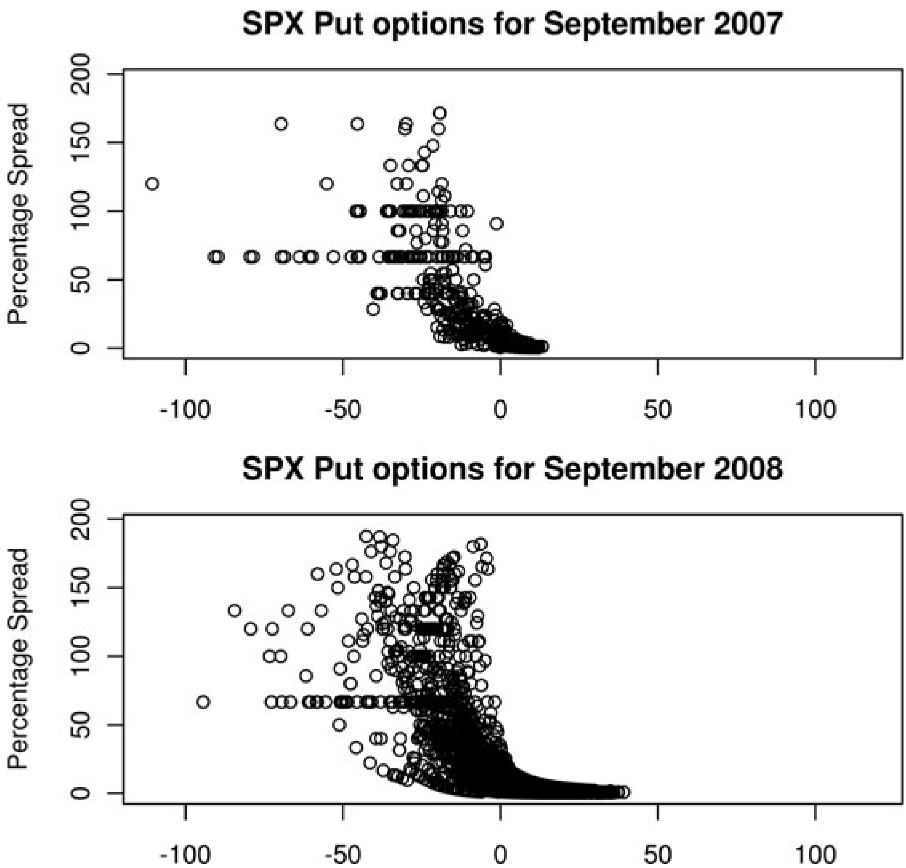

FIGURE 1

Variation of put option spreads, United States. These graphs show the relationship between percentage spread and moneyness for the U.S. market index options markets for the month of September 2007 (precrises) versus September 2008 (crises). The first is the plot of the *put options market* on the S&P 500 at the Chicago Board Options Exchange (CBOE) with near-month expiry for the month of September 2007. On the y-axis is the percentage spread (%) and on the x-axis is the moneyness of the option, calculated as (Strike − Current index level)/(Strike) and also expressed in percentage. Similarly, the second is the plot of the *put options market* for the market index at CBOE with near-month expiry for the month of September 2008. The graphs show that put spreads worsened during the crises period in the U.S. options market.

results may be biased toward understating the gains from bringing liquidity considerations integrally into the construction of an implied volatility index.

Further, in all four figures, we show that in the crisis period (September 2008), option illiquidity was a much bigger issue when compared with a tranquil period (September 2007). This suggests that the importance of this work would be enhanced under stressed market conditions.

