

The imprecision of volatility indexes

Rohini Grover* Ajay Shah[†]

August 2, 2014

Abstract

Concerns about sampling noise arise when a VIX estimator is computed by aggregating several imprecise implied volatility estimates. We propose a bootstrap strategy to measure the imprecision of a model based VIX estimator. We find that the imprecision of VIX is economically significant. We propose a model selection strategy, where alternative statistical estimators of VIX are evaluated based on this imprecision.

Keywords: Implied Volatility; Volatility index; Imprecision
JEL Classification: G12; G13; G17

*We are grateful to Jayanth Varma and Susan Thomas for discussions that led up to this paper. Rohini Grover is with Finance Research Group, Indira Gandhi Institute of Development Research, Gen A.K. Vaidya Marg, Goregaon (E), Mumbai, India 400065, email: grover.rohini@gmail.com, phone: +91-22-2841-6592.

[†]Macro/Finance Group, National Institute of Public Finance and Policy, 18/2 Satsang Vihar Marg, Special Institutional Area, New Delhi, India 110067, email ajahshah@mayin.org, phone +91-11-2656-9303.

Contents

1	Introduction	3
2	Concerns about measurement	5
2.1	Model based approach	6
2.2	Model free approach	6
3	Measuring the imprecision of a volatility index	7
3.1	A bootstrap inference strategy for the VIX	7
4	Data description	8
4.1	S&P 500 index (SPX)	8
4.2	S&P CNX Nifty index	9
5	Two empirical examples	11
5.1	Estimating confidence interval using SPX options	11
5.2	Estimating confidence interval using Nifty options	12
6	How large is the imprecision of VVIX	14
6.1	Imprecision of VVIX	14
6.2	Comparing against earlier study	14
6.3	Implications	16
7	Is this measure of imprecision a mere proxy for spot market liquidity?	17
8	Using this measure of imprecision for model selection	20
8.1	Methodology	20
8.2	Pairwise comparisons of VIXs	21
9	Reproducible research	23
10	Conclusion	24

1 Introduction

The volatility index (VIX) is a measure of market's expectation of volatility. The usefulness of the forward-looking VIX, over volatility forecasts based on historical data, has been explored extensively in the literature [Blair et al., 2001, Jiang and Tian, 2005, Corrado and Miller, 2005, Giot, 2005]. This has led to applications of VIX in areas such as option pricing and value-at-risk calculations. The negative correlation of VIX with the market index has also been employed as a market timing device for trading strategies that result in positive returns by switching between two portfolios based on whether VIX rises or falls [Copeland and Copeland, 1999, Arak and Mijid, 2006, Cipollini and Manzini, 2007]. Derivatives on VIX exploit this negative correlation to allow investors to obtain direct exposure to the market's volatility and diversify portfolios. A recent literature has started exploring the effects of macroeconomic shocks to aggregate uncertainty, and VIX is a prominent tool for the measurement of this uncertainty [Bloom, 2009, Bekaert et al., 2013].

Implied volatility indexes are in use worldwide. In the US, CBOE introduced VIX futures and options on March 26, 2004 and February 24, 2006 respectively. In 2012, the open interest for these contracts was at 326,066 contracts and 6.3 million respectively.

In academic research and in real-world applications, VIX estimators are now widely used. Each of these applications treats the observed VIX value as a hard number. However, every VIX estimate is a statistical aggregation of a set of underlying option prices. This raises concerns about the imprecision of VIX [Latane and Rendleman, 1976, Chiras and Manaster, 1978, Jiang and Tian, 2007]. In the previous literature, Hentschel [2003] measures the precision of IV and the old CBOE VIX by estimating confidence intervals derived from the Black-Scholes formula.

We propose viewing each traded option as a noisy scientific instrument, that gives one imprecise reading of the implied volatility. Multiple different traded option series are tantamount to having multiple scientific instruments, each of which gives one noisy reading of the true implied volatility at that point in time. Each reading of implied volatility is noisy, owing to problems such as microstructure effects, and the violations of pricing efficiency that cannot be profitably arbitrated away. From this perspective, the VIX methodology is a location estimator. This motivates a new bootstrap strategy for estimating the distribution of one model based VIX estimator – the vega weighted VIX (VVIX).

It can be argued that this measure of imprecision is merely a transform of illiquidity in the spot market, and thus does not constitute a novel addition to our information set. Hence, we explore relationships between impact cost on the spot market and our imprecision measure. We find that there is a low relationship between spot market liquidity and the imprecision of VIX. The proposed new measure is, thus, not a mere transform of impact cost on the spot market.

Our empirical examples suggest that the imprecision of VIX, that is estimated by our methods, is economically significant in absolute terms and when compared with the overnight change in VIX.

As an example, in the empirical analysis shown here, when the point estimate of the VVIX was 17.82%, the 95% confidence interval ran from 16.03% to 19.91%, which is a fairly wide confidence interval. Imprecision in VIX translates into imprecision in option prices that are computed using the VIX. In the empirical example, the point estimate for an out of the money call option was 1.92. However, applying the two end-points of the 95% confidence interval for VIX yields a range of option prices from 0.89 and 3.86.

We find that the median value of the width of the 95% confidence interval for VVIX is a large number when compared with the one-day change in VVIX. This suggests that caution should be exercised in interpreting changes in VIX over short periods. Research and real-world applications which utilise VIX may benefit from taking this imprecision into account.

