Counterparty Credit Risk Backtest

Backtesting is the comparison of forecasts to realised outcomes. This comparison is either the comparison of a distribution with a single realised value at a point in time, as for market risk factor

Backtesting is the comparison of forecasts to realised outcomes. This comparison is either the comparison of a distribution with a single realised value at a point in time, as for market risk factor or exposure distribution backtesting, or the comparison of a single predicted value against some realised value at a point in time, as for backtesting EPE or pricing models.

VaR backtesting is a particular example of the former comparison of testing forecast distributions against realised outcomes. This paper argues that the VaR approach is inappropriate for backtesting the internal models used for counterparty credit risk calculations and suggests approaches that are more suitable.

The methodologies we discuss are based on comparison between the internal model forecasted probability distribution of exposure at various time horizons (calculated for representative counterparty portfolios) and the actual exposures that would have occurred on each portfolio at each time horizon by using in computation historical data on movements in market risk factors.

A forecast distribution of market risk factors or exposures has a number of properties. Forecasts are initialised at a particular point in time. The initialisation point is the date and time that a forecast is launched or issued.

Each forecast distribution has a time horizon, the time between initialisation and the realisation of the forecast. A forecast initialised on January 1st that realises on January 15th has a 14 day time horizon, a two week forecast. Note that forecasts with different time horizons can have the same initialisation date, ie two week and four week forecasts that realise on 15th and 29th January respectively would both have been initialised on the same date, 1st January.

Backtesting is a statistical test with the significance of any result depending on the amount of data used. A backtesting data set is a set of forecasts and the corresponding realisations of those forecasts, ie what actually occurred. This backtesting data set can be put together in a number of ways.

Backtesting using data from a single counterparty over a short period of time may not produce a meaningful conclusion about the quality of the EPE models and its sub-components used to generate that exposure. Firms with advanced model permission have addressed the data requirement problem by aggregating backtesting data across a number of dimensions. The possible dimensions are discussed below.

All choices over how pricing and simulation are performed should be inherited from the job spec, using the appropriate bi-temporal version of the job spec as of the anchor timestamp. It should not be bi-temporal in nature as a specific job should only be run once.

Under normal operation, many jobs will be created each day. The anchor timestamp will be constant across many jobs, as market data will be held constant throughout the day, only updating periodically (e.g. at end of day). The trade timestamp may be different every job as new trades will need to be picked up with each job to satisfy intra-day requirements.

The compute framework reads the info contained in a job and triggers all required simulation and pricing tasks as specified in the job. The compute framework has no business logic embedded; it is completely metadata driven and simply performs the tasks specified in the job and stores the output for use in aggregation.

A underlying market factor (UMF) represents a market factor taken as a whole, for example the USD libor curve or a volatility surface. The economic properties of a UMF are a system-wide definition – for example, the fact that USD libor is an interest rate is not a configuration choice but a fact. However, each job spec may specify a different way to model that particular UMF through specification of the calibration set.

A market data path represents a possible evolution of market data through time. Generally, all paths start at the same place with the real world market data, but evolve differently to each other over time. Future market data points on a path may be generated either through a simulation model (Monte Carlo paths), through application of pre-specified 'shocks' to each market data point, or may be real world values if the path is being generated retrospectively (e.g. for back testing).

Reference:

https://finpricing.com/lib/IrCurveIntroduction.html