

Dr Andrew Aziz, Executive Vice President, Risk Solutions, Algorithmics, Justyn Trenner, CEO of ClientKnowledge and Steve White, Executive Director, RiskCare discuss the importance of forward-looking risk frameworks that link transaction level decisions with high level strategy and their applicability to the FX market.



Dr Andrew Aziz

Justyn Trenner

Steve White

Outside the comfort zone: FX e-commerce operates in a new era of Risk

Do we live in a new world of foreign exchange risk? The thing is, it is much like the old world, only a great deal faster and a great deal busier. So the old paradigms of risk seem to apply – counterparty risk, processing or settlement risk, technology risk, market risk, criminal risk and of course, the more general risk of losing your clients. However, read into this volumes that are three

times those at the start of the decade, messaging volumes that are more than ten-fold, frequency of trading that is multiples per second and technology specifications leading to microsecond measurements (as against old-school refresh rates measured in the once-a-second frame).

Speed does funny things to people's thinking. It makes them

focus with remarkable intensity on some elements of the spectrum – the one that showed itself last month at your competitor, or the one that seems to be generating previously wholly unpredictable losses, or the one that results in this year's trades being less profitable at the same apparent spreads than last year's, or the one revealed over time by lack of investment. However, the risks

that have not hurt us for long enough become under-attended.

ClientKnowledge has recently undertaken research and analysis among sell-side firms in two different areas. One has focused on liquidity management, and in particular, how best to transfer risk, once captured, into a market-tradeable format, minimizing latency. The other has examined how sell-sides manage and check counterparty credit risk and limits. Of course, if the goal is to minimize latency and credit is easy, two solutions suggest themselves; either eliminate pre-trade checks or make limits easily queried and approximately right, leaving the heavy-lifting analysis for overnight runs.

The argument for the first approach is that not much money has been lost in forex working like this – so far. The problem with the second approach is often systemic – the pricing against which today's overall limits are set are based on last night's closing prices. It is clear that there should be a better approach. This article explores credit risk in a way middle-office credit managers will relate to and in a manner that front office should be able to work with.

Over the last five years many so-called market experts have tried to forecast the future of e-trading

from the demise of single bank portals to the re-emergence of single bank portals, from the expansion of e-trading from cash products to simple and ever more exotic options to the superiority of the hybrid model, from double-digit growth rates brought on by algorithmic trading and increasing use of electronic channels by retail investors to ... well you get the idea.

As Niels Bohr, a famous twentieth-century physicist and one of the fathers of quantum mechanics once remarked "prediction is very difficult, especially if it's about the future."

But if the last twenty years of financial markets are any guide it is a pretty safe bet, overlooking short-term quantum fluctuations, that e-trading volumes will increase, more complex derivatives will be sold over electronic channels, the jockeying between single and multi-bank portals will continue, dealers will continue to compete by offering a broader set of products across asset classes and within product categories, and that there will be periods of market volatility and reduced credit quality and liquidity just at the very moment it seems least

likely. What is also clear is that the risk management frameworks and systems of dealers and clients alike will determine their ability to navigate and prosper in this uncertain future.

Frameworks and systems

So how well equipped are dealer's credit risk management frameworks and systems to cope with the simultaneous pressures of higher volumes, more complex products across asset classes, and credit quality and liquidity issues that can and will appear just at the moment when they are least expected. The Senior Supervisor's Group of Basel in March of this year published a report "Observations of Risk Management Practices during the Recent Market Turbulence" that assesses which risk management practices worked well and which did not during the recent period of market turmoil. A key conclusion of the report is that firms that had forward-looking, scenario-based frameworks that fully incorporated expert judgement on exposures,

limits, reserves, and capital did far better than those that did not.

The remainder of this article contrasts the two primary approaches to modeling counterparty credit exposure, the add-on approach and the forward-looking Monte Carlo simulation approach. We then look at how leading banks, leveraging advances in computational power and disk speeds, can use these simulation-based approaches not only in a defensive manner to control risk but also offensively to link transaction-level risk to high-level strategy.

Credit Exposure Models

Credit exposure models can be assessed based on their ability to capture natural portfolio offsets and credit mitigation in the form of break clauses, give-up agreements, close-out netting, and collateralization where appropriate.

