

Financial Price Volatility Calibration

Financial data are known to be far from normal and replete with outliers, many of which result from “dirty” data – data that contain errors. Such errors introduce extreme or aberrant data points, which can significantly distort parameter estimation results. Therefore, robust estimations are required to achieve stable and accurate results. As a matter of fact, robust estimation will be particularly applicable for financial data, which often features the three situations we are protecting against: occasional rogue values (outliers), many small errors and underlying non-normality.

Robust estimation is an estimation technique which is insensitive to a small number of large departures from model assumptions, such as outliers. Robust means that changing a few samples in the data, even by a large amount, does not cause a large change in estimates.

Let X_t be the set of clean returns for a non-proxy risk factor. Certain risk factors exhibit periods of stale data (sequences of 0 return) followed by a non-zero return. This is indicative of a sample distribution with a large peak at 0. To smooth out the return series, the return following a series of zero returns is spread over the previous returns with a sequence of positive and negative values.

If a return series has a sequence of 0's at the end of the return series, these 0's are replaced with NaN (not a number). Let Y_t be the adjusted ('processed') return series of X_t . If the return series X_t is good (no 0's), $Y_t = X_t$. NaN data in Y_t is then filtered out when computing the mean and volatility. The mean and volatility for each non-proxy risk factor is computed using the adjusted return series Y_t (not X_t) and stored.

If r_n is the first return after a series of N stale or missing data points, the missing data is filled-in as follows:

- First $\text{round}(\sqrt{N})$ points replaced with $\frac{r_n}{\sqrt{N}}$
- Remaining points with $\frac{-r_n}{\sqrt{N}}$

The robust estimate of volatility methodology is described. The distribution is estimated as a t-distribution with degrees of freedom parameter $\nu = 4.5$. The value of ν was selected by calculating the estimation error in recovering a known volatility from simulated data from a t-distribution. The estimation error was shown to exhibit a small spread across all input distributions with $\nu > 2$ for values of the estimation parameter $4 \leq \nu \leq 5$.

Consider a stochastic representation of a t-distributed random variable,

$$X_t = \mu + \sigma \frac{Z_t}{S_t}; t = 1, 2, \dots, n$$

where Z_t are independent standard normal variables that are also independent of S_t , and S_t are independent $\chi^2_{\nu}/(\nu-2)$ random variables. We can determine an iterative scheme for calculating the mean and variance of the distribution. The following set of equations is solved iteratively, converging to $\hat{\sigma}$, $\hat{\mu}$ for sufficiently large k ,

$$\hat{\sigma}_{k+1}^2 = \frac{1}{n} \sum_{t=1}^n w_t^k (X_t - \hat{\mu}_k)^2$$

$$\hat{\mu}_{k+1} = \frac{\sum_{t=1}^n w_t^k X_t}{\sum_{t=1}^n w_t^k}$$

$$w_t^{k+1} = \frac{\nu + 1}{\nu - 2} \left(1 + \frac{(X_t - \hat{\mu}_{k+1})^2}{(\nu - 2) \hat{\sigma}_{k+1}^2} \right)^{-1}$$

where v is the degree of freedom and is set to 4.5

The original estimate does not affect the final result of the iterations, but only the number of iterations (rate of convergence). The current version uses the sample median and sample standard deviation as the starting point for the mean and volatility respectively. The iterations are repeated until the following convergence criterion is met:

$$\left| \frac{\hat{\sigma}_{k+1}^2 - \hat{\sigma}_k^2}{\hat{\sigma}_k^2} \right| \leq 1e^{-5}$$

If the estimate does not converge within 10,000 iterations, a warning is issued. If the variance converges but is less than 1×10^{-12} , then a value of zero is assigned.

The standard form of MLE places the same weight on all data points. One can use a weighted MLE with weights y_i to place more or less weight on different parts of the data series. In particular for time series, an exponential weighting can be applied so that more recent observations receive greater weight. This makes the volatility estimate more responsive to recent market conditions (see <https://finpricing.com/lib/EqCallable.html>).

We can determine an iterative scheme for calculating the mean and variance of the distribution. The following set of equations is solved iteratively, converging to $\hat{\sigma}$, $\hat{\mu}$ for sufficiently large k :

$$\hat{\sigma}_{k+1}^2 = \frac{1}{\sum_{t=1}^n y_t} \sum_{t=1}^n w_t^k (X_t - \hat{\mu}_k)^2$$

$$\hat{\mu}_{k+1} = \frac{\sum_{t=1}^n w_t^k X_t}{\sum_{t=1}^n w_t^k}$$

$$w_t^{k+1} = y_t \frac{\nu+1}{\nu-2} \left(1 + \frac{(X_t - \hat{\mu}_{k+1})^2}{(\nu-2)\hat{\sigma}_{k+1}^2} \right)^{-1}$$

where $y_t = \frac{1-\lambda}{\lambda} \lambda^t$ with $t \in (1, n)$

The lambda value is set at 0.969 to provide a 3 month effective window (i.e. the area under a uniform 3 month window is equivalent to the area under an exponential decay curve with lambda equal to 0.969). The original estimate does not affect the final result of the iterations, but only the number of iterations (rate of convergence). The current version uses the sample median and sample standard deviation as the starting point for the mean and volatility respectively. The iterations are repeated until the following convergence criteria is met:

$$\left| \frac{\hat{\sigma}_{k+1}^2 - \hat{\sigma}_k^2}{\hat{\sigma}_k^2} \right| \leq 1e^{-5}.$$

If the estimate does not converge within 10,000 iterations, a warning is issued. If the variance converges but is less than 1e-12, then a value of zero is assigned.

The capped volatility is defined as:

$$\sigma = \min \left\{ \sigma_{AVG} * cap, \max \left\{ \sigma_{AVG}, \sigma_{EXP(3mo)} \right\} \right\}$$

Where σ_{avg} represents the average volatility based on the robust t estimate with uniform weighting,

$\sigma_{exp(3mo)}$ represents the exponentially weighted volatility based on the robust t-estimate with weighting to produce a 3 month effective window, $cap = 1.25$. NOTE: in the future it may be desirable to have a different cap value for each asset class (or curve).

This produces 3 regimes depending on the calculated value of the average and exponentially weighted volatility:

$$\sigma = \begin{cases} \sigma_{AVG}, & \sigma_{EXP(3mo)} \leq \sigma_{AVG} \\ \sigma_{EXP(3mo)}, & \sigma_{AVG} < \sigma_{EXP(3mo)} \leq cap * \sigma_{AVG} \\ cap * \sigma_{AVG}, & cap * \sigma_{AVG} < \sigma_{EXP(3mo)} \end{cases}$$

The advantages of using the capped volatility are as follows:

1. Reacts as fast as exponentially weighted volatility for periods of increasing volatility above average volatility
2. Has a stable peak – behaves like average volatility when exponentially weighted volatility is much larger than average volatility
3. Provides buffer above average volatility for periods of higher volatility (buffer is a parameter)