APPENDIX 10A: Exposure and swaption analogy.

Sorensen and Bollier (1994), effectively calculate the CVA of a swap position and show this can be written as:

$$CVA_{swap} = LGD \sum_{i=1}^{n} V_{swaption}(t; t_i, T)\, PD(t_{i-1}, t_i).$$

Whilst more details on CVA computation are given in Chapter 12 and Appendix 12, we can note that in the above formula the expected exposure (EE) is represented by $V_{swaption}(t; t_i, T)$ which is the value today of a European swaption on the underlying swap. The swaption has an exercise date of t_i, which is the potential default time in the discretised CVA formula and the underlying swap is defined by the period $[t_i, T]$.[1]

The above shows that the EE for the purpose of calculating the CVA of the swap can be represented as a series of swaption values. The intuition is that the counterparty has the "option" to default at any point in the future and therefore effectively cancel the swap. Not only is this formula useful for analytical calculations, it is also quite intuitive for explaining CVA.

An interest rate swaption can be priced in a modified Black-Scholes framework via the formula:

$$[F\Phi(d_1) - X\Phi(d_2)]D(t^*, T) \qquad \text{(payer swaption)}$$

$$[-F\Phi(-d_1) + X\Phi(-d_2)]D(t^*, T) \qquad \text{(receiver swaption)}$$

$$d_1 = \frac{\ln\left(\frac{F}{X}\right) + 0.5\sigma_S^2(t^* - t)}{\sigma_S\sqrt{t^* - t}} = d_2 + \sigma_S\sqrt{t^* - t}$$

Where F is the forward rate of the swap, X is the strike (the fixed swap of the underlying swap), σ_S is the swap rate volatility, t^* is the maturity of the swaption (the time horizon of interest). The function $D(t^*, T)$ represents the underlying swap duration (annity value) for which the maturity is $(T - t^*)$. The exposure of the swap will be defined by the interaction between two factors: the swaption payoff, e.g. $F\Phi(d_1) - X\Phi(d_2)$, and the duration. These quantities respectively increase and decrease monotonically over time. The overall swaption value therefore peaks somewhere in-between as illustrated in Figure 10.1 and Spreadsheet 10.1.

[1] There could be a question of whether the swaption is cash or physically settled. This relates to the close out discussion in Chapter 14.

APPENDIX 10B: Comments on exposure models by asset class

i) Interest rates

In the Hull and White model, the "short rate" (short-term interest rate) is assumed to follow the following process:

$$dr_t = [\theta(t) - ar_t] + \sigma_r dW_t,$$

which is a Brownian motion with mean reversion. Mean reversion dictates that when the rate is above some "mean" level, it is pulled back towards that level with a certain force according to the size of parameter a. The mean reversion level $\theta(t)$ is time-dependent which is what allows this model to be fitted to the initial yield curve. The mean reversion has the effect of damping the standard deviation of discount factors, $B(t, T)$, written as:

$$\sigma(B(t,T)) = \sigma_r \left[\frac{1 - \exp\big(-a(T-t)\big)}{a(T-t)}\right].$$

Although the yield curve is not modelling directly, it can be "reconstructed" at any point, given knowledge of the above parameters and the current short rate. Hence, using such an approach in a Monte Carlo simulation is relatively straightforward. Using historical data (real world calibration), one can estimate the standard deviation (volatility) of zero-coupon bond prices of various maturities, then it is possible to estimate values for σ_r and a. For risk-neutral pricing, these parameters would be calibrated to the prices of interest-rate swaptions with a time dependent volatility function $\sigma_r(.)$. Spreadsheet 10.2 uses this model for a flat yield curve (Vasicek model).

ii) FX

FX is typically modelled via a geometric Brownian motion (GBM):

$$\frac{dX_t}{X_t} = k[\theta - \ln(X_t)] + \sigma_{FX}(t)dW_t,$$

where k is the rate of mean reversion to a long-term mean level θ. For real world calibration then the volatility function will be flat and mean reversion may be relevant to avoid the FX rate "exploding" which is particularly important for long time horizons. For risk-neutral calibration then the mean reversion is less important since the volatility function, $\sigma_{FX}(t)$, will be calibrated directly. It is generally necessary to make assumptions about long-dated volatility (e.g. above 5-years) in this respect.

iii) Equity

A standard geometric Brownian motion (GBM) is defined by:

$$\frac{dS_t}{S_t} = \mu(t)dt + \sigma_E(t)dW_t.$$

where S_t represents the value of the equity in question at time t, $\mu(t)$ is the drift, $\sigma_E(t)$ is the volatility and dW_t is a standard Brownian motion. It is generally preferable to simulate single stocks via their relationship (beta) to indices.

iv) Commodities

A simple and popular model (see Geman 2005) is:

$$\ln(S_t) = f_t + Z_t.$$

$$dZ_t = [\alpha - \beta Z_t]dt + \sigma_C(t)dW_t.$$

Where f is a deterministic function, which may be expressed using sin or cos trigonometry functions to give the relevant periodicity and the parameters α and β are mean reversion parameters.

v) Credit

A typical model for credit including jumps could be:

$$d\lambda_t = \theta[\eta - \lambda_t]dt + \sigma_\lambda\sqrt{\lambda_t}dW_t + jdN,$$

where λ_t is the intensity (or hazard rate) of default[2] and θ and η are mean reversion parameters. This model (depending on the calibrated parameters) can prevent negative hazard rates as required. Additionally, dN represents a Poisson jump with jump size j. This jump size can itself be random such as following an exponential distribution.

