Scenario-Generation with IVTS Block-Sampling

Chrilly Donninger
Chief Scientist, Sibyl-Project
Sibyl-Working-Paper, Oct. 2012
http://www.godotfinance.com/

Abstract:
A novel method for generating scenarios for portfolio optimization and risk models is presented. The method is simple, fast and flexible. The resulting data are closer to the stylized facts of financial time-series than comparable methods.

Introduction:

There are two broad methods for generating asset-scenarios. In the simplest case one assumes a brownian motion of the returns. One draws from a normal distribution with a given volatility/variance and adds a trend. The two parameters are estimated from a historic-time-window. To model fat-tails the normal distribution can be replaced by a Student-t (or a similar fat-tailed distribution). The distribution of most assets is skewed and the fat tails are more pronounced on the left/downside (in case of inverse or VIX-based products it is of course the other way round). One way to overcome this problem is to sample from the empirical distribution.
An important stylized fact of financial time-series is volatility-clustering. Both methods can not handle this effect. The generate implicitly or explicit scenarios with the mean-volatility of the historic-data (if one draws from a distribution, the volatility can also be set to an arbitrary value).
One way to overcome this problem are GARCH related schemes. One forecasts the volatility and scales the drawn sample accordingly. This can also be done with empiric sampling. The returns are before normalized by the current GARCH-volatility. The sample is then multiplied by the volatility forecast. This method is not without it's problems. There are dozens of GARCH models. Simple GARCH models do not fully capture the dynamics of the volatility process. For more complicated one's the parameter estimation is rather unstable. Another nasty problem is the correlation of volatility and trend. Low volatility regimes have usually a positive, high-volatility regimes a negative trend. In case of empirical sampling this relation is destroyed by the normalization process. One samples in a high volatility scenario a normalized return from a low volatility-regime. The normalized positive trend is magnified by the high rescaling factor. To overcome this problem, one has first to remove this effect. But it is not clear how to do this. The relation is non-linear and does not only depend on the volatility level. After a crash volatility is declining slowly from a high-level. The trend is in this phases – per definition – high. I am not aware of any appropriate model.

The Block-Sampling-Method:

Block-Sampling uses the empirical distribution. But instead of sampling day by day one can also sample a block of days. One selects e.g. 2010-05-05 as the starting date and uses the next k-days (e.g. till 2010-05-19) for generating a MC-path. If one wants to generate a trajectory 21 trading days ahead, one could use in the simplest case a block-size of 21.

Each path repeats a sequence in the past. Obviously one preserves in this way volatility-clustering. But there are 2 drawbacks. One has only a very limited bag of different paths available. A block starting at 2010-05-06 differs only in the first and last day from a block starting at 2010-05-05. Additionally a block sampled at 2008-10-05 would generate a very dramatic crash-scenario. If the market is currently quiet it is unrealistic that the returns behave in the next 21 trading days like in October 2008. Besides to reactions to natural disasters (e.g. the Japanese earthquake) crashes are not bolts from the blue. They build up in successive phases.
To solve the first problem one has to use smaller block-sizes. But this makes the second problem even worse. One introduces at the block-gaps huge volatility and regime jumps. A quite phase is followed by a crash-phase which is followed by a very quite phase … This can not happen in a realistic sequence.
Block sampling is also used in Bootstrap Methods (see [1]). There are two known solutions. Use either long blocks or calculate first a model and use only the model-residuals for the Bootstrap. If the model works reasonable, the residual have almost no structure like autocorrelation. One can use long blocks in the Bootstrap, because one wants to replicate the whole historic time-series. The purpose of the Bootstrap is to get from the same data a somewhat different but statistically equivalent time series. In this way one can estimate the variance of the time series parameters. This solution is – as described above – for a simulation of the next 21 trading days not feasible. Using a model is going back to square one. It's the purpose of the approach to avoid a model at all.