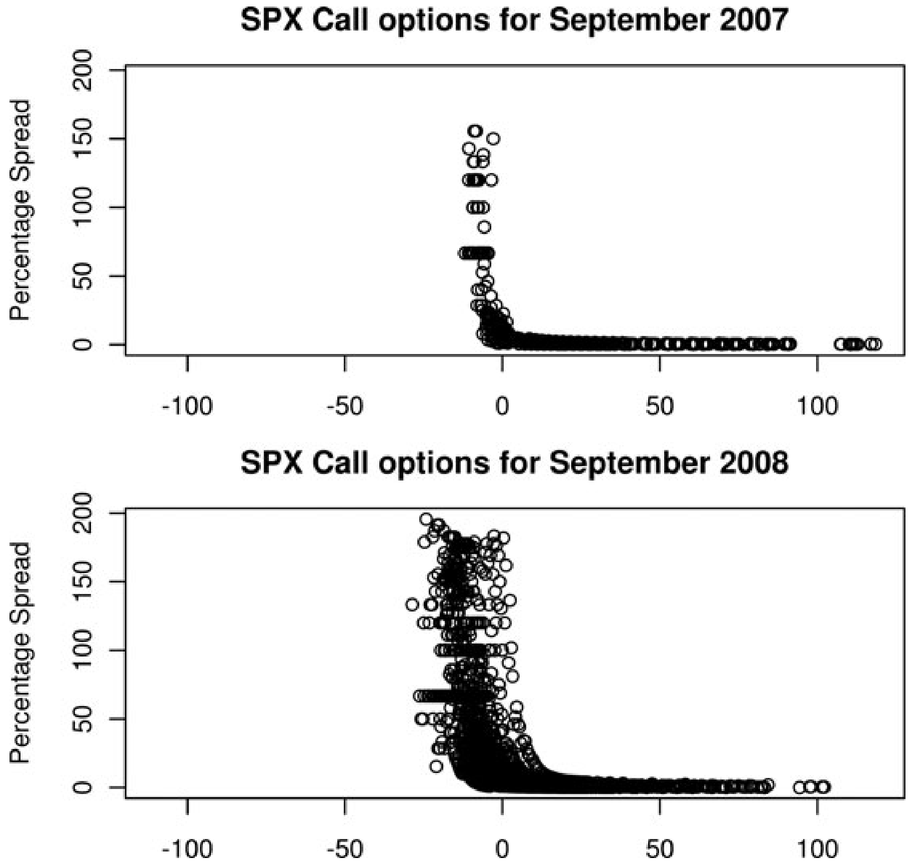

FIGURE 2

Variation of call option spreads, United States. These graphs show the relationship between percentage spread and moneyness for the U.S. market index options markets for the month of September 2007 (precrises) versus September 2008 (crises). The first is the plot of the *call options market* on the S&P 500 at the Chicago Board Options Exchange (CBOE) with near-month expiry for the month of September 2007. On the y-axis is the percentage spread (%) and on the x-axis is the moneyness of the option, calculated as (Current index level − Strike)/(Strike) and also expressed in percentage. Similarly, the second is the plot of the *call options market* for the market index at CBOE with near-month expiry for the month of September 2008. The graphs show that call spreads worsened during the crises period in the U.S. options market.

6. EMPIRICAL RESULTS

Four alternative weighting schemes have been proposed through which IV is estimated. The first weighting scheme uses the Black–Scholes vega (VVIX), the second uses volatility elasticity (EVIX), the third uses the bid–offer spread in order to construct the weights (SVIX), and the fourth uses traded volumes (TVVIX) for computing a VIX.

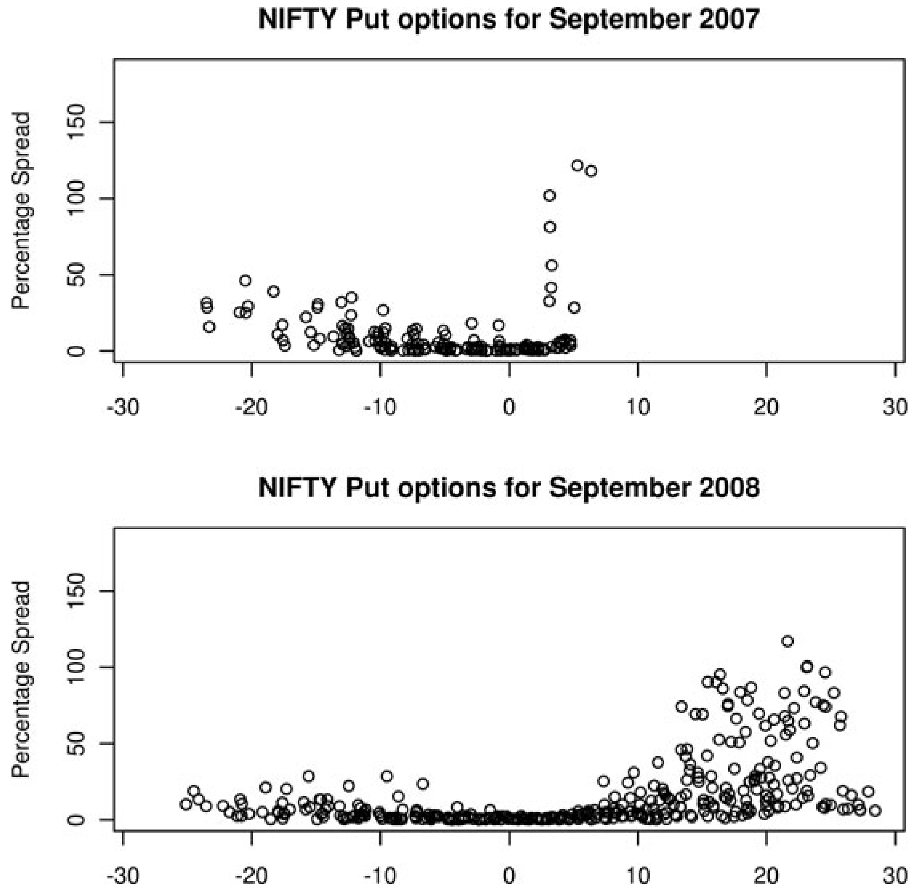

FIGURE 3

Figure 3 Variation of put option spreads, India. These graphs show the relationship between percentage spread and moneyness for the Indian market index options markets for the month of September 2007 (precrises) versus September 2008 (crises). The first is the plot of the *put options market* on the Nifty with near-month expiry for the month of September 2007. On the y-axis is the percentage spread (%) and on the x-axis is the moneyness of the option, calculated as (Strike − Current index level)/(Strike) and also expressed in percentage. Similarly, the second is the plot of the *put options market* on the Nifty with near-month expiry for the month of September 2008. The graphs show that put spreads worsened during the crises period in the Indian options market.

