

Forecasting the VIX to improve VIX-Derivatives Trading
Chrilly Donniger
Chief Scientist, Sibyl-Project
Sibyl-Working-Paper, April 2016

*It is best to read the weather forecast before praying for rain.
(Mark Twain).*

Abstract:

In [1] Konstantinidi et. al. state: “The question whether the dynamics of implied volatility indices can be predicted has received little attention”. The overall result of this and the papers quoted in [1] is: The VIX is too a very limited extend (R^2 is typically 0.01) predictable, but the effect is economically not significant. This paper confirms this finding if (and only if) the forecast horizon is limited to one day. But there is no practical need to do so. One can – and usually does – hold a VIX Future or Option several trading days. It is shown that a simple model has a highly significant predictive power over a longer time horizon. The forecasts improve realistic trading strategies.

The Model:

Konstantinidi et. al. investigate in [1] different models for forecasting several volatility indexes one day ahead. They call their first model forecasting with economic variables. The variables are the positive and negative returns of the underlying index (in case of VIX the SPX), the log-differences of the one-month interest rate, the EUR/USD exchange rate, the WTI, the 30 days historical volatility, the interest rate yield spread, the log-difference of the index-futures trading volume. All these factors are for the VIX not significant, the adjusted R^2 is 0.002. An AR(1) model does not perform significantly better. A VAR model with several volatility indexes works fine for the VCAC and the VSTOXX, but not for the VIX. The volatility of European indexes is strongly influenced by the US-market. But Granger-causality is a one way street. “When Wall Street sneezes, the rest of the world catches a cold”. An ARIMA(1,1,1) and an ARFIMA(1,d,1) model are also of very limited use. K. Ahoniemi reports in [2] somewhat better results for an augmented ARIMA(1,1,1) model. The error terms are modeled with GARCH(1,1). Other explanatory variables do not improve heir forecasts.

There is no practical need to restrict the forecast to one day. The one day convention is for trading purposes unusual. One either trades intraday or over a longer time horizon. It is well known that the VIX has a mean-reverting behavior. Mean-reversion is swamped in the short run by the high volatility of the index. But it should be possible to exploit mean-reversion in the long run. The best – and most practical – model I have found is:

$$VIX_{ret}(h) = a_0 + a_1 * VIX(t) + a_2 * VXV(t) + a_3 * IVTS(t) \quad (1)$$

$VIX_{ret}(h)$ is $\log(VIX(t+h)) - \log(VIX(t))$ where h is the forecast horizon in trade days.

$VIX(t)$ is the current VIX-value.

$VXV(t)$ is the 3-months volatility index.

$IVTS(t)$ is the implied-volatility-term-structure defined as $VIX(t)/VXV(t)$.


The model uses the current VIX level, VXV can be interpreted as a smoothed version of the VIX. The IVTS is a measure of the current term-structure.

(1) has for a one-day forecast the same low predictive power then the models presented above. I got slightly better results when adding the current positive and negative return of the SPX. But the effect

disappears for $h \geq 2$.

Graphic-1 shows the actual (red) and the forecasted (yellow) $VIX_{ret}(10)$ value in the last 5 years. The basic benchmark is the random-walk model. The best guess for VIX_{ret} is for this model zero. The ratio of the mean-square-error of the basic model and (1) is 0.925. Another benchmark is the current value of the VIX-Future. This ratio is 0.931. The VIX Future has with a ratio of 0.994 to the random-walk model practically no additional predictive power. The current VIX forecasts the 10 days ahead value as well. But it should be noted that there is usually no Future with a maturity of exactly 10 trading days available.

The forecasts are all out of sample. The parameters are re estimated every 21 trading days. A higher frequency does not improve the fit significantly. The model is also used for real-time trading. In this case the parameters are calculated daily at start up.




Graphic-1: $VIX_{ret}(10)$. Actual (red), forecast (yellow) from 2011-04-25 to 2016-04-25


Graphic-2 shows the behavior for a 21-trading days (1 month) forecast. The error ratio of model (1) is 0.897 for the random walk and 0.913 for the VIX-Futures value. The VIX-Future improves the random-walk by 0.981. It is usually possible to construct from the 1st and 2nd Future a synthetic Future with a mean-maturity of 21 trading days. But this improves the forecast quality of the VIX-Future only marginally. The Futures have a large positive bias aka risk premium. The performance of (1) can be slightly improved if one estimates the parameters with the Lasso or Ridge-Regression instead of plain OLS. Both methods contain a penalty for larger parameter values. The Lasso uses the absolute, the Ridge-Regression the squared parameter values. The weight lambda is in both cases set to 0.05.

Note: The estimation of the Lasso and Ridge-Regression parameters is according to [3] quite sophisticated. I implemented instead the simple differential-evolution heuristics. One uses the OLS values as the starting point of the heuristic. Any other heuristic like simulated annealing should do well (see [4]). There is no practical need to use the sophisticated methods.