Once a measure of the imprecision of a statistical estimator is in hand, a natural next step is to utilise this for the comparison of alternative estimators. We apply this model-selection strategy to a group of model based VIX estimators from [Grover and Thomas \[2012\]](#). Pairwise comparisons of VIX estimators shows that the vega weighted VIX is the best estimator, followed by the liquidity weighted VIXs, and the elasticity weighted VIX.

In a simple Black/Scholes setting, economic agents have a complete information set as a consequence of knowing the price and the implied volatility at all points in time. In recent years, there has been considerable interest in ‘ambiguity’: where economic agents are exposed to risk but do not know the probability distributions of the shocks that they face. Many recent papers have proposed measures of ambiguity: [Brenner and Izhakian \[2011\]](#) use the variance of probability of gain or loss, [Baltussen et al. \[2013\]](#) use the volatility of implied volatility, and [Ehsani et al. \[2013\]](#) use the mean divergence between implied probability distributions. The proposed measure of imprecision of VIX, of this paper, might prove to be useful as quantifying the

extent of ambiguity that is present at a point in time.

The remainder of this paper is organised as follows: Section 2 discusses the sources of errors that may hamper the estimation of a volatility index. Section 3 discusses the model-free strategy to measure the imprecision of a volatility index. Section 4 describes the data. Section 5 shows the working of the methodology through an example. Section 6 discusses the results for the entire period of analysis. Section 7 examines the relationship between imprecision and liquidity. Section 8 discusses the use of this measure of imprecision for model selection. Section 10 concludes.

2 Concerns about measurement

Volatility indexes are computed using a chain of option prices at different strikes and maturities. There are two broad ways to compute a volatility index. The first involves summarising implied volatilities of each observed option price, using an option pricing formula.¹ The second involves deriving a formula for VIX from the concept of fair value of a variance swap. This is termed the ‘*model free* method’, and is currently employed by CBOE to compute VIX every 15 seconds.²

The approach of this paper is to see alternative methods for computing volatility indexes as competing statistical estimators. Microstructure noise and the limits of arbitrage introduce imprecision into each observed option price. Every observed option price suffers from a certain degree of measurement error. In the presence of this noise, every methodology for computing the volatility index is a statistical location estimator that works with noisy data. This merits an examination of the imprecision of the location estimator, and an assessment about the extent to which this imprecision is important in various applications. Before our treatment of this statistical problem, we briefly discuss the sources of noise in observed option prices.

¹The most common example of an index that employs this procedure is the old CBOE VIX (VXO) [Whaley, 1993]. It is computed from eight near-the-money options for the two nearest maturities. Other examples of the model based approach include aggregating IVs weighted by vega, liquidity, and volatility elasticity [Grover and Thomas, 2012].

²The concept of a model-free implied variance was first introduced by Dupire [1993] and Neuberger [1994] and further developed by Carr and Madan [1998], Demeterfi et al. [1999] and Britten-Jones and Neuberger [2000]. In contrast to the VXO, VIX uses a wide range of out-of-the-money (OTM) options to estimate the expected volatility of the market.

2.1 Model based approach

[Hentschel \[2003\]](#) discuss measurement error in option prices that result from finite-quote precision, bid-ask spreads and non-synchronous prices of options and underlying. He finds that fairly small errors in measured option prices result in large errors in IVs. This effect is more pronounced for OTM options. For measuring the precision of IV estimates, [Hentschel \[2003\]](#) estimates confidence bands derived from the Black-Scholes formula based on several assumptions about the distribution of the errors. He assumes that the errors stemming from bid-ask quotes are normally distributed with zero means and standard deviation proportional to the magnitude of the bid-ask spread. He shows that for an ATM stock option with twenty days to expiry, the ninety-five percent confidence intervals are of the order ± 6 percentage points. In contrast, he finds VXO to be an efficient estimator with 95% confidence intervals of the order ± 25 basis points based on a given price level and a set of measurement error assumptions, computed for a pair of strike prices and maturities sampled from a range of strike prices and maturities.

2.2 Model free approach

[Jiang and Tian \[2007\]](#) discuss methodological errors that may induce inefficiencies in the model free estimates of implied variance. They discuss four types of errors: truncation, discretisation, expansion, and maturity interpolation:

1. Only a finite range of strikes prices are available. This induces truncation error in VIX and causes a downward bias in the estimate.
2. The absence of a continuum of strike prices as a result of availability of strikes only in discrete increments leads to discretisation error. This mainly stems out from using a numerical integration procedure to approximate the integrals in the VIX formula.
3. The Taylor series expansion of the log function in the VIX formula induces a negative error in the expansion.
4. The last error arises from the linear maturity interpolation scheme used to compute VIX. The evidence on the term structures of the implied variances reveals that they may not be linear or monotonic in option maturity.

Building on these ideas, [Andersen et al. \[2011\]](#) show that varying ranges of strike prices induce considerable error in intra-day estimates of VIX.

3 Measuring the imprecision of a volatility index

We treat each option price as an imprecise transformation of the true implied volatility index. In this perspective, the computational method for the volatility index is a statistical estimator. We then propose a bootstrap strategy for estimating the imprecision of the estimator.

While this approach is fully general, for the purposes of this paper, we analyse one specific volatility index estimator: the vega adjusted volatility index (VVIX). This model-based estimator is interesting, when compared with the model-free estimator, in the interests of greater generality. While the model-free approach is used at the CBOE, in many real-world settings, this is constrained by the illiquidity in the options market, and a wide inter-strike spacing.