The performance of these four VIXs, along with the old CBOE methodology (VXO) is compared against the benchmark of RV. The performance of the VIXs is compared in the following two ways:

1. Forecasting regressions that test the information content of the VIXs. Instrumental variable regressions are also run to correct for potential errors-in-variables problems in IV estimates.

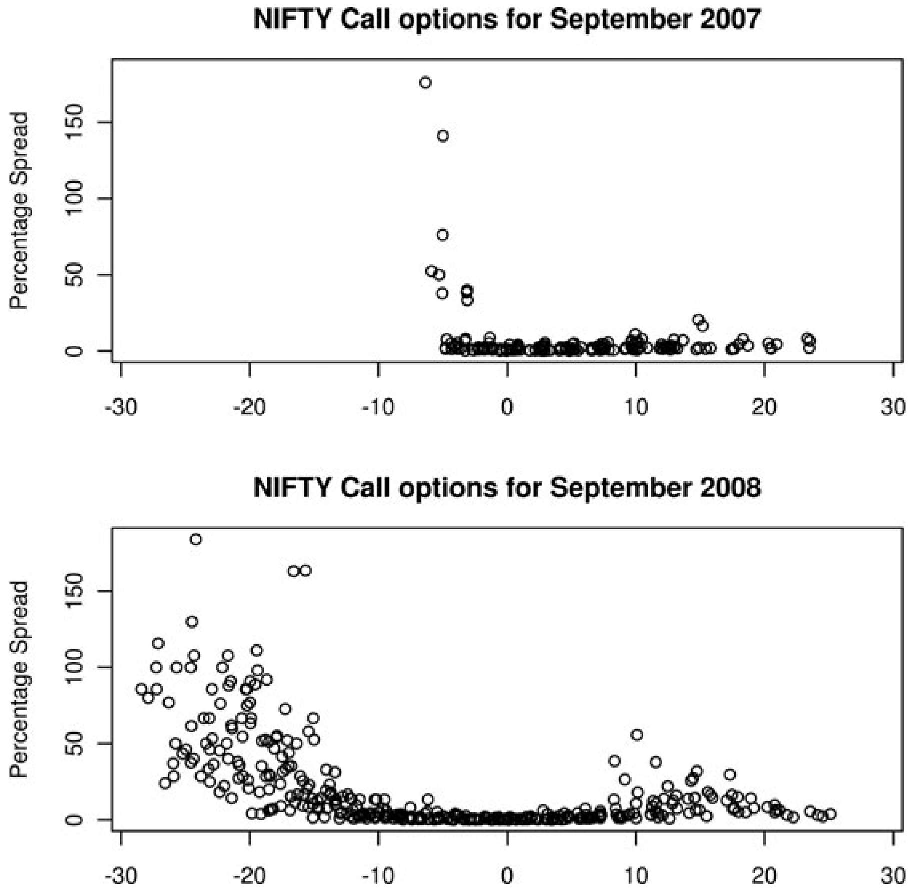

FIGURE 4

Variation of call option spreads, India. These graphs show the relationship between percentage spread and moneyness for the Indian market index options markets for the month of September 2007 (precrises) versus September 2008 (crises). The first is the plot of the *call options market* on the Nifty with near-month expiry for the month of September 2007. On the y-axis is the percentage spread (%) and on the x-axis is the moneyness of the option, calculated as (Strike − Current index level)/(Strike) and also expressed in percentage. Similarly, the second is the plot of the *call options market* on the Nifty with near-month expiry for the month of September 2008. The graphs show that put spreads worsened during the crises period in the Indian options market.

2. The MCS methodology that allows comparison of multiple volatility forecasts and chooses the volatility forecast that is best in tracking RV.

Figure 5 shows how each VIX tracks the RV. A common feature is that all the candidates appear to be an overestimate of volatility, which is measured as RV. One possible reason for the bias is that RV is computed as a sum of intraday squared returns from opening of trading to the closing[11] and does not

[11] See Section 4.2.

