A Neural Network Approach to Credit Risk and xVA

Even the most simple derivative contracts are today priced using complex Valuation Adjustment (xVA) methods and large scale Monte Carlo simulations. We present a proof of concept artificial neural network model trained on generated realistic data, capable of reproducing credit exposure profiles from market and trade data. The model errors were comparable to those expected from a Monte Carlo simulation with 5K paths. With regards to computational efficiency, the proposed model showed great potential in outperforming traditional numerical methods. The model could thus potentially be used as a shortcut in quick xVA front office pricing as it generates an accurate estimate of a simulated exposure profile in an instant.

The global financial crisis 2007-2009 exposed critical flaws in the valuation methods used in the derivatives markets. The default of many large derivative counterparties during the financial crisis led to an increased demand for counterparty credit risk assessment in the valuation of derivatives contracts. The framework for Credit Valuation Adjustments (CVA) was rapidly developed and increased awareness and regulation has since then led to several valuation adjustments.

An important measure in management of counterparty credit risk is the Expected Exposure (EE). The expected exposure of a trade defines the expected loss that one party would face in the event of the other's default. An exposure profile over the lifetime of a contract may be used to gain information about the xVA and credit risk associated with a trade.

With regards to computational efficiency, neural networks demonstrate significant advantages over Monte Carlo methods. Our model generates predictions in less than 1 ms, which is significantly more efficient than the corresponding Monte Carlo simulation used in this example. Furthermore, the computational efficiency of the Monte Carlo method is greatly dependent on the number of performed simulations. The possible benefits of a machine learning approach may thus be even greater than demonstrated in this example. Another advantage of the model was the rather short offline training time. The specific model was trained on a data set of 200K samples within a workday. It is fair to assume that more data and longer training would further increase the performance of the network.

To be useful in an industry context, the machine learning model may be applied to observed market data and exposure data generated within the bank. Reasonable standards for data gathering and model retraining should be established. The computational costs of offline training increase with the amount of data and retraining frequency. The performance of the model would presumably also increase. The trade-off between performance and costs needs further assessment.

In its current state, the proposed model has its greatest advantages in quick single-contract exposure evaluations that could be used in front office xVA solutions. With further development, the proposed architecture may prove especially useful for contracts with high-rated counterparties, traded in a normal and liquid market. The model could be used as an efficient pricing tool and complement to existing pricing engines. If a client requests xVA pricing for a single trade, the model could give a quick indication of the price level. The area of application could be expanded by extending the model, for example by training also on quantiles of the exposure.

Isabelle Frodé & Viktor Sambergs