First generation models to estimate this future potential credit exposure are based on an add-on or factor-based approach. In this approach the potential exposure of each trade is estimated based on the currency, product type, maturity, and potentially other characteristics of the trade and then the aggregate exposure at the counterparty, industry or country level is calculated as the simple sum of the trade level exposures. Add-on methods are simple, fast, and reasonably accurate for deal-level exposures. But clearly they can dramatically overstate the exposure of the credit capital that must be held at the counterparty level. Although more sophisticated add-on methodologies have been

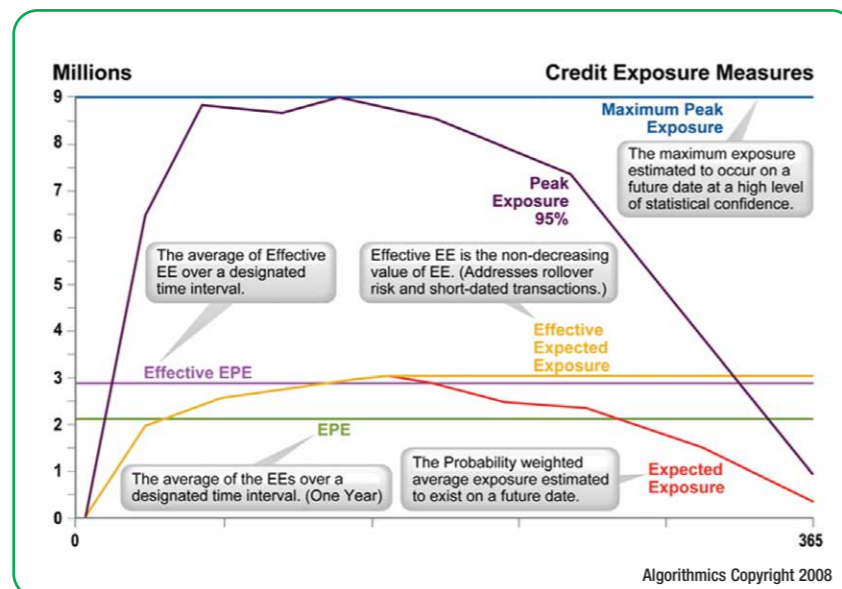

developed it is now widely recognized that add-ons cannot capture the stochastic nature of exposures and the richness of netting and collateral behaviour. In spite of the analytic limitations, add-on models are still widely used in pre-deal limit systems due to their inherent speed and simplicity. But as we shall discuss later, this is changing.

The current gold standard for estimating the potential future

exposure of OTC derivatives including FX products is a multi-step Monte Carlo simulation approach. Monte Carlo simulation is a mathematical process that is used to imitate the “real-life” behaviour of random processes such as financial prices and rates.

Multi-step Monte Carlo simulation is an extension of a single-step Monte Carlo simulation to cover an array of future time points.

Calculating Credit Exposure using Monte Carlo Simulation

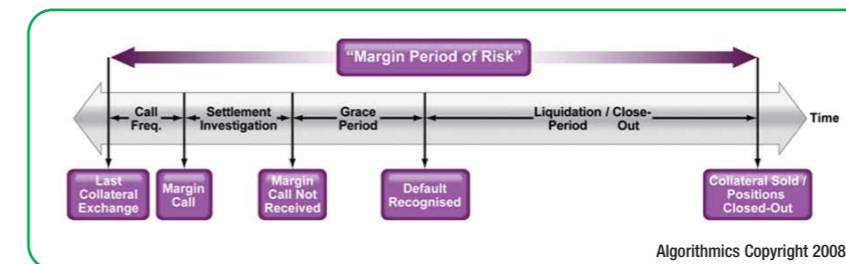


The methodology calculates exposures in a way that takes into account market risk consistency across transactions, portfolios, and asset classes. The pricing of transactions under the Monte Carlo scenarios also considers the path dependent nature of certain products like barriers and swaptions.

Modeling of Collateralization

The use of collateral in mitigating counterparty risk is a technique widely used in the market. A Monte Carlo simulation-based approach takes into account credit support annex rules while calculating the exposures at future time steps. Netted exposure along with the opening collateral balance at each scenario and time step is used to simulate future margin calls. The calculations also incorporate an additional element of risk corresponding to a possible drift in the portfolio value between a possible default and eventual close-out of the positions – the so-called margin period of risk.