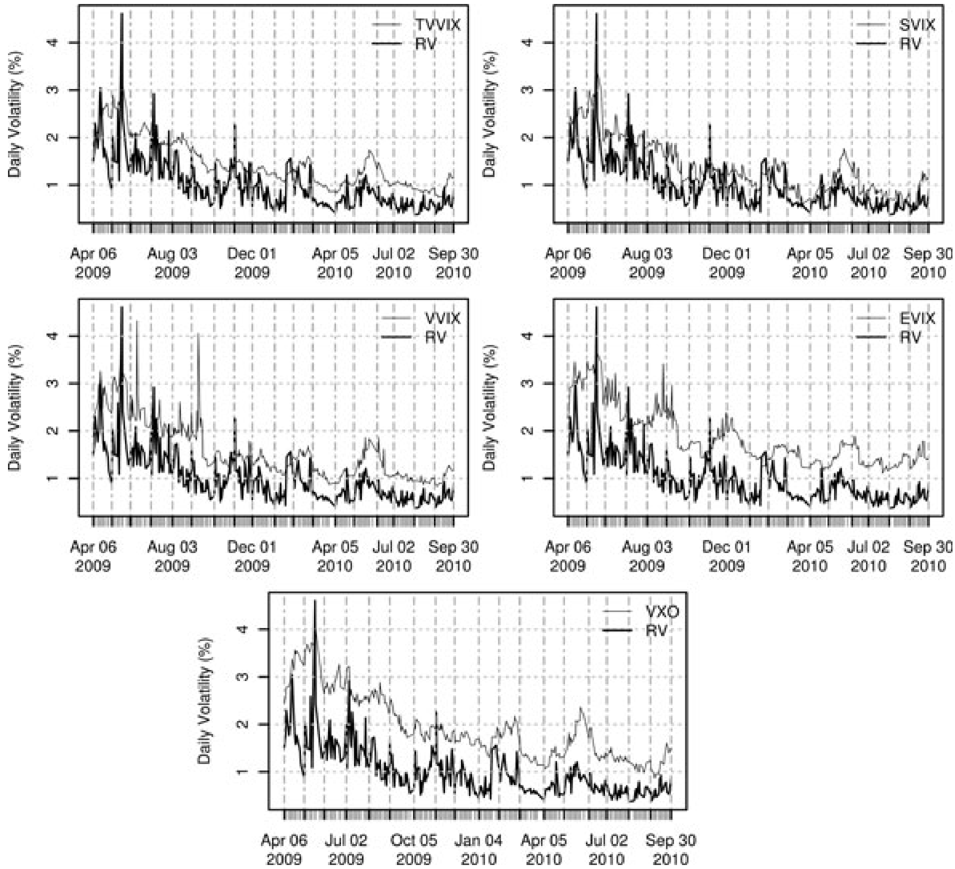

FIGURE 5
Volatility indexes versus RV.

TABLE IV
RV versus |*R*|

Nifty	*2009Q3*	*2009Q4*	*2010Q1*	*2010Q2*
$\overline{\lvert R\rvert}$	15.21	14.31	15.05	12.07
$\overline{RV}$	15.23	11.57	11.00	8.25

Note. The difference between $\overline{RV}$ and $\overline{|R|}$ for market index returns in four quarters of data is presented in this table. 2009Q3 is the quarter covering the period October, November, and December 2009 whereas 2010Q2 includes July, August, and September 2010. *RV* is computed as explained in Section 4.2 whereas $\overline{|R|}$ is computed as $|R_t| = |\ln(p_t/p_{t-1})|$, where p_t refers to closing price of the market index on day t.

include the close-to-open volatility. This is unlike the assumption about the IV as a forecast of the volatility of returns, which is calculated as price change from closing to closing of the day. For example, data on $|r|$ compared to *RV* for the four quarters in Table IV show that volatility of closing-to-closing returns tends to be higher on average than the *RV*.

TABLE V
Summary Statistics of *RV* and *IV*

	RV	*VXO*	*VVIX*	*EVIX*	*TVVIX*	*SVIX*
No. of observations	294	294	294	294	294	294
Min.	0.37	0.87	0.84	1.11	0.73	0.44
Max.	4.61	4	4.32	3.65	3.08	3.41
Mean	1.01	1.95	1.61	1.89	1.48	1.4
Kurtosis	11.22	2.86	4.74	3.29	3.16	3.31
Skewness	2.09	0.77	1.26	1.05	0.87	0.79
Standard deviation	0.52	0.68	0.62	0.59	0.49	0.56
Q1	0.64	1.39	1.14	1.43	1.07	0.99
Q2	0.90	1.80	1.43	1.71	1.36	1.29
Q3	1.25	2.46	1.97	2.23	1.84	1.81

The graph of SVIX clearly indicates that it is best in tracking *RV* followed by TVVIX and VVIX. The differences between *RV* at time t and *IV* observed at time $(t - 1)$ represent observed forecast errors.

In Table V, we give the summary statistics for each type of VIX. They are all higher on average than the corresponding *RV* series. The *IV* values are thus biased forecasts of *RV*.

6.1. Volatility Forecast Regressions

Univariate regressions of *RV* are run on each of the VIXs separately to test several hypotheses associated with the information content of the volatility measures. Regressions run on log volatility ensure that the probability density of the error term is close to the normal density and is less sensitive to outliers. However, Hansen and Lunde (2006) have shown that R^2 from log volatility, regressions cannot be used to rank models. Therefore, regressions are run on volatility series rather than the log series. If the volatility forecast contains no information about the future volatility then the slope coefficient would be zero.

We consider

$$RV_t = a_0 + a_1 VIX_{i(t-1)} + \varepsilon_t.$$

Here *VIX* belongs to the set *VXO*, *VVIX*, *EVIX*, *TVVIX*, *SVIX*.

In Table VI, we summarize the regression results. The slope coefficient is positive and significant at 1% for all VIXs indicating that all of them contain important information about future volatility.

