

Counterparty credit risk – let's get serious

A call for the introduction of simulation-based methods to quantify counterparty credit exposure. By David Rowe

Last month, I argued that counterparty credit exposure estimates based on mark-to-market plus add-ons are inherently inconsistent. I expect that few would argue with this contention. But then why are so many major institutions still using this method as their sole approach to quantifying such exposure? The answer mainly lies in the difficulty of assembling enterprise-wide information. A proper simulation of counterparty exposure requires considerable detail on each trade. In effect, it demands enough structural detail to support reasonably accurate valuation under alternate market conditions. The data requirements for Bank for International Settlements-style calculations are much less severe. Also, as noted last month, the specific BIS calculation must be supported to meet regulatory requirements. Building internal counterparty exposure calculations on a similar basis is often viewed as an attractively economical extension to such a system.

Requirements

The pernicious impact of a simplistic approach to counterparty exposure assessment occurs at many levels. The net result, however, is sub-optimal credit decisions. Unfortunately, the cost of such poor decisions is much harder to identify and isolate than is the direct cost of building improved measurement systems. To make progress, someone must be able to judge the difficult-to-quantify, but very real, benefits of improved credit decisions against the immediate costs necessary to achieve such improvement, and be able to act on that judgement.

A major source of difficulty is fragmented trading systems in most organisations. Reliable counterparty exposure estimates require simultaneous analysis of all the trades done with any customer across all the active trading systems. New XML-based techniques make assembly of the necessary trade details easier than it once was. Nevertheless, the required level of effort and co-operation across departments is significant. Indeed, the data integration challenge is an important reason why many first-tier institutions have not succeeded in going beyond an add-on approach while some second-tier institutions have done so.

In addition to data integration, there is the need to develop and implement analytical tools to perform the simulations and presentation tools to deliver the resulting information to decision-makers in a timely and easily understood fashion.

For all the above reasons, senior management understanding and sponsorship is essential.

Correctly capturing the effects of close-out netting is a definite requirement for accurate ex-



David Rowe is president of the Infinity business unit at SunGard Trading & Risk Systems
e-mail: david.rowe@risk.sungard.com

posure estimation. Doing so requires that simulations be carried out at the transaction level or, at worst, at the legal agreement level. This allows application of the appropriate netting rules within each hypothetical set of market conditions. There simply is no other reliable way to assure that netting is captured accurately. Shortcut methods, such as factoring down exposure based on the ratio of net-to-gross market values, may be marginally acceptable for aggregate portfolio exposure. They are completely inadequate for counterparty-specific exposure estimation.

Another difficulty in implementing an effective counterparty simulation system is the amount of computer processing required. Like market value-at-risk calculations, counterparty exposure simulation requires a significant number of statistically relevant scenarios. These also must be performed at multiple dates over a horizon measured in years. Finally, as noted above, this analysis must be carried out at the transaction level, or at least at the legal agreement level, if netting effects are to be captured consistently.

In purely technical terms, this is not an overwhelming problem. Many scientific applications, such as weather simulation, require much more extensive processing resources than does counterparty exposure simulation. The real constraint is economic and commercial. Improved counterparty exposure estimates must be expected to pay for themselves in higher long-term profits. Achieving this requires a firm resistance to unnecessary analytical complexity. Many technical analysts tend toward what a former colleague of

mine, Hans Fokkema, called "refining within the margin of error". This tendency is especially problematic in counterparty exposure estimation where hypothetical market conditions must be analysed many years into the future. The volatilities and correlations used in constructing such long-term scenarios are themselves surrounded with considerable uncertainty. Constructing highly accurate trade valuations conditional on such uncertain market conditions certainly qualifies as refining within the margin of error.

The answer is to utilise some form of grid pricing. In this approach, a small number of full revaluation results are produced on the basis of controlled perturbations in future market conditions relative to status quo values. These are stored in a structured table or price grid. A Monte Carlo simulation, based on random perturbations, can then be performed in which valuation of the individual trades is based on interpolation from these price grids. This can improve the computational efficiency of the analysis by several thousand times compared to a full valuation approach. Indeed, by saving and reusing the price grids it is quite possible to update simulation-based exposure profiles in near real time (or at least as fast as the trading systems can transmit new trade details to the simulation engine).

Benefit

The central benefit of a simulation approach to this problem is consistency in the exposure estimates across counterparties, regardless of how complex their trading pattern may be. Exposure based on mark-to-market plus add-ons is inherently inconsistent in this sense. It is simply unreasonable to expect consistent risk-versus-reward decisions from credit staff if the exposures they are asked to approve are measured in fundamentally inconsistent ways.

Furthermore, results of a simulation approach clearly reveal the crucial time dimension of where exposure can be expected to arise. Peak exposure on a par interest rate swap will occur about 40% of the way through its remaining life. On the other hand, peak exposure on a currency swap (with principal exchange) will occur just prior to maturity. Even if these peak exposures are the same size, the two transactions represent very different credit decisions. Simulation-based methods make this distinction clear by quantifying exposure at regular intervals over the future life of the longest deal in a counterparty's portfolio. The credit officer does not have to make complex mental adjustments to compensate for inconsistencies in the exposure estimates. The result is bound to be better informed and more consistent credit decisions, more efficient weighing of risk versus reward, and greater long-term profitability. ■