

Predicting financial crashes

Financial crashes are a paradigm example of extreme events. They are typically highly correlated with spikes in volatility, the standard measure of financial risk. Here I will focus on the implied volatility index VIX for the US equity index S&P500 and I will consider as financial extreme events all spikes of this index above the value 30. The question I address here is if such VIX spikes can be predicted and, if yes, with which accuracy.

Background

It is by now well established that extreme events in financial processes have fat tail distributions, i.e. their occurrences are much more frequent than what implied by the typical Gaussian distributions used in modeling. There are two schools of thought about these extreme events. The first (see e.g. N. Taleb, “Fooled by Randomness”) considers them as purely and entirely random events, by now generally called black swans. The second school of thought (see e.g. D. Sornette, “Dragon Kings, Black Swans and the Prediction of Crises”), instead, calls them dragon kings and posits that they arise due to a complex and chaotic, but fundamentally deterministic non-linear dynamics and that, as a consequence they display a certain degree of predictability. It should be clear that identifying any degree of predictability in financial extreme events can lead to potentially very large gain opportunities.

Extreme events appear not only in finance but in many other realms of nature in which a large number of constituent elementary agents conspires to create a complex dynamics of the whole. One particularly spectacular realization is that of rogue (or freak) waves in the seas and in non-linear optics. The idea of my approach is to consider volatility spikes as financial rogue waves and to use methods of non-linear physics and quantum mechanics to identify typical randomness patterns that can predict their appearance.

Results

The results are shown in the two figures below in which VIX data for the last 20 years are shown together with the predictive signals generated by my method. There are two types of signals: level 1 signals constitute a “yellow” alarm for approaching danger while level 2 signals are “red” alarms confirming the approach of a financial rogue wave.



Out of 11 events in which VIX spiked over 30, 10 are clearly predicted (the notable exception being the spike of '10), giving a recall measure of 90%. Of the 12 predicted events, 10 are evidently relevant ones, giving a precision measure of 83%. Together, this amounts to a F1 accuracy of 86%. Note however that the only two false positives are “yellow” alarms, all level 2 alarms are relevant. In general the prediction times are typically within a couple of months or less. The only notable exception is the spike of '00, which is anyhow a much more protracted event with an unclear “beginning”.