If a given volatility forecast is an unbiased estimator of future *RV*, the slope coefficient should be one and the intercept should be zero. The null hypothesis

TABLE VI
Regression Results

Volatility Indexes	a_0	a_1	*Adj.* R^2	χ^2	*DW*
VXO	−0.14	0.59	0.59	67.7	1.68
	(0.09)	(0.00)		(0.00)	
VVIX	−0.01	0.64	0.57	21.3	1.59
	(0.94)	(0.00)		(0.00)	
EVIX	−0.16	0.62	0.51	28.4	1.37
	(0.19)	(0.00)		(0.00)	
TVVIX	−0.19	0.81	0.59	6.7	1.64
	(0.03)	(0.00)		(0.01)	
SVIX	0.03	0.70	0.57	31.4	1.72
	(0.55)	(0.00)		(0.00)	

Note. For each regression, the *T*-statistic is computed by following a robust procedure taking into account the heteroskedastic and autocorrelated error structure (Newey and West, 1987). *p*-Value is reported in the parentheses below each coefficient.

of no bias is tested using the Newey–West covariance matrix. χ^2 statistics and *p*-values are reported in Table VI. The null hypothesis is rejected at the 5% significance level in all cases with estimated coefficients ranging from 0.59 to 0.81. The Durbin-Watson (*DW*) statistic is significantly different from 2 indicating that the residuals still reflect dependence across time points.

This result is not surprising because summary statistics in Table V indicate that all the VIXs are on average greater than the *RV*. The evidence is also consistent with the existing option pricing literature that documents that stochastic volatility is priced with a negative market price of risk (or equivalently a positive risk premium). The volatility implied from option prices is thus higher than their counterpart under the objective measure due to investor risk aversion (Jiang & Tian, 2005b).

These regressions can be used to compare the performance of volatility forecasts by comparing the R^2 estimated from regression of each volatility forecast on the volatility proxy. However, it is not the best measure to compare multiple models because it does not penalize biased forecasts.[12] Regression-based comparison is also biased to the use of volatility proxy. R^2 can be used to compare the performance of multiple models only when $a_0 = 0$ and $a_1 = 1$ is satisfied. Hence, we use robust loss functions to directly compare multiple volatility forecasts in Section 6.3. However, to improve our regressions such that the estimates are robust to the error-in-variables problem, instrumental variable regressions are performed, which is discussed in Section 6.2.

[12]See Hansen and Lunde (2005).

TABLE VII
Instrumental Variable Regression Results

Dependent Variable: $\overbrace{VIX_{i(t-1)}}$					Dependent Variable: RVt					
1st Stage	a_0	a_1	*Adj. R^2*	*DW*	*2nd Stage*	b_0	b_1	*Adj. R^2*	χ^2	*DW*
VXO_{t-2}	0.04 (0.09)	0.98 (0.00)	0.95	2.28	$\overbrace{VXO_{(t-1)}}$	−0.11 (0.16)	0.58 (0.00)	0.55	1466.8 (0.00)	1.59
$VVIX_{t-2}$	0.19 (0.00)	0.87 (0.00)	0.77	2.71	$\overbrace{VVIX_{(t-1)}}$	−0.06 (0.55)	0.67 (0.00)	0.49	518.4 (0.00)	1.60
$EVIX_{t-2}$	0.11 (0.01)	0.94 (0.00)	0.88	2.62	$\overbrace{EVIX_{(t-1)}}$	−0.19 (0.11)	0.64 (0.00)	0.47	643.8 (0.00)	1.37
$TVVIX_{t-2}$	0.02 (0.20)	0.98 (0.00)	0.96	1.91	$\overbrace{TVVIX_{(t-1)}}$	−0.15 (0.09)	0.79 (0.00)	0.54	360.8 (0.00)	1.55
$SVIX_{t-2}$	0.08 (0.00)	0.94 (0.00)	0.88	2.61	$\overbrace{SVIX_{(t-1)}}$	0.01 (0.80)	0.71 (0.00)	0.53	249.2 (0.00)	1.59

6.2. Instrumental Variable Regressions

The instrumental variable approach is adopted when there may be possible errors in explanatory variables. Many studies such as Christensen and Prabhala (1998), Jiang and Tian (2005b), Corrado and Miller (2005) have discussed the possible reasons that may result in the error-in-variable problem in IV estimates, which may further bias the slope coefficient in the univariate regressions discussed earlier. Following Christensen and Prabhala (1998), a two-stage least squares regression is applied to implement the instrumental variable estimation procedure. The lagged IV is used as an instrument for IV. In the first stage, each VIX is regressed on the instrumental variable. In the second stage, RV is regressed on the fitted values obtained from the regression in the first stage.

We consider the following:

$$\overbrace{VIX_{i(t-1)}} = a_0 + a_1 VIX_{i(t-2)},$$

$$RV_t = b_0 + b_1 \overbrace{VIX_{i(t-1)}} + \varepsilon_t.$$

In Table VII, we summarize the results for the instrumental variable regressions. The null hypothesis of no bias is rejected at 1% for all VIXs. However a decrease in value of adjusted R^2 is seen for the instrumental variable regressions.