The VVIX is computed from all option prices as follows:

1. Estimation of IVs using the Black-Scholes model for the two nearest maturities.
2. Computation of the average weighted IV for each maturity i :

$$IV_i = \frac{\sum_{j=1}^n w_{ij} IV_{ij}}{\sum_{j=1}^n w_{ij}}$$

where, IV_{ij} refers to a vector of IVs for $j = \{1 \dots n\}$ and two nearest maturities, $i = \{near, next\}$, w_{ij} refers to the vega weight for the corresponding IV_{ij} .

3. The vega weighted average IVs are interpolated to compute the 30 day expected volatility, VVIX.

3.1 A bootstrap inference strategy for the VIX

The question of the imprecision of a location estimator with noisy data is faced in LIBOR estimation. In that setting, there is a true unobserved price on an OTC market. Each dealer, who is polled for the measurement of LIBOR, reports a noisy estimate of this true price. The LIBOR methodology is a robust estimator that utilises these pieces of information to construct a more precise estimate [Cita and Lien, 1992, Berkowitz, 1999, Shah, 2000]. Bootstrap inference is used to obtain confidence intervals.

By analogy, in our setting, each option provides an imprecise IV number, several of which are aggregated into a volatility index. We propose a bootstrap estimator through which we construct the distribution of VVIX as follows.

1. At any point in time, we observe a chain of option series for different strikes and two nearest maturities.
2. We estimate the implied volatilities for every option using the Black-Scholes model.
3. At each maturity, we sample with replacement among all the observed implied volatilities to construct a bootstrap replicate. This gives two bootstrap datasets for each of the two nearest maturities.
4. Each of these datasets is summarised into a vega weighted average IV.
5. The vega weighted IVs are interpolated to obtain the VVIX estimate.
6. This procedure is repeated R times to create a bootstrap distribution of the statistic.
7. The standard deviation (σ) and confidence bands are computed using the adjusted bootstrap percentile method [Efron, 1987].

This is a direct adaptation of bootstrap inference for the problem at hand. We now turn to examining how this estimator works with some real world data.

4 Data description

4.1 S&P 500 index (SPX)

The end-of-day data on SPX index options at the CBOE obtained from Bloomberg is used to estimate the volatility index values. The data is available for the months of September, October and November 2010. This includes: transaction date, expiry date of the options contract, strike price, type of the option i.e. call or put, price of the underlying index, best bid and ask price of options. The risk-free rates used are the one and three month US Treasury bill rates provided by the US department of the Treasury.

Options with zero bid/ask prices, and negative bid-ask spreads are discarded from the dataset. Further, we also check whether the market price of options lie within the no-arbitrage band of the Black-Scholes model. Options that violate this condition are also discarded from the set.

Table 1 Global exchanges: number of contracts traded and/or cleared

Rank	Exchanges	Jan-Dec 2012	Annual % change
1	CME Group	2,890,036,506	-14.7%
2	Eurex	2,291,465,606	-18.8%
3	National Stock Exchange of India	2,010,493,487	-8.6%
4	NYSE Euronext	1,951,376,420	-14.5%
5	Korea Exchange	1,835,617,727	-53.3%

Source: FIA, <http://www.futuresindustry.org/volume-.asp>

Table 2 Ranked by the number of contracts traded in equity index

1	Kospi 200 Options, KRX	3,671,662,258	1,575,394,249	-51.7%
2	S&P CNX Nifty Options, NSE India	868,684,582	803,086,926	-7.6%
3	SPDR S&P 500 ETF Options, CME	729,478,419	585,945,819	-19.7%
4	E-mini S&P 500 Futures, CME	620,368,790	474,278,939	-23.5%
5	RTS Futures, Moscow Exchange	377,845,640	321,031,540	-15.0%

Source: FIA, <http://www.futuresindustry.org/volume-.asp>

4.2 S&P CNX Nifty index

Data on the NSE-50 (Nifty) index options at the National Stock Exchange (NSE) of India Ltd is used to estimate the volatility index values. NSE is the fourth largest derivative exchange in the world in terms of number of contracts traded (Table 1). It is also the second largest exchange in terms of number of contracts traded in equity index (Table 2). NSE is thus a source of high quality data about exchange-traded derivatives.


The options data is available at the tick-by-tick frequency. This includes: transaction date, expiry date of the options contract, strike price, type of the option i.e. call or put, price of the underlying index, best buy and ask price of options. The risk-free rates used are the one and three month MIBOR rates provided by NSE.

A volatility index number is estimated for a given trading interval. For example, the CBOE disseminates real-time values of VIX updated every fifteen seconds. The VIX may be contaminated with market microstructure noise if calculated at a very high frequency. Thus, we also conduct all estimations with data sampled at a frequency of fifteen seconds. For a robustness check, sampling frequencies of thirty and sixty seconds are also used.