The model works – in relation to the random-walk – best for a forecast horizon of 31 trading days. The error ratio is 0.860 for the random walk and 0.893 for the VIX Futures. The VIX Future improves the random-walk by 0.962. The VIX Future ratio increases also for longer horizons. The mean reversion effect is probably “sucked dry” at this forecast horizon.



Graphic-2: VIXret(21). Actual (red), forecast (yellow) from 2011-04-25 to 2016-04-25



Graphic-3: VIXret(31). Actual (red), forecast (yellow) from 2011-04-25 to 2016-04-25

One can replace in equation (1) the VIX and the VXV by their log values. This is (surprisingly) slightly worse than using the indexes directly.

An advantage of equation (1) is that the values are easily available at real-time. The CBOE calculates the VIX and VXV every 30 seconds. One has therefore – during regular trading hours – a timely estimate of the future VIX at hand.

Note: Some of the fully automated trading strategies in the Sibyl-project are based on the IVTS. The IVTS has spikes which can last up to 1 minute. The 3-month VXV options are less liquid. The VXV lags behind the VIX in turbulent markets. An IVTS signal must therefore be active for some time (e.g. 120 secs) before it actually triggers a trading action.

VIX-Futures basis trading:

Campanano & Simon proposed in [5] a simple VIX Futures strategy to exploit the positive bias. A refined version of this idea was published in [6].

The daily roll of a VIX-Future is defined as.

$$R(t) = (VXF(t) - VIX(t)) / TTS(t) \quad (2)$$

VXF is the VIX Futures Price.


TTS are the Trade-days Till Settle (expiry).

One enters a short VIX Future position if $R(t)$ is above a given threshold and sells the Futures back if the basis is either below a lower threshold or one is close to the expiry.

One can replace in (2) the current VIX value with the VIX forecast at expiry. At of this writing (2016-04-26) the VIX is at 13.91, the June 2016 Future VXM6 at 18.00. The VIX forecast for the VXM6 expiry at 2016-06-15 is 16.38. This is a typical pattern in a contango situation. The forecast is higher than the current VIX but below the Futures price. In backwardation it is usually (but not always) the other way round. The daily roll $R(t)$ is in the typical contango situation less.

Graphic-4 shows the performance of basis-trading with the plain VIX value (red) and the VIX forecast (yellow). The thresholds are set like in [7] to 0.07 and 0.03 for the plain value. If one replaces the VIX by the forecast the thresholds have to be reduced to 0.03 and 0.01.

One closes a position at least 20 trade-days before expiry. The maturity at entry must be between 30 and 93 trade-days. One selects the Future with the highest daily roll. This is usually the most nearest one (2nd or 3rd). Trading starts with an index value of 500.000\$. This is a convention from previous working papers. The volume is determined by $(\text{index} * 0.4) / \$$ -multiplier. One shorts initially 20 Futures. The quantity is increased proportionally with a higher index value. The trading costs are the typical VIX Futures spread of 0.05 or 50\$ per trade and Future. Graphic-4 shows the performance in the last 5 years. The strategy with the plain VIX has a P&L of 110.2% with a Sharpe-Ratio of 0.93 and a maximum relative drawdown of 18.2%. The forecast improves this to a P&L of 156.2%, a Sharpe-Ratio of 1.12 and a drawdown of 16.8%.




Graphic-4. Basis Trade with plain VIX (red) and VIX-Forecast (yellow) 2011-04-25 to 2016-04-25.

The position is always closed at least 20 trading days before expiry. The 1st Future is avoided. One

could argue that the forecast should not be until expiry but to this sell-back date. But in this case one needs additionally a model of the Futures term structure for a given VIX-value and maturity. I have developed in [7] such a model. But it makes things quite complicated and does not improve the trading behavior. Usually the daily roll is highest for the 1st Future. But the 1st Future has also a (much) higher beta (see [7]).

The yellow chart in Graphic-5 shows the effect of setting the minimum maturity to 5 trade days (this value is used by Campasano&Simon). The overall P&L is with 156.1% identical to the 20 days rule. But the Sharpe-Ratio is with 0.95 considerable lower and the max. relative drawdown with 22.1% significantly higher. The additional risk of entering and/or keeping the 1st Future does not pay off.



Graphic-5. Basis Trade with 20 (red) and 5 days min-Maturity (yellow) 2011-04-25 to 2016-04-25.

Monthly VIX -Calls Writing:

I analyzed in [7] also a strategy for writing monthly VIX Calls. The strategy uses the IVTS as an additional signal. A position is only entered if the IVTS is below a given threshold. It is well known that the premium is higher in relative quiet markets. If the IVTS is above a higher threshold the position is closed. Additionally there is a conventional Stop&Loss criterion: If the underlying moves a given factor above the strike the position is also closed. One can define also a maximum loss. But these two stop-signals have the same effect. There is one big question: What is the underlying of a VIX option? According the contract specification it is the VIX. But there are a lot of warnings in trading-blogs for using the VIX as the underlying.