TABLE VIII
Loss Function Rankings

VIX	*MSE*	*MAD*	*QLIKE*
VXO	14.328	3.029	1.504
EVIX	11.191	2.713	1.496
VVIX	6.499	1.757	1.244
TVVIX	3.288	1.319	1.167
SVIX	**3.267**	**1.205**	**1.103**

Note. Models with best performance are in bold.

6.3. MCS Results

In Table VIII, we report the rankings of all VIXs based on their performance. The performance is measured by three types of loss functions, which are used to rank the models. The loss functions are defined as follows:

$$MSE : L_{(RV,VIX)} = \frac{\sum_i \left(VIX^2_{i(t-1)} - RV^2_{it}\right)^2}{n},$$

$$MAD : L_{(RV,VIX)} = \frac{\sum_i \left|VIX^2_{i(t-1)} - RV^2_{it}\right|}{n},$$

$$QLIKE : L_{(RV,VIX)} = \frac{\sum_i \left(\log VIX^2_{i(t-1)} + \dfrac{RV^2_{it}}{VIX^2_{i(t-1)}}\right)}{n}.$$

The *MSE* and *QLIKE* are robust loss functions and are not sensitive to use of volatility proxy. The *QLIKE* loss function penalizes underpredictions more than overpredictions. Patton and Sheppard (2009) show that the Diebold–Mariano test has highest power under the *QLIKE* loss function thus suggesting that *QLIKE* should be used over *MSE* to compare the performance of volatility forecasts. *QLIKE* and *MAD* are less sensitive to outliers than the *MSE*. The *MSE*, *QLIKE*, and *MAD* loss functions show that *SVIX* has the least error. However, under *MSE*, the *TVVIX* has an error very close to that of *SVIX*. The *QLIKE* and *MAD* loss functions show that *SVIX* has a much lower error compared to all other VIXs.

In Table IX, we report the MCS p-values for the given loss functions. Under the *MSE* loss function, it turns out that the model with the largest range statistic T is *VXO*. The p-value in the first reduction is 0.007. As it is eliminated in the first round, this automatically determines that the MCS p-value for *VXO* is 0.007.

TABLE IX
MCS Results

VIX	p_{T_r}	$MCS\ (p_{T_r})$	$p_{T_{SQ}}$	$MCS\ (p_{T_{SQ}})$
MSE				
VXO	0.007	0.007	0.000	0.000
EVIX	0.004	0.007	0.000	0.000
VVIX	0.078	0.078	0.033	0.033
TVVIX	**0.991**	**0.991**	**0.916**	**0.916**
SVIX	–	**1.000**	–	**1.000**
MAD				
VXO	0.000	0.000	0.000	0.000
EVIX	0.000	0.000	0.000	0.000
VVIX	0.000	0.000	0.000	0.000
TVVIX	0.002	0.002	0.002	0.002
SVIX	–	**1.000**	–	**1.000**
QLIKE				
VXO	0.000	0.000	0.000	0.000
EVIX	0.000	0.000	0.000	0.000
VVIX	0.000	0.000	0.000	0.00
TVVIX	0.000	0.000	0.000	0.00
SVIX	–	**1.000**	–	**1.000**

Note. Models with best performance are in bold.

In the second round, *EVIX* is eliminated with a p-value of 0.004. Because this p-value is smaller than the MCS p-value of model previously dropped, hence the MCS p-value for *EVIX* is 0.007.

In the third round, *VVIX* is eliminated with a p-value of 0.078. As this p-value is larger than the MCS p-value of model(s) previously dropped the MCS p-value for *VVIX* is 0.078. The results remain the same when we look at the semiquadratic statistic T_{SQ} rather than T_r.

The VIXs that survive in the MCS are SVIX and TVVIX whereas all other VIXs are dropped at 10% level of significance for the T_r and at 5% level of significance for the T_{SQ}.

However, when we look at the *MAD* and *QLIKE* loss functions, we see that the only VIX that survives in the MCS is SVIX whereas all other VIXs are dropped at 1% level of significance for the T and T-statistic. This shows that using loss functions such as *QLIKE* and *MAD*, which are superior to the *MSE*, clearly result in SVIX outperforming all other VIXs.

6.4. Sensitivity Analysis

The main strategy of this study has involved the bid/offer spread as a measure of option liquidity, and weights that vary inversely with the spread. There is a role for exploring alternatives to both these foundations of the research.

In recent work, Chaudhury (2011) proposes the following two alternative measures of option liquidity:

$$s = \frac{\text{ask} - \text{bid}}{\text{vol}},$$

$$\text{vol} = \frac{S\sigma}{\sqrt{252}},$$

where S refers to underlying asset price, σ refers to IV of option.

And

$$s = \frac{\text{ask} - \text{bid}}{\sigma \dfrac{\delta V}{\delta \sigma}},$$

where V refers to the mid-price of option and σ refers to IV of option.