The tick-by-tick data on options, which has approximately 200,000 observations per day, is sampled at a frequency of fifteen seconds using the procedure outlined in Andersen et al. [2011] to estimate a volatility index. The sam-

Figure 1 The distribution of vVIX: One example, SPX

The kernel density plot captures the distributions of the weighted average IVs for the two nearest maturities referred to as IV_{Near} and IV_{Next} respectively and the corresponding vVIX. The distributions are estimated by drawing 1000 bootstrap replicates from a particular sample of SPX options, end-of-day on 2010-09-17. The dashed lines indicate the point estimate of vVIX and its 95% confidence bounds (LB, UB). The distribution of vVIX overlaps with the distribution of IV_{Near} . The interpolation scheme used in the computation of vVIX assigns higher weight to IV_{Near} when first month options have higher liquidity. This results in the distribution of vVIX being biased towards IV_{Near} . This pattern reverses when first month options near expiration and IV estimates from these options are more noisy.



pling procedure involves constructing fifteen seconds series for each individual option using the *previous tick method*. If no new quote arrives in a fifteen second interval, the last available quote prior to the interval is retained. If no quote is available in the previous interval, the last available quote from the last five minutes is retained.

In addition to the above, options with zero traded volume, zero bid/ask prices, and negative bid-ask spreads are discarded from the dataset. Further, we also check whether the market price of options lies within the no-arbitrage band of the Black-Scholes model. Options that violate this condition are also discarded from the set.

5 Two empirical examples

As an example of the working of this methodology, we estimate the distribution of $VVIX$ at one point in time for a sample of SPX and Nifty options respectively.

5.1 Estimating confidence interval using SPX options

As the first example, we show the estimation of a confidence interval for $VVIX$ from a particular sample of end-of-day SPX options, on 2010-09-17³.

Table 3 Near-the-money SPX options for the first month maturity

The underlying is at 1125.59, the number of days to expiry is 29 and the risk-free rate is 0.12%.

Strike	Type	Mid-Quote	IVol (%)	Strike	Type	Mid-Quote	IVol (%)
1095	c	42.50	19.31	1095	p	14.00	21.29
1100	c	38.00	18.30	1100	p	15.25	20.87
1105	c	35.25	18.71	1105	p	16.65	20.48
1110	c	31.75	18.35	1110	p	18.05	19.99
1115	c	28.60	18.14	1115	p	19.70	19.59
1120	c	25.75	18.05	1120	p	21.55	19.24
1125	c	22.75	17.70	1125	p	24.55	19.68
1130	c	19.35	16.90	1130	p	26.30	19.00
1135	c	16.85	16.66	1135	p	28.10	18.22
1140	c	14.00	15.98	1140	p	30.85	18.05
1145	c	12.35	16.12	1145	p	33.65	17.79
1150	c	10.50	15.94	1150	p	36.45	17.37
1155	c	8.55	15.49	1155	p	39.75	17.22

Note: We define near-the-money-options as call and put options with strike-to-spot ratio between 0.97 and 1.03 [Pan and Poteshman, 2006].

In order to fix intuition, Table 3 shows the near-the-money SPX options for the near month only. As we see, there is considerable variation in the IV. If a sample mean of these were computed it would have a 95% confidence interval from 17.65 to 18.84 which shows a wide band of width 1.19 percentage points.

The $VVIX$ estimated from the original sample is 21.53%. Bootstrap replicates are drawn from the original sample and the distribution of $VVIX$ is

³The numerical values in this sample and a bootstrap replicate of this sample are available at http://ifrogs.org/releases/GroverShah2014_volatilityIndexes.html.

estimated as described in Figure 1. The lower and upper 95% confidence bounds estimated from the bootstrap distributions of vVIX are 20.8% and 22.32% respectively.

This measure of imprecision, of roughly 1.5 percentage points, is an economically significant one. For a sense of scale, the one-day change in vVIX is smaller than 1.5 percentage points on 62% of the days.⁴ This suggests that on 62% of the days, we know little about whether vVIX went up or down when compared with the previous day.

The above example demonstrates that the level of imprecision found in a vIX estimate computed from SPX options may not be trivial. The level of imprecision in vVIX captured by its wide confidence band of width 1.5 percentage points is economically significant and cannot be ignored.

5.2 Estimating confidence interval using Nifty options

As the second example, we show the estimation of a confidence interval for vVIX from a particular sample of Nifty options, from 14:59:45 to 15:00:00 on 2010-09-01⁵.

Table 4 Near-the-money NIFTY options for the first month maturity

Strike	Type	Underlying	Mid-Quote	Maturity (Days)	Risk-free (%)	IVol (%)
5500	c	5464.75	72.57	29	6.29	12.43
5600	c	5464.75	30.55	29	6.29	11.57
5400	c	5464.75	133.95	29	6.29	13.11
5600	p	5464.75	160.65	29	6.29	15.75
5400	p	5464.75	68.75	29	6.29	17.80
5500	p	5464.75	105.30	29	6.29	16.48

Note: We define near-the-money-options as call and put options with strike-to-spot ratio between 0.97 and 1.03 [Pan and Poteshman, 2006].


Table 4 shows the near-the-money Nifty options for the near month only. As we see, there is considerable variation in the IV. If a sample mean of these were computed it would have a 95% confidence interval from 12.53 to 16.52 which shows a wide band of width 4 percentage points.

⁴Out of a total of 60 days.

⁵The numerical values in this sample and a bootstrap replicate of this sample are available at http://ifrogs.org/releases/GroverShah2014_volatilityIndexes.html.

Figure 2 The distribution of vVIX: One example, NIFTY

The kernel density plot captures the distributions of the weighted average IVs for the two nearest maturities referred to as IV_{Near} and IV_{Next} respectively and the corresponding vVIX. The distributions are estimated by drawing 1000 bootstrap replicates from a particular sample of NIFTY options, from 14:59:45 to 15:00:00 on 2010-09-01. The dashed lines indicate the point estimate of vVIX and its 95% confidence bounds (LB, UB). The distribution of vVIX overlaps with the distribution of IV_{Near} . The interpolation scheme used in the computation of vVIX assigns higher weight to IV_{Near} when first month options have higher liquidity. This results in the distribution of vVIX being biased towards IV_{Near} . This pattern reverses when first month options near expiration and IV estimates from these options are more noisy.