The option greeks for VIX options shown by most brokers are wrong. Most options chains that brokers provide assume the VIX index is the underlying security for the options, in reality the appropriate volatility future contract is the underlying. [8].

The author advises to use the corresponding VIX Future instead.


The CBOE answers this question with:

The Price Relationship of VIX[®] Options to the VIX[®] Index and VIX[®]Futures

At times, VIX option contracts may appear to be mispriced relative to the VIX index. This phenomenon is the result of the anticipatory nature of forward implied volatility relative to the current level of the VIX index. Although VIX options settle to the VIX index at expiration, prior to expiration prices of VIX options are based on expectations of what the VIX index will be at expiration. This expectation of what the VIX index will be at expiration also appears in the pricing of VIX futures contracts. As a result of this market dynamic, traders are advised to consider the price of the corresponding VIX futures contract when making trading decisions regarding VIX options. [9]

The argument “*This expectation of what the VIX index will be at expiration also appears in the pricing of VIX futures contracts*” is not very convincing. The VIX Future is a biased and poor forecast of the VIX at expiry. One can argue that hedging can only done with the Future. But this argument is mute if one does not want to hedge at all. It was already shown that model (1) forecasts considerable better. So one could use this forecast as an underlying. The options can be mispriced too. According my calculations this does not happen in the strict logical sense. The Black-Scholes IV can always be calculated for the considered options. But the IV is for Calls sometimes quite high. One has to be aware that mispricings are an artifact which can not be directly traded.

The parameters of Call-writing strategy are as follows: The maturity at entry is between 12 and 42 trading days (but if everything runs fine one keeps the options till expiry). If an option is closed before expiry one enters another position at the first moment the entry conditions are met (one does not sync with the expiry cycle). The IVTS must be below 0.94. The delta of the options must be between 0.35 and 0.45. If there are several options in this delta-range one selects the one with the highest IV. The position is closed if the IVTS is above 1.03 or if the underlying is above $1.10 \cdot \text{strike}$ (Futures) or $1.05 \cdot \text{strike}$ (Forecast). This is the only difference between these two strategies. The trading costs are the usual spread of 0.10 (or 10\$) per option and trade. The quantity is $(\text{index} \cdot 0.02) / \-multiplier . The index starts as before with 500.000. One writes in this case 100 calls. Trading is done between 2010-03-01 and 2015-03-26. The time range is dictated by the available options data.



Graphic-6: Call Writing with Future (red) and Forecast (yellow) as underlying.

Graphic-6 shows the performance if one uses the Future (red) and the Forecast (yellow) as underlying. The final P&L for the Future based strategy is 76.9% with a Sharpe-Ratio of 1.11 and a maximum relative drawdown of 6.8%. With the Forecast as underlying the P&L is 97.6% with a Sharpe-Ratio of 1.15 and a drawdown of 9.4%. The Forecast based strategy is clearly more profitable but only slightly better in risk-adjusted terms. When the IVTS is below 0.94 the Futures are in contango and the forecast is lower than the Futures price. The strikes of the Calls are hence lower and the premium higher (this effect is somewhat mitigated by the higher IV). One could compensate this effect by choosing a larger Delta range for the Futures based strategy. But this improves neither the final P&L nor the risk-adjusted performance.

Conclusion:

Forecasting the VIX is clearly possible if one does not stick to the one-day ahead convention. There are a lot of possibilities to exploit the mean-reverting behavior of the VIX. Instead of (1) one could also use an AR(1) model with – iterated – daily forecasts. But the presented model behaves superior and is as easy to implement. It is convenient to use in real-time trading. There is also some evidence (but of course no definite proof) that the forecast boosts existing trading strategies.

References:

- [1] E. Konstantinidi., G. Skiadopoulos, E. Tzagkaraki: Can the Evolution of Implied Volatility be Forecasted? Evidence from European and U.S. Implied Volatility Indices. Draft from 18/12/2007
- [2] K. Ahoniemi: Modeling and Forecasting Implied Volatility: An Econometric Analysis of the VIX Index. Working Paper, Helsinki School of Economics, 2006.
- [3] T. Hastie, R. Tibshirani, J. Friedman: The Elements of Statistical Learning, 2nd Ed. Ch. 3.4
- [4] M. Gilli M., D. Maringer, E. Schumann: Numerical Methods and Optimization in Finance.
- [5] J. Campasano, D. Simon: The VIX Futures Basis: Evidence and Trading Strategies. June 27, 2012
- [6] Ch. Donninger: VIX Futures Basis Trading: The Calvados-Strategy, Rev. 1, Jan. 2014
- [7] Ch. Donninger: Modeling and Trading the VIX and VSTOXX with Futures, Options and the VXX.
- [8] V. Harwood: Thirteen Things You Should Know About Trading VIX options
sixfigureinvesting.com/2010/01/trading-vix-options/
- [9] CBOE: VIX Options – Price Relationship. www.cboe.com/strategies/vix/optionsintro/part4.aspx