The analysis of this study was repeated using both these measures. The VIX computed using either of these measures is inferior to our main work.[13]

Another direction of exploration is the variation of the weight by option spread. The main work of this study has employed weights $w = 1/s$. This is an ad-hoc specification lacking theoretical rationale. Hence, we also explore two alternative specifications as follows:

$$w = \frac{1}{s^2}$$

and

$$w = \frac{1}{\sqrt{s}}.$$

The former has weights that rapidly drop off, when the spread widens, and the latter has weights that drop off relatively slowly. Neither of these alternatives yield an improvement when compared with the main work.[14]

7. CONCLUSION

The VXO and VIX are widely accepted VIXs and are computed by many exchanges across the world. However, options markets show substantial

[13]Detailed results are available on request from the authors.
[14]Detailed results are available on request from the authors.

cross-sectional variation in liquidity. This cross-sectional variation is accentuated in crisis periods. Price information for illiquid options is less informative. The present strategies for construction of VIXs err in treating all price data as equally informative.

The contribution of our study lies in isolating this issue, and proposing a VIX where the option IV, which is computed using the midpoint quote, is weighted by the inverse of the bid–offer spread of the option.

Our work falls under the larger theme of bringing microstructure considerations more integrally into the utilization of information from financial markets (Shah & Thomas, 1998). Some markets that are highly liquid in industrial countries may be relatively illiquid in emerging markets. While some traded products (e.g., ATM options) might be highly liquid, other traded products might be illiquid. Although some markets may be ordinarily highly liquid in ordinary times (e.g., the U.S. Treasury Inflation Protected Securities [TIPS] market), they may become illiquid under stressed conditions. This microstructure perspective can be useful with many applications of financial market data.

Our results indicate that the liquidity-weighted VIX (SVIX) outperforms other VIXs when compared against future RV. In an ideal world, if all option series are identically liquid, then the SVIX would yield an answer that is no different from the conventional scheme: our proposed scheme does no harm at times when all options are highly liquid, but it improves matters when cross-sectional variation in option liquidity occurs. This improved methodology is thus potentially useful in improving measurement of IV at option exchanges worldwide.

In this study, the simplest strategy—weighting by the inverse spread—proved to yield a VIX that was superior to traditional methods. More generally, a superior VIX might involve utilizing information in both vega and the bid–offer spread, and can be an interesting avenue for future research.

BIBLIOGRAPHY

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001). The distribution of realized exchange rate volatility. Journal of the American Statistical Association, 96, 42–55.

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

Becker, R., Clements, A. E., & White, S. I. (2008). Are combination forecasts of S&P 500 volatility statistically superior? International Journal of Forecasting, 24, 122–133.

Beckers, S. (1981). Standard deviations implied in option prices as predictors of future stock price variability. Journal of Banking and Finance, 5, 363–381.

Blair, B. J., Poon, S. H., & Taylor, S. J. (2001). Forecasting S&P 100 volatility: The incremental information content of implied volatilities and high-frequency index returns. Journal of Econometrics, 105, 5–26.

Bollen, N. P. B., & Whaley, R. E. (2004). Does net buying pressure affect the shape of the implied volatility functions? Journal of Finance, 59, 711–753.

Brenner, M., Eldor, R., & Hauser, S. (2001). The price of options illiquidity. Journal of Finance, 56, 789–805.

Brenner, M., & Galai, D. (1993). Hedging volatility in foreign currencies. Journal of Derivatives, 1, 53–59.

Britten-Jones, M., & Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. Journal of Finance, 55, 839–866.

Brous, P., Ufuk, I., & Ivilina, P. (2010). Volatility forecasting and liquidity: Evidence from individual stocks. Journal of Derivatives & Hedge Funds, 16, 144–159.

Campa, J. M., & Chang, P. H. K. (1995). Testing the expectations hypothesis on the term structure of volatilities in foreign exchange options. Journal of Finance, 50, 529–547.

Canina, L., & Figlewski, S. (1993). The information content of implied volatility. The Review of Financial Studies, 6, 659–681.

Carr, P., & Madan, D. (1998). Towards a theory of volatility trading. In R. Jarrow (Ed.), Volatility: New estimation techniques for pricing derivatives (pp. 417–427). London: Risk.

Cetin, U., Jarrow, R., Protter, P., & Warachka, M. (2006). Pricing options in an extended Black-Scholes economy with illiquidity: Theory and empirical evidence. The Review of Financial Studies, 19, 493–529.

Chaudhury, M. (2011). Option bid-ask spread and liquidity (working paper). Quebec, Canada: McGill University.

Chiras, D. P., & Manaster, S. (1978). The information content of option prices and a test of market efficiency. Journal of Financial Economics, 6, 213–234.

Chou, R. K., Chung, S. L., Hsiao, Y. J., & Wang, Y. H. (2009). The impact of liquidity risk on option prices (technical report). Taiwan: National Central University and National Taiwan University.

Christensen, B. J., Hansen, C. S., & Prabhala, N. R. (2001). The telescoping overlap problem in options data (working paper). Denmark: University of Aarhus; and College Park, MD: University of Maryland.

Christensen, B. J., & Prabhala, N. R. (1998). The relation between realized and implied volatility. Journal of Financial Economics, 50, 125–150.

Corrado, C. J., & Miller, T. W., Jr. (2005). The forecast quality of CBOE implied volatility indexes. Journal of Futures Markets, 25, 339–373.

Cox, J., & Rubinstein, M. (1985). Options markets. Englewood Cliffs, NJ: Prentice-Hall.