The vVIX estimated from the original sample is 17.82%. Bootstrap replicates are drawn from the original sample and the distribution of vVIX is estimated as described in Figure 3. The lower and upper 95% confidence bounds estimated from the bootstrap distributions of vVIX are 16.03% and 19.91% respectively.

This measure of imprecision, of roughly 4 percentage points, is an economically significant one. For a sense of scale, the one-day change in vVIX is smaller than 4 percentage points on 92% of the days.⁶ This suggests that on 92% of the days, we know little about whether vVIX went up or down when compared with the previous day. This level of imprecision is higher in comparison to the imprecision number estimated using SPX options.

⁶Out of a total of 379 days.

6 How large is the imprecision of VVIX

The subsection above was about only one example. We now present evidence about imprecision that runs over a bigger dataset for Nifty options. We also discuss the differences of our methodology against the earlier study and the resulting difference in inferences.

6.1 Imprecision of VVIX

For the period February 2009 - September 2010, the median imprecision observed in VVIX measured in terms of confidence interval width is 2.9 percentage points (Table 5), which is a large number when compared with the one-day change in VVIX which has a median value of 1.18 percentage points. In addition, the median standard deviation of the bootstrap estimates for the period is 0.739 percentage points (Table 5).

Table 5 Measures of imprecision for VVIX

	Width of CI	σ
No. of Obs.	1516	1516
Min	1.03	0.25
1st Qu	2.27	0.58
Median	2.92	0.74
Mean	3.91	0.94
3rd Qu	4.06	1.02
Max	50.49	11.78
Std Dev	3.64	0.75

Note: Both in percentage points.

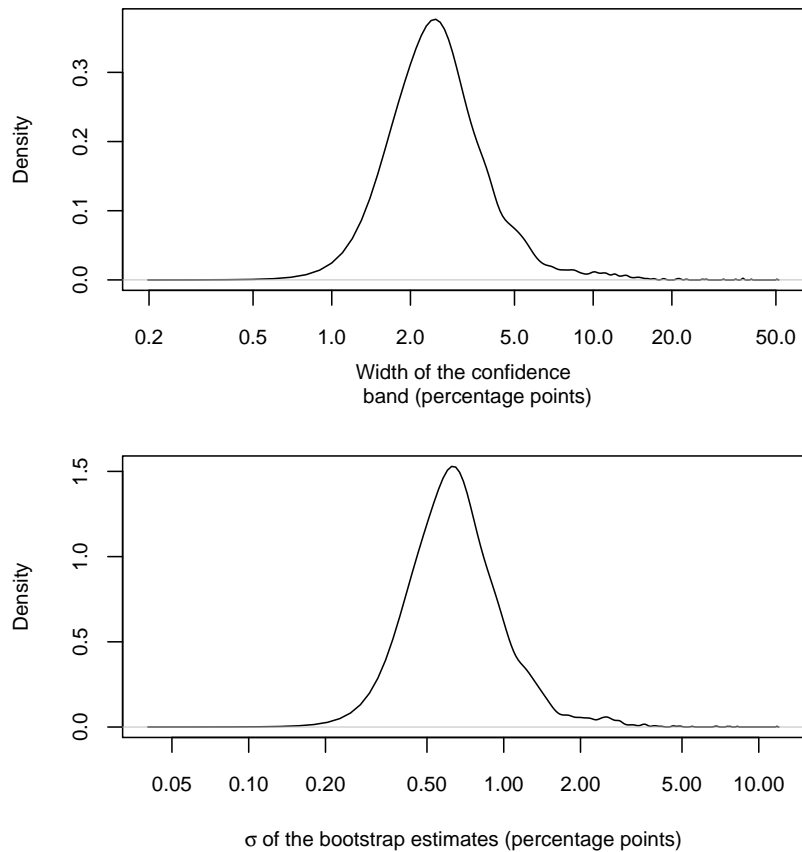
Figure 3 captures the distributions of the measures of imprecision: the width of the confidence band and the standard deviation of the bootstrap estimates for the VVIX estimator for the period February 2009 - September 2010.

6.2 Comparing against earlier study

Based on several assumptions about the distribution of the errors, [Hentschel \[2003\]](#) shows that for an ATM stock option with twenty days to expiry, the ninety-five percent confidence intervals are of the order of ± 6 percentage points. This is an economically significant scale of imprecision. While individual option prices are observed with low precision, [Hentschel \[2003\]](#) argues

Figure 3 Measures of imprecision for VVIX

The figure plots the kernel density for the width of the confidence band and the standard deviation of the bootstrap estimates measured in percentage points. The 95% confidence interval and the standard deviation are computed four times a day, for the period February 2009 - September 2010. For each point in time, the confidence interval and the standard deviation are computed from the bootstrapped sampling distribution of VVIX.



that the old CBOE VIX is fairly efficient, with 95% confidence intervals of the order of ± 25 basis points. This is because the VXO is estimated from near-the-money options that tend to be less contaminated with errors, low weights are attached to options near expiry, and errors in underlying prices cancel when averaging implied volatilities from puts and calls. By this argument, even though each option price is measured with imprecision, volatility indexes are fairly precisely measured. We on the contrary find a substantially higher level of imprecision in the volatility indexes. This may be attributed to the following reasons.