The novel approach is to divide the time-series into regimes. One samples a block only from the same regime, follows the past-history of the block and calculates the regime of the last block-trading-day. This is the new regime-state. One selects a block from the same regime …. E.g. one starts in a quite regime. If nothing has happened in the historic block, one selects another block from a quiet phase. If one had selected the block starting at 2010-05-03, one would repeat at 2010-05-06 the flash crash and would select the next block from a volatile regime. If in this new block the market calms down somewhat, one would select the next block from the middle-volatile basket …
The only remaining question is: How to define regimes? In a previous working paper (see [2]) I used the VIX. The practical experience with this method was mixed. It makes a big difference if VIX is going up from 15 to 20 or falling from 25 to 20. This are clearly two different regimes (GARCH scaled sampling has the same problem, see above). Hence besides three VIX levels there was the additional classification VIX-Up and VIX-Down. So in total there are 6 regimes. The problem with 6 regimes is that some of them contain only a few historic data. One gets the same problem than with very long blocks.

The Implied-Volatility-Term-Structure:

S&P defined in [3] the Dynamic VIX-Futures Index. The index is a weighted mean of the short-term (ETF VXX) and mid-term VIX-Futures index (ETF VXZ). The weight depends on the implied volatility term structure (IVTS).

```
IVTS(t) = VIX(t)/VXV(t).
```

The VXV is the implied volatility of S&P-options with 3 months maturity. Beside the different maturity the calculation is the same than for the VIX.
In a series of recent working-papers the IVTS was used to develop and improve trading-strategies (see [4],[5],[6]). It was shown that the IVTS contains information.
In the context of block-sampling the IVTS has several advantages over the VIX. There is no need to differentiate between a rising and falling VIX. This information is already in the term-structure. If there are signs for a crash, the VIX rises much faster than the VXV. In case of a recovery the VIX also falls faster then the VXV. The VIX is mean-reverting. But the mean is not constant over time. The IVTS is robust to the drift of the mean. If the overall level of the VIX is going up/down, the VXV goes up/down too. The only practical downside of IVTS is: The ^VXV time-series starts at finance.yahoo at 2010.08.13. This restricts seriously backtests.

Like for the VIX the IVTS-Block-Sample uses 3-regimes. The reasonable results were found for a threshold of **0.89** for the low-Regime and a threshold of **0.96** of the high-regime.
These values are approximately the IVTS median and the 8th quantile. Hence about half of the historic-days are in the low-, 30% in the mid- and 20% in the high-regime. The 0.89 threshold is also used in the trading-strategies of [4],[5],[6]. But these strategies use for the high-regime 1.0 as the threshold. This corresponds approx. to the highest quantile. For the purpose of block-sampling a somewhat larger high-volatility regime seems to be preferable.

