

VIX is not a hard number



VIX is not a hard number

- Grover and Shah have written an interesting paper with valuable insights.



VIX is not a hard number

- Grover and Shah have written an interesting paper with valuable insights.
- I agree completely with their conclusion that VIX is not a hard number. Probably, this can be generalized beyond VIX:
 - Wall Street tends to create derivatives on derivatives on derivatives to the n 'th order.
 - Modern financial econometrics tends to create estimates of estimates of estimates and treat them as data.
 - In both cases, the entire structure is often a bit shaky.

VIX is not a hard number

- Grover and Shah have written an interesting paper with valuable insights.
- I agree completely with their conclusion that VIX is not a hard number. Probably, this can be generalized beyond VIX:
 - Wall Street tends to create derivatives on derivatives on derivatives to the n 'th order.
 - Modern financial econometrics tends to create estimates of estimates of estimates and treat them as data.
 - In both cases, the entire structure is often a bit shaky.
- Morgenstern wrote a whole book about this half a century ago (Oskar Morgenstern (1950) *On the Accuracy of Economic Observations*. Princeton: Princeton University Press.) in which he quoted Norbert Wiener: “Economics is a one or two digit science”

Optimal Aggregation of Noisy Estimators



Optimal Aggregation of Noisy Estimators

- “We propose viewing each traded option as being akin to a noisy scientific instrument, that gives one imprecise reading of the forecasted implied volatility.” Grover and Shah (2014)



Optimal Aggregation of Noisy Estimators

- “We propose viewing each traded option as being akin to a noisy scientific instrument, that gives one imprecise reading of the forecasted implied volatility.” Grover and Shah (2014)
- The optimal solution to the aggregation problem is to weight each instrument in proportion to its precision (in inverse proportion to its variance).

Optimal Aggregation of Noisy Estimators

- “We propose viewing each traded option as being akin to a noisy scientific instrument, that gives one imprecise reading of the forecasted implied volatility.” Grover and Shah (2014)
- The optimal solution to the aggregation problem is to weight each instrument in proportion to its precision (in inverse proportion to its variance).
- If you have an informative prior, then you must weight with that also. If you are not a Bayesian, you will not go down this path.

Optimal Aggregation of Noisy Estimators

- “We propose viewing each traded option as being akin to a noisy scientific instrument, that gives one imprecise reading of the forecasted implied volatility.” Grover and Shah (2014)
- The optimal solution to the aggregation problem is to weight each instrument in proportion to its precision (in inverse proportion to its variance).
- If you have an informative prior, then you must weight with that also. If you are not a Bayesian, you will not go down this path.
- Grover and Shah (2014) implicitly use equal weighting of all instruments – they use the same weighting schemes that have been proposed for non noisy option prices.

Optimal Aggregation of Noisy IVs




Optimal Aggregation of Noisy IVs

- If the standard deviation of the noise in the option price is σ_ϵ , then the approximate standard deviation of the noise in the implied volatility can be obtained by dividing by the vega (\mathcal{V})

$$\sigma_{IV} \approx \frac{\sigma_\epsilon}{\mathcal{V}}$$

The approximation is in assuming that vega is constant over the range of pricing errors.



Optimal Aggregation of Noisy IVs

- If the standard deviation of the noise in the option price is σ_ϵ , then the approximate standard deviation of the noise in the implied volatility can be obtained by dividing by the vega (\mathcal{V})

$$\sigma_{IV} \approx \frac{\sigma_\epsilon}{\mathcal{V}}$$

The approximation is in assuming that vega is constant over the range of pricing errors.

- If all options have the same σ_ϵ , then the optimal aggregation would be to weight by the square of the vega (\mathcal{V}^2).

Optimal Aggregation of Noisy IVs

- If the standard deviation of the noise in the option price is σ_ϵ , then the approximate standard deviation of the noise in the implied volatility can be obtained by dividing by the vega (\mathcal{V})

$$\sigma_{IV} \approx \frac{\sigma_\epsilon}{\mathcal{V}}$$

The approximation is in assuming that vega is constant over the range of pricing errors.

- If all options have the same σ_ϵ , then the optimal aggregation would be to weight by the square of the vega (\mathcal{V}^2).
- Near money options that have high vega are also more liquid and probably have low price error as well and should be weighted even more than what was suggested above – maybe (\mathcal{V}^3).

Optimal Aggregation of Noisy IVs

- If the standard deviation of the noise in the option price is σ_ϵ , then the approximate standard deviation of the noise in the implied volatility can be obtained by dividing by the vega (\mathcal{V})

$$\sigma_{IV} \approx \frac{\sigma_\epsilon}{\mathcal{V}}$$

The approximation is in assuming that vega is constant over the range of pricing errors.