First, we employ a model-free strategy to gauge the level of imprecision in the implied volatility index. On the other hand, [Hentschel \[2003\]](#) derives confidence bands from the Black-Scholes model. Second, our methodology does not require any assumptions on the distribution of errors. In contrast, [Hentschel \[2003\]](#) assumes that the errors stemming from bid-ask quotes are normally distributed with zero mean and standard deviation proportional to the magnitude of the bid-ask spread. Third, [Hentschel \[2003\]](#) finds VXO to be an efficient estimator with 95% confidence intervals of the order ± 25 basis points based on a given price level and a set of measurement error assumptions, computed for a pair of strike prices and maturities sampled from a range of strike prices and maturities. We on the contrary find VVIX to be observed with significant imprecision. Our analysis is based on a large set of data which includes 1,516 samples of the Nifty index options. Each sample consists of options for a range of strike prices, and two maturities.

In addition to the above, there is a difference in the choice of options used for the estimation of VXO and VVIX. Our non-parametric methodology applies to a large set of model based volatility indexes estimated from all options that give a market price. The VXO however is estimated from a fixed choice of eight options corresponding to two strikes and two maturities with limited information content.

6.3 Implications

In recent years, numerous applications of volatility indexes have sprung up. Our bootstrap strategy for understanding the imprecision of VVIX suggests that the observed estimates of the volatility index contain significant imprecision. This evidence of the imprecision of volatility indexes may suggest greater caution in the use of these indexes.

7 Is this measure of imprecision a mere proxy for spot market liquidity?

It can be argued that the measures of imprecision of the volatility index is a mere transform of the liquidity of the underlying asset. We analyse the relationship between the imprecision indicators and liquidity of the Nifty index to understand whether or not these measures are a mere proxy for liquidity. We measure liquidity by computing the impact cost of the Nifty index for an order size of Rs.2.5 million. For the computation of impact cost, data on bid and ask prices of the Nifty stocks available at daily and fifteen seconds interval from the orders dataset, are used. The imprecision indicators are also computed for the daily and fifteen seconds frequency using the data inputs described in Section 4. For the daily analysis, the measures are computed half an hour before the close of the market i.e. at 3 p.m. for the period July 2012 - March 2014. The intra-day analysis is conducted for the month of March 2014.

Figure 4 plots the linear relation between the daily imprecision indicators and the liquidity of the market index. There is a positive relationship between the imprecision and liquidity measures. Similarly, Figure 5 captures the intra-day relationship between these measures. Contrary to our daily analysis, we find a negative relationship between the imprecision and liquidity measures.

We test the relationship between the imprecision indicators and liquidity by using the following robust regression:

$$\text{Imp}_{it} = \alpha_0 + \alpha_1 \text{IC}_t + \epsilon_t$$

where Imp_{it} are the imprecision indicators: width of the confidence interval (Width of CI) and standard deviation of the bootstrap estimates (σ). IC is the impact cost of the Nifty index.

The daily regression results presented in Table 6 show that the slope coefficient is positive and significant at 10%. The adjusted R^2 is low at 0.01%. For the intra-day regressions, the slope coefficient is negative and also significant. The adjusted R^2 is 0.03%. The extremely low R^2 values imply a weak relationship between the imprecision indicators and liquidity.

Figure 4 Daily imprecision measures vs. impact cost

The figure plots the linear relationship between the daily imprecision indicators and the liquidity of the Nifty index. The liquidity of the index is captured by computing the impact cost of the Nifty stocks. All the measures are computed half an hour before the end-of-day i.e. at 3 p.m. for the period July 2012 - March 2014. The plot indicates a weak positive relationship between the imprecision indicators and the liquidity of the index.




Figure 5 Intra-day imprecision measures vs. impact cost

The figure plots the linear relationship between the intra-day imprecision indicators and the liquidity of the Nifty index. The liquidity of the index is captured by computing the impact cost of the Nifty stocks. All the measures are computed at a fifteen seconds frequency for March 2014.

The plot indicates a negative relationship between the imprecision indicators and the liquidity of the index.




Table 6 Regression results

The table presents robust regression results for the model:

$$\text{Imp}_{it} = \alpha_0 + \alpha_1 \text{IC}_t + \epsilon_t$$

where Imp_{it} are the imprecision indicators: width of the confidence interval (Width of CI) and standard deviation of the bootstrap estimates (σ). IC is the impact cost of the Nifty index. It captures the liquidity of the index.

p-Value is reported in the parentheses below each coefficient.

	Daily		Intra-day	
	Width of CI	σ	Width of CI	σ
α_0	3.2 (0.05)	0.8 (0.04)	13.5 (0.00)	3.4 (0.00)
IC	126.1 (0.06)	29.7 (0.07)	-262.3 (0.00)	-69.4 (0.00)
Adj. R ²	0.01	0.01	0.03	0.03

8 Using this measure of imprecision for model selection

Weighting options by their vega to estimate a volatility index is one of the several approaches available to compute a volatility index. The alternatives include weighting options by volatility elasticity, liquidity etc [Grover and Thomas, 2012]. Traditionally, volatility indexes have been judged on the extent to which they forecast future realised volatility. However, another approach to model selection can be rooted in the problem of statistical imprecision. Just as model selection for LIBOR has been done by identifying the estimator which has the lowest bootstrap standard deviation, we could do model selection for the volatility index by choosing the index which has the lowest imprecision.