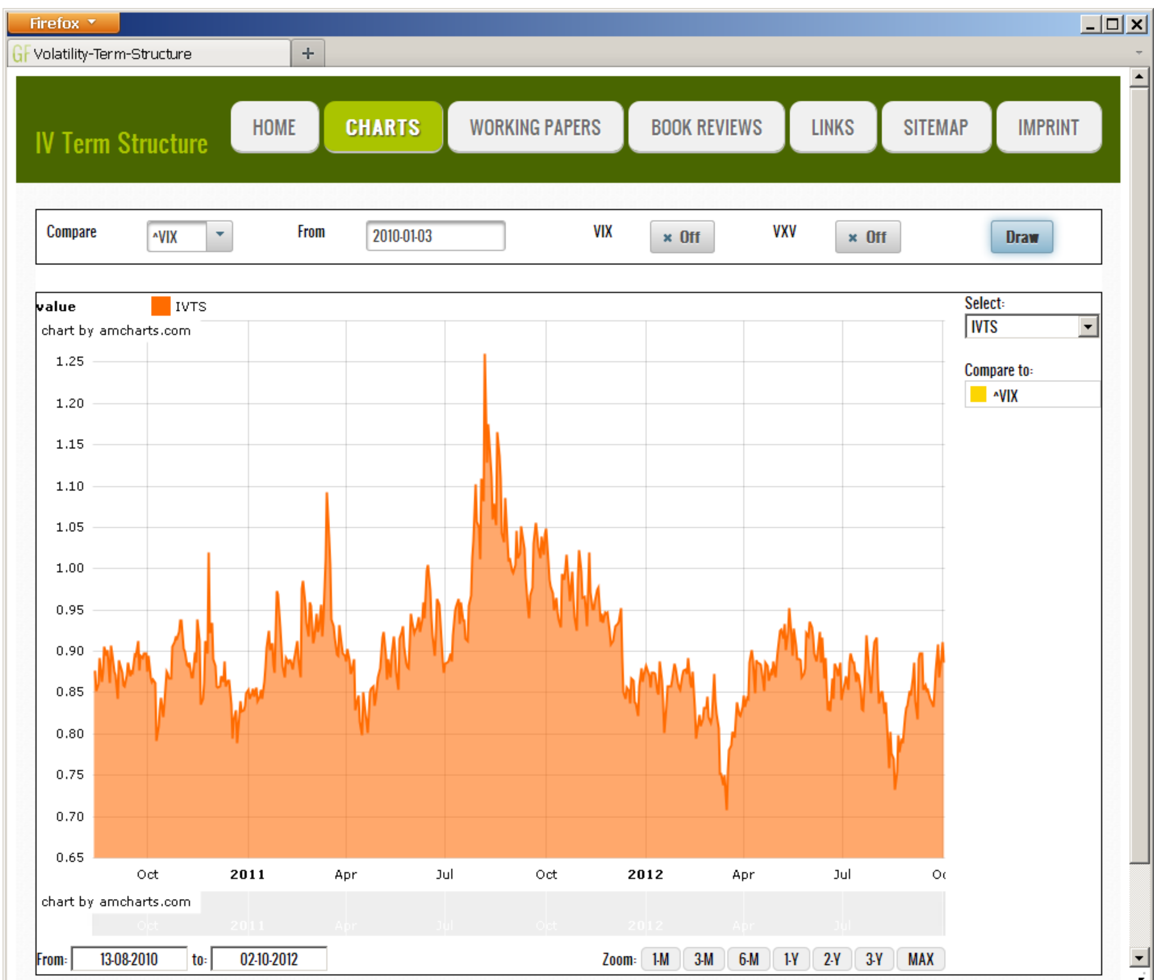

Graphic-1: IVTS from 2010.08.13 to 2012.10.02

Graphic-2: Performance IVTS (orange) to VIX (yellow) from 2010.08.13 to 2012.10.02

The second important parameter is the window length for sampling. The window is restricted by the availability of the VXV data on yahoo-finance (the VXV existed long before 2010.08.13, but I have not found a source for older data).
The short IVTS-time-series is for current (Oct. 2012) simulation no restriction. Financial time-series are not stable. It does not make sense to sample (very) old data. According previous results with VIX based regimes is a window-length of 1.5 years (378 trading-days) a reasonable choice. The short IVTS-time-series restricts with this settings systematic backtests.

The third parameter is the Block-Sample length. In the previous implementation the block length was selected randomly in a range of 4 to 8 trading-days. This should eliminate/wash-out short term return-patterns. In the current implementation the block-length is set fixed to 5. This setting handles all weekly-patterns. I could not detect any significant difference between the variable and fixed setting. So the simpler and somewhat faster fixed-length approach is used.