Day, T. E., & Lewis, C. M. (1988). The behavior of the volatility implicit in the prices of stock index options. Journal of Financial Economics, 22, 103–122.

Demeterfi, K., Derman, E., Kamal, M., & Zou, J. (1999). A guide to volatility and variance swaps. Journal of Derivatives, 6, 9–32.

Dennis, P., Mayhew, S., & Stivers, C. (2006). Stock returns, implied volatility innovations, and the asymmetric volatility phenomenon. Journal of Financial and Quantitative Analysis, 41, 381–406.

Dumas, B., Fleming, J., & Whaley, R. E. (1998). Implied volatility functions: Empirical tests. Journal of Finance, 53, 2059–2106.

Dupire, B. (1993). Model art. Risk, 6, 118–124.
Etling, C., & Miller, T. W., Jr. (2000). The relationship between index option moneyness and relative liquidity. Journal of Futures Markets, 20, 971–987.
Galai, D. (1989). A proposal for indexes for traded call options. Journal of Finance, 34, 1157–1172.
Garleanu, N., Pedersen, L. H., & Poteshman, A. M. (2009). Demand-based option pricing. The Review of Financial Studies, 22, 4259–4299.
Gastineau, G. L. (1977). An index of listed option premiums. Financial Analyst Journal, 33, 70–75.
Hansen, P. R., & Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1, 1)? Journal of Applied Econometrics, 20, 873–889.
Hansen, P. R., & Lunde, A. (2006). Consistent ranking of volatility models. Journal of Econometrics, 131, 97–121.
Hansen, P. R., Lunde, A., & Nason, J. M. (2003). Choosing the best volatility models: The model confidence set approach. Oxford Bulletin of Economics and Statistics, 65, 839–861.
Heynen, R., Kemna, A., & Vorst, T. (1994). Analysis of the term structure of implied volatilities. Journal of Financial and Quantitative Analysis, 29, 31–56.
Jackwerth, J. C., & Rubinstein, M. (1996). Recovering probability distributions from option prices. Journal of Finance, 51, 1611–1631.
Jiang, G. J., & Tian, Y. S. (2005a). Gauging the investor fear gauge: Implementation problems in the CBOE's new volatility index and a simple solution (working paper). University of Arizona.
Jiang, G. J., & Tian, Y. S. (2005b). The model-free implied volatility and its information content. The Review of Financial Studies, 18, 1305–1342.
Jorion, P. (1995). Predicting volatility in the foreign exchange market. Journal of Finance, 50, 507–528.
Lamoureux, C. G., & Lastrapes, W. D. (1993). Forecasting stock-return variance: Toward an understanding of stochastic implied volatilities. The Review of Financial Studies, 6, 293–326.
Latane, H. A., & Rendleman, R. J., Jr. (1976). Standard deviations of stock price ratios implied in option prices. Journal of Finance, 31, 369–381.
Mayhew, S., & Stivers, C. (2003). Stock return dynamics, option volume, and the information content of implied volatility. Journal of Futures Markets, 23, 616–646.
Nandi, S. (2000). Asymmetric information about volatility: How does it affect implied volatility, options prices and market liquidity? Review of Derivatives Research, 3, 215–236.
Neuberger, A. (1994). The log contract: A new instrument to hedge volatility. Journal of Portfolio Management, 2, 74–80.
Newey, W. K., & West, K. D. (1987). Simple positive definite heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica, 55, 703–708.
Patton, A. J., & Sheppard, K. (2009). Evaluating volatility and correlation forecasts. In T. G. Andersen, R. A. Davis, J. P. Kreiss, & T. Mikosch (Eds.), The handbook of financial time series. Berlin: Springer Verlag.
Poteshman, A. M. (2000). Forecasting future volatility from option prices (working paper). University of Illinois at Urbana-Champaign.

Rubinstein, M. (1994). Implied binomial trees. The Journal of Finance, 49, 771–818.
Schmalensee, R., & Trippi, R. (1978). Common stock volatility expectations implied by option premia. Journal of Finance, 33, 129–147.
Shah, A., & Thomas, S. (1998). Market microstructure considerations in index construction. In proceedings of CBOT research symposium, Chicago Board of Trade, Chicago, pp. 173–193.
Trippi, R. R. (1977). A test of option market efficiency using a random-walk evaluation model. Journal of Economics and Business, 29, 93–98.
Tzang, S. W., Hung, C. H., Wang, C. W., & Shyu, D. S. D. (2010). Do liquidity and sampling methods matter in constructing volatility indices? Empirical evidence from Taiwan. International Review of Economics and Finance, 20, 312–324.
Whaley, R. E. (1982). Valuation of American call options on dividend-paying stocks: Empirical tests. Journal of Financial Economics, 10, 29–58.
Whaley, R. E. (1993). Derivatives on market volatility: Hedging tools long overdue. Journal of Derivatives, 1, 71–84.
Whaley, R. E. (2009). Understanding VIX. Journal of Portfolio Management, 35, 98–105.
Xu, X., & Taylor, S. J. (1994). The term structure of volatility implied by foreign exchange options. Journal of Financial and Quantitative Analysis, 29, 57–74.