8.1 Methodology

We benchmark the performance of the four volatility indexes computed in Grover and Thomas [2012]. They are: vega weighted VIX (VVIX), liquidity weighted VIX (SVIX, TVVIX), and volatility elasticity weighted VIX (EVIX). The liquidity weighted VIXs incorporate the spread and traded volume weights

in the computation of SVIX and TVVIX respectively. The VIXs are computed four times a day, for the period February 2009 - September 2010. The performance indicators utilised are the standard deviation of the bootstrap estimates and the width of the confidence band. The performance of the volatility index estimators are compared pairwise by employing the Wilcoxon signed rank test. The test is based on the null hypothesis that the median difference between a pair of volatility indexes is zero.

We present the results for the pairwise comparisons of VIXs in the next section. The results are robust to the size of sampling frequency adopted, thus we present results only for the fifteen seconds frequency.

8.2 Pairwise comparisons of VIXs

Table 7 Measures of imprecision

Width of the confidence band (percentage points)				
	SVIX	TVVIX	VVIX	EVIX
No. of Obs	1516	1516	1516	1516
Min	0.93	1.36	1.03	2.18
1st Qu	2.71	2.74	2.27	6.20
Median	3.55	3.42	2.92	7.37
Mean	4.54	4.02	3.91	8.24
3rd Qu	4.84	4.44	4.06	9.26
Max	52.94	23.79	50.49	51.08
Std Dev	3.80	2.11	3.64	3.93
σ of the bootstrap estimates (percentage points)				
	SVIX	TVVIX	VVIX	EVIX
No. of Obs	1516	1516	1516	1516
Min	0.24	0.34	0.25	0.57
1st Qu	0.71	0.71	0.58	1.58
Median	0.91	0.88	0.74	1.87
Mean	1.14	1.03	0.94	2.05
3rd Qu	1.25	1.14	1.02	2.32
Max	13.39	4.77	11.78	10.58
Std Dev	0.82	0.53	0.75	0.88

Table 7 reports the summary statistics for the width of the 95% confidence interval for the four VIXs. The EVIX has the widest 95% band of 7.37 percentage points and largest σ of 1.87 percentage points in the median case. Thus, implying lowest precision among the volatility indexes. On the other

Figure 6 Width of the confidence band

The kernel density plot captures the distributions of the width of the confidence band for four VIXs. The 95% confidence band for a VIX is estimated estimated four times a day, for the period February 2009 - September 2010. For each point in time, the confidence band is computed from the bootstrapped sampling distribution of the VIX




Figure 7 Standard deviation of the bootstrap estimates

The kernel density plot captures the distributions of the standard deviation of the bootstrap estimates for four VIXs. The standard deviation for a VIX is computed four times a day, for the period February 2009 - September 2010. For each point in time, the standard deviation is calculated from the bootstrapped sampling distribution of the VIX.




Table 8 Wilcoxon signed rank test

	Width of the confidence band		σ of the bootstrap estimates	
	Median Diff	Pval	Median Diff	Pval
EVIX - SVIX	3.745	0.000	0.923	0.000
EVIX - TVVIX	3.846	0.000	0.962	0.000
EVIX - VVIX	4.326	0.000	1.097	0.000
SVIX - TVVIX	-0.004	1.000	0.013	0.341
SVIX - VVIX	0.641	0.000	0.183	0.000
TVVIX - VVIX	0.618	0.000	0.165	0.000

hand, VVIX has the highest precision with a confidence interval width of 2.92 percentage points and a σ of 0.74 percentage points in the median case.

Figure 6 represents the kernel density plot for the width of the 95% confidence interval for the four VIXs, estimated four times a day, for the period February 2009 - September 2010. The distributions of the VVIX, SVIX, and TVVIX are similar, while the EVIX is shifted heavily towards the right of the plot implying larger confidence bands than the rest. A similar pattern is seen for the distributions of the standard deviation of the bootstrap estimates across the VIXs (Figure 7).

Table 8 reports the median differences and the p-values for each pair of VIXs. The pairwise comparisons imply that VVIX outperforms the other three VIXs in terms of precision. It has the highest precision, whereas EVIX has the lowest precision. The difference in precision for the liquidity VIXs is statistically insignificant. Thus, ranking the VIXs based on their performance, we get: VVIX, SVIX and TVVIX, and EVIX.

9 Reproducible research

An R package named [ifrogs](#) has been released into the public domain, with an open source implementation of the methods of this paper. An example shown in this package replicates all the calculations of this paper for a sample of index options for two series: S&P 500 and Nifty.

10 Conclusion

Volatility indexes are being used in a large number of applications. In this paper, we suggest that there is a need for greater caution in their use. The volatility index is a location estimator constructed off noisy data; it is not a hard number. The imprecision of this estimator has an economically significant magnitude. The median value of the width of the 95% confidence interval for VVIX , which is the most precise volatility index, is 2.9 percentage points, which is a large number when compared with the one-day change in VVIX which has a median value of 1.18 percentage points.

For a sample of SPX options, we find that the confidence interval of VVIX estimated has a width of 1.5 percentage points which is an economically significant number when compared with the one-day change in VVIX which is smaller than this number for 62% of the days in our sample. This finding has important consequences for real-world applications of volatility indexes. As an example, exchanges such as CBOE may find it useful to disseminate a confidence interval for VIX instead of only showing a point estimator.