- If all options have the same σ_ϵ , then the optimal aggregation would be to weight by the square of the vega (\mathcal{V}^2).
- Near money options that have high vega are also more liquid and probably have low price error as well and should be weighted even more than what was suggested above – maybe (\mathcal{V}^3).
- The superiority of the VVIX suggests that vega weighting is closer to the optimal than other weighting schemes. But weighting with the square or even higher powers of vega may be even better.

Optimal Sampling for the Bootstrap



Optimal Sampling for the Bootstrap

- How do we justify the bootstrap standard error as a measure of the imprecision of the estimate?



Optimal Sampling for the Bootstrap

- How do we justify the bootstrap standard error as a measure of the imprecision of the estimate?
- The observed option prices (traded options) are a sample from a much larger universe of potential options that could have been traded. We might as well have observed a different sample which could have given us a different VIX estimate.

Optimal Sampling for the Bootstrap

- How do we justify the bootstrap standard error as a measure of the imprecision of the estimate?
- The observed option prices (traded options) are a sample from a much larger universe of potential options that could have been traded. We might as well have observed a different sample which could have given us a different VIX estimate.
- But then uniform sampling is not optimal. Some strikes are so liquid that they are highly likely to be in any plausible sample.

Optimal Sampling for the Bootstrap

- How do we justify the bootstrap standard error as a measure of the imprecision of the estimate?
- The observed option prices (traded options) are a sample from a much larger universe of potential options that could have been traded. We might as well have observed a different sample which could have given us a different VIX estimate.
- But then uniform sampling is not optimal. Some strikes are so liquid that they are highly likely to be in any plausible sample.
- Optimal sampling would be to sample with probabilities proportional to trading frequency or inversely proportional to inter-trade times.

Do we know nothing about VIX changes?



Do we know nothing about VIX changes?

- “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. This suggests that on 62% of the days, we know little about whether VVIX went up or down when compared with the previous day.” Grover and Shah, 2014



Do we know nothing about VIX changes?

- “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. This suggests that on 62% of the days, we know little about whether VVIX went up or down when compared with the previous day.” Grover and Shah, 2014
- This is not necessarily true! **Counterexample:** Ten options are observed on both days and in each of them, the IV on the second day is higher than on the first day by amounts ranging from 0.1% to 3% and the VIX is higher by 1% (which is well below the imprecision of 1.5%). Yet, by the binomial test, we can say with 99.9% confidence that VIX is higher on second day.

Do we know nothing about VIX changes?


- “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. This suggests that on 62% of the days, we know little about whether VVIX went up or down when compared with the previous day.” Grover and Shah, 2014
- This is not necessarily true! **Counterexample:** Ten options are observed on both days and in each of them, the IV on the second day is higher than on the first day by amounts ranging from 0.1% to 3% and the VIX is higher by 1% (which is well below the imprecision of 1.5%). Yet, by the binomial test, we can say with 99.9% confidence that VIX is higher on second day.
- Counterexample is contrived, but the key issue is one identified in Morgenstern’s book: the accuracy of the first difference of a noisy variable depends not only on the error variance but also (and more importantly!) on the autocorrelation of the noise.

Criteria for the best VIX estimator



Criteria for the best VIX estimator

- “Traditionally, volatility indexes have been judged on the extent to which they forecast future realised volatility. However, ... we could do model selection for the volatility index by choosing the index which has the lowest imprecision.” Grover and Shah, 2014



Criteria for the best VIX estimator

- “Traditionally, volatility indexes have been judged on the extent to which they forecast future realised volatility. However, ... we could do model selection for the volatility index by choosing the index which has the lowest imprecision.” Grover and Shah, 2014
- I agree strongly with the idea that ability to forecast future realised volatility is not a good criterion. VIX is an estimate more of risk aversion than of risk and more of tail risk than of variance. That is why it proves so useful in asset classes like currency that are far removed from equities.

Criteria for the best VIX estimator

- “Traditionally, volatility indexes have been judged on the extent to which they forecast future realised volatility. However, ... we could do model selection for the volatility index by choosing the index which has the lowest imprecision.” Grover and Shah, 2014
- I agree strongly with the idea that ability to forecast future realised volatility is not a good criterion. VIX is an estimate more of risk aversion than of risk and more of tail risk than of variance. That is why it proves so useful in asset classes like currency that are far removed from equities.
- It may be econometrically impractical, but to me the best VIX is the one that is most closely correlated with the first principal component of a basket of tail risk indicators.