Block-Sampling for Portfolio-Optimization:

The Block-Sampling Method was initially developed for an Options-Model. But this model was - due to administrative reasons - so far not used practically. Instead (VIX based) block-sampling was used with good results in constructing an ETF-Portfolio (see [7]).
There are two ways to generate with block-sampling a portfolio-scenario.
A) One generates for each asset in isolation a scenario. These independent scenarios are then multiplied by the Cholesky of the correlation matrix.
B) One does not generate in the block-sampling-step return-scenarios. Along a sample-path only the trading-days are sampled (and stored in a matrix of Julian dates). The asset-returns are then all sampled on each path at the same block-dates. This preserves the assets-correlations.
The final results of A) and B) differ nevertheless. The correlation-matrix is usually generated by the Risk-Metrics method (see [7]). The correlation is exponentially weighted with an alpha of 0.97. The result reflects mainly the correlation of the last quarter. Approach B) uses in contrast the correlation of the full window-length of 1.5 years. The weight is not time but regime based. The first block is from the same regime than the current one. The regime can be different at the end of the block and hence the following blocks can be from different regimes. But especially for short paths historic periods with the same regime than the current one get a significant higher weight. This is after all the purpose of the regime-based block-sample. It is well known that the correlation depends on the regime. In times of trouble correlation increases. Approach B) models this behavior.

The basic assumption and alternatives:

The IVTS period-classification assumes that the S&P-500 defines for all considered assets the regime. This is certainly true for US-stocks. It is also reasonable to assume that international stock-markets follow Wall Street ("When Wall Street sneezes the rest of the world catches a flue"). But it is less clear for assets like bonds or currencies which have a low correlation with the S&P-500.
Another approach is to use for each asset it's own realized-volatility. The downside of this method is the problem of measuring volatility. If one uses daily data, one needs a time-window of about 1 month. This introduces a lag in the volatility-measure. A sudden increase is only "seen" after several trading days.
One can measure realized-volatility more accurately with high-frequency data. The access to this data is considerable slower than for daily-data. For portfolio-optimization purposes it does not matter, because the optimization of a 100+ assets portfolio takes much longer than the retrieval and volatility-calculation of HF-data. If one uses the scenario for option-pricing, the time factor becomes an issue.
But using HF-Data is not without hassle. It is not clear how one should handle the overnight return. The overnight return is significant for the overall performance of an asset. But the effect is not persistent. A large overnight jump today does not mean that there will be large jumps in the future. So for volatility modeling and forecasting it is preferable to ignore the overnight-return at all. The VIX reacts on an overnight jump. But it usually reverts during the day if there is no further large movement.
There are also effects due to low volume or special trading-days/periods. In a previous approach the median daily HF-volatility of time t-1, t, t+1 was used. Statistically speaking salt and pepper was filtered away. This works reasonable for single daily-spikes, but if fails to address e.g. the low realized-volatility in the last week of December. A VIX-based classification is less sensible to such effects. The VIX is the expected mean-variance over the next month (plus a risk-premium). Traders know, that the calm last December-days are usually followed by a relative stormy January. They adjust the implied volatility accordingly.

Conclusion:
The presented method is simple and fast. Although detailed backtest results could not be presented, it seems to be an attractive methodology for generating asset-scenarios. An especially attractive feature is the joint scenario-generation for portfolios.

References:
[1] A.Davison, D.Hinkley, Bootstrap Methods and their Application, chap. 8, Complex Dependence.
[2] Ch. Donninger: Using Monte-Carlo Basket-Block-Sampling for Options-Evaluation and Risk-Models, Sibyl-Working-Paper, Feb. 2011
[3] S&P: S&P 500 Dynamic VIX Futures Index Methodology, Aug. 2011
[4] Ch. Donninger: Improving the S&P Veqtor-Index: The Daiquiri Strategy, Sibyl-Working-Paper, Revision-1, July 2012
[5] Ch. Donninger: Improving the S&P Dynamic VIX-Futures Strategy: The Mojito Strategy, Sibyl-Working-Paper, Revision-1, July 2012
[6] Ch. Donninger: Trading with Low-Volatility ETF's: The Irish-Cream-Strategy, Sibyl-Working-Paper, Aug. 2012
[7] Ch. Donninger: ETF-Portfolio-Construction: The Camel-Strategy, Sibyl-Working-Paper, Revision-1, July 2012