This paper also relates to the literature that develops empirical proxies to measure ambiguity. Uncertainty in VIX may be another way of capturing ambiguity. Imprecision indicators proposed in the paper can be computed to measure ambiguity for a wide range of stocks and at higher frequencies. This measure of ambiguity can further be used to test the relationship between ambiguity and expected stock returns. Traditional pricing factors discussed in the literature can also be included in the analysis to check whether uncertainty in VIX captures information beyond that found in these factors.⁷

The research of this paper was centred around volatility index estimation that is based on the Black-Scholes option pricing formula. Future extensions of this work are required in order to obtain inference procedures for model-free estimators such as the CBOE VIX .

⁷Similar analysis has been conducted in [Brenner and Izhakian \[2011\]](#), [Baltussen et al. \[2013\]](#), and [Ehsani et al. \[2013\]](#).

References

- T. G Andersen, O Bondarenko, and M. T. Gonzalez-Perez. A corridor fix for high-frequency vix. *Working Paper, Northwestern University*, 2011.
- M Arak and N. Mijid. The vix and vxm volatility measures: Fear gauges or forecasts? *Derivatives Use, Trading Regulation*, 12(1-2):1–2, 2006.
- Guido Baltussen, Sjoerd Van Bakkum, and Bart Van der Grient. Unknown unknowns: Vol-of-vol and the cross-section of stock returns. In *AFA 2013 San Diego Meetings Paper*, 2013.
- Geert Bekaert, Marie Hoerova, and Marco Lo Duca. Risk, uncertainty and monetary policy. *Journal of Monetary Economics*, 60(7):771–788, 2013.
- Jeremy Berkowitz. Dealer polling with noisy reporting of interest rates. *Journal of Fixed Income*, 9(1):47–54, 1999.
- B. J Blair, S. H Poon, and S. J. Taylor. Forecasting S&P 100 volatility: The incremental information content of implied volatilities and high-frequency index returns. *Journal of Econometrics*, 105:5–26, 2001.
- Nicholas Bloom. The impact of uncertainty shocks. *econometrica*, 77(3):623–685, 2009.
- Menachem Brenner and Yehuda Izhakian. Asset prices and ambiguity: Empirical evidence. *Stern School of Business, Finance Working Paper Series, FIN-11-10*, 2011.
- M Britten-Jones and A. Neuberger. Option prices, implied price processes, and stochastic volatility. *Journal of Finance*, 55:839–866, 2000.
- P Carr and D. Madan. Towards a theory of volatility trading. In *R. Jarrow (Ed.), Volatility: New Estimation Techniques for Pricing Derivatives*. London: Risk, 1998.
- D. P Chiras and S. Manaster. The information content of option prices and a test of market efficiency. *Journal of Financial Economics*, 6:213–234, 1978.
- A Cipollini and A. Manzini. Can the vix signal market’s direction? an asymmetric dynamic strategy. *An Asymmetric Dynamic Strategy (April 2007)*, 2007.

- John Cita and Donald Lien. Constructing accurate cash settlement indices: The role of index specifications. *Journal of Futures Markets*, 12(3):339–360, 1992.
- M. M Copeland and E. T. Copeland. Market timing: Style and size rotation using the vix. *Financial Analysts Journal*, 55(2):73–81, 1999.
- C. J Corrado and T. W., Jr. Miller. The forecast quality of CBOE implied volatility indexes. *Journal of Futures Markets*, 25:339–373, 2005.
- K Demeterfi, E Derman, M Kamal, and J. Zou. A guide to volatility and variance swaps. *Journal of Derivatives*, 6:9–32, 1999.
- B. Dupire. Model art. *Risk* 6, pages 118–124, 1993.
- B. Efron. Better bootstrap confidence intervals. *Journal of the American Statistical Association*, 82(397):171–185, 1987.
- Sina Ehsani, Timothy Krause, and Donald Lien. An empirical index of knightian uncertainty. 2013.
- P. Giot. Implied volatility indexes and daily value-at-risk models. *Journal of Derivatives*, 12:54–64, 2005.
- R Grover and S. Thomas. Liquidity considerations in estimating implied volatility. *Journal of Futures Markets*, 32(8):714–741, 2012.
- L. Hentschel. Errors in implied volatility estimation. *Journal of Financial and Quantitative Analysis*, 38:779–810, 2003.
- G. J Jiang and Y. S. Tian. The model-free implied volatility and its information content. *The Review of Financial Studies*, 18:1305–1342, 2005.
- G. J Jiang and Y. S. Tian. Extracting model-free volatility from option prices: An examination of the vix index. *Journal of Derivatives*, 14:1–26, 2007.
- H. A Latane and R. J., Jr. Rendleman. Standard deviations of stock price ratios implied in option prices. *The Journal of Finance*, 31:369–381, 1976.
- A. Neuberger. The log contract: A new instrument to hedge volatility. *Journal of Portfolio Management*, 2:74–80, 1994.
- Jun Pan and Allen M Poteshman. The information in option volume for future stock prices. *Review of Financial Studies*, 19(3):871–908, 2006.

- A. Shah. Improved methods for obtaining information from distributed dealer markets, 2000.
- R. E. Whaley. Derivatives on market volatility: Hedging tools long overdue. *Journal of Derivatives*, 1:71–84, 1993.