[2] This means that the default probability in a period dt conditional on no default before time t is $\lambda_t dt$. The hazard rate is related to the credit spread as is explained in Chapter 12.

APPENDIX 10C: Computation of incremental and marginal exposure

i) Incremental exposure

In order to calculate the incremental exposure we simply need to add the simulated values for a new trade (i) to those for the rest of the netting set. Working from equation (10.1), we can write

$$V_{NS+i}(s,t) = \sum_{k=1}^{K} V(k,s,t) + V(i,s,t) = V_{NS}(s,t) + V(i,s,t),$$

giving the future value of the netting set, including the new trade in each simulation (s) and at each time point (t). From this, it is easy to calculate the new EE, which can be compared with the existing EE as required by equation (10.2). What is helpful here is that we need only know the future value of the netting set, not the constituent trades. From a systems point of view this reduces the storage requirements from a cube of dimension $K \times S \times T$ (which could be extremely costly) to a matrix of dimension $S \times T$.

Typically, systems handle the computation of incremental exposure by calculating and storing the netting set information $V_{NS}(s,t)$ (often in an overnight batch) and then generating the simulations for a new trade, $V(i,s,t)$ "on-the-fly" as and when required. The "reaggregation" is straightforward and recalculation of measures such as EE is then a quick calculation.

ii) Marginal exposure

Suppose we have calculated a netted exposure for a set of trades under a single netting agreement. We would like to be able write the total EE as a linear combination of EEs for each trade, i.e.:

$$EE_{total} = \sum_{i=1}^{n} EE_i^*.$$

If there is no netting then we know that the total EE will indeed be the sum of the individual components and hence the marginal EE will equal the EE $(EE_i^* = EE_i)$. However, since the benefit of netting is to reduce the overall EE, we expect in the event of netting that $EE_i^* < EE_i$. The aim is to find allocations of EE that reflect a trade's contribution to the overall risk and sum up to the total counterparty level EE (EE_{total}).

As described in Chapter 10, this type of problem has been studied for other metrics such as value-at-risk (VAR). In the absence of a collateral agreement, EE (like VAR) is homogenous of degree one which means that scaling the size of the underlying positions by a constant will have the same impact of the EE. This is written as:

$$\boldsymbol{\alpha} EE(x) = EE(\boldsymbol{\alpha} x),$$

where $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_n)$ is a vector of weights. By Euler's theorem we can then define the marginal EE as:

$$EE_i^* = \frac{\partial EE_{total}(\boldsymbol{\alpha})}{\partial \alpha_i}.$$

One way to compute the above partial derivative is to change the size of a transaction by a small value and calculate the marginal EE using a finite difference. This does not require any additional simulation but just a rescaling of the future values of one trade by an amount $(1 + \varepsilon)$ followed by a recalculation of the EE for the netting set[3]. The marginal EE of the trade in question is then given by the change in the EE divided by ε. The sum of the marginal EEs will sum to the total EE[4].

Alternatively, as shown by Rosen and Pykhtin (2010), it can be also computed via a conditional expectation:

$$EE_i^* = E[\max(V_i, 0) \,|V_{NS} > 0] = S^{-1} \sum_{k=1}^{S} \max\left(V_{i,s}, 0\right) I(V_{NS} > 0)$$

where $V_{i,s}$ represents the future value for the transaction i in simulation s (ignoring the time suffix) and $V_{NS} = \sum_{i=1}^{n} V_i$ is the future value for the relevant netting set. The function $I(.)$ is the indicator function that takes the value unity if the statement is true and zero otherwise. Such calculations are illustrated in Spreadsheet 10.6. More detail, including discussion on how to deal with collateralised exposures can be found in Rosen and Pykhtin (2010). The intuition behind the above formula is that the future values of the trade in question are added only if the netting set has positive value at the equivalent point. A trade that has a favourable interaction with the overall netting set may then have a negative marginal EE since its future value will be more likely to be negative when the netting set has a positive value.

Whilst marginal EE is easy to calculate as defined above, it does require full storage of all the future values at the trade level. From a systems point of view, marginal EE could be calculated during the overnight batch with little additional effort whereupon the trade-level future values can be discarded. However, for analysing the change in marginal EE under the influence of a new trade(s) then, unlike incremental EE, all trade-level values must be retained.

[3] ε is a small number such as 0.001.
[4] At least in the current case where no collateral is assumed as discussed below.