The margin period of risk is depicted below:



Post-Trade vs. Pre-Trade, Linking Transaction Level Risk to High Level Strategy

These forward looking, Monte Carlo-based simulation approaches are now commonly used on an overnight batch basis to calculate credit exposures for all traded positions and counterparties. But

COUNTERPARTY CREDIT RISK is the risk that a counterparty to a financial transaction will default prior to the expiration of the contract and will not make all the payments required by the contract. If a counterparty in a derivative contract defaults, the maximum possible loss to the surviving contract is determined by the contract’s risk-free value at the time of default. If the derivative assumes an exchange of cash flows between two counterparties (as for example in a swap) the value can be positive or negative. And since the contract value changes unpredictably over time with the level of interest rates and FX rates and other market factors, only the current exposure is known with certainty, while future exposure is uncertain.

the dream has always been to shift this analysis from a once a day, post-trade mode, to a mode where risk and capital are measured before each and every transaction. Today, that dream is becoming a reality. Organizations are beginning to implement systems that enable them to make risk-adjusted reward decisions before they transact. No longer are risk systems only high level strategic support tools. Now every tactical decision at the micro level can be aligned with the organization’s strategic objectives, even in high volume businesses like FX.

Initially, there were major technological and cultural constraints to doing more than

collateral and complex netting agreements, to mitigate credit exposures.

At that time, pre-deal assessment in real time using full simulation, yet still incorporating all the mitigation mechanisms consistently, was inconceivable. But it was not just the machines that needed to improve. What was required were technological frameworks that could gather all the relevant data and consolidate it, and perform the thousands of scenario simulations to price risk, not only of the market risk of the instruments, but also the credit risk of the counterparty, including risk mitigants. Organizations with operations around the globe required systems that could run 24x7 and still incorporate new scenarios as markets closed, without disrupting continuous trade processing. Now we have these frameworks. The hardware is very much more powerful, and the software and architectures no longer rely on a big black box in the middle to do all the number crunching. Instead, an organization can distribute scenario simulations across many machines in various locations. With a service-oriented architecture (SOA), an organization can now have a simulation engine

wherever it requires, enabling it to calculate risk as and when it happens – say when a new trade is about to take place, or a when a counterparty nears a credit limit. There have also been cultural constraints to bringing risk into the decision-making process in the trading room, but these are beginning to disappear. There is a growing acceptance in the front office of risk-adjusted reward calculations, and now most organizations accept the concept of limits based on stochastic exposures rather than notional limits.

Improvements in speed and accuracy have also helped convince traders of the value of pre-deal risk calculations. A key motivator for the front office is to allow trading 'closer to the limit', allowing traders to avoid the false positive limit breaches that can impede trading. In fact, various levels of speed and accuracy are possible. As products become commoditized, so margins fall and it becomes necessary to calculate risk to finer degrees of accuracy. For example, a system might be able to represent collateral, but not the risk during the margin period.

These kind of pre-deal risk-reward calculations are being carried out even in the high speed, high volume world of electronic trading now. Real time in this world is measured in double-digit milliseconds where as recently as 2003, 'real time' meant a couple of seconds. Organizations are implementing algorithmic trading systems that may be minus the human trader, but still include a risk reward assessment before every trade. By implementing such a framework with an appropriate

level of accuracy, an organization can start to make risk based decisions at the micro level across the board. And once it can do that, then it becomes possible to link those decisions to the organization's macro strategy in terms of the risk profile it wants to achieve, and how it wants to allocate its capital. Where once there were very separate trading silos in organizations, each making their own risk decisions, now organizations can have the ability to link all decision making –

across the enterprise and at the transaction level – to their overall strategy, appropriately capturing the impact of correlation and natural offsets across the former silos. Some of the most interesting steps in this direction are being taken by organizations whose use of risk technology is not driven by compliance. They have invested in risk technology in order to make the most intelligent risk- and reward-based allocations and trades. They allocate capital based on a consistent approach to measuring risk across all businesses and risk factors, and assessing the incremental impact to aggregate risk. In doing so, they are creating value for their businesses. In this new line up, the quantification of risk occurs upstream, prior to any trading activity. Economic capital is on the desk top, part of every

trading decision. P&L attribution is a natural by-product, rather than an add-on. In effect, some of the concepts and techniques used by prime brokers providing VaR-based margining to their hedge fund clients are being used internally to distribute risk-weighted among their various product lines. The building blocks of a bank changes from "Capital > Trade > Risk" to become "Capital > Risk > Trade".

Conclusion

We have reached the stage where if you were to create a financial organization from scratch today, you would build it so that every trade that is done, and every position that is taken, contributes not just to some limit system that manages the exposures according to the risk profile the organization wants to take, but also ultimately contribute to its risk-adjusted return on capital calculations. This leads to a real integration of enterprise risk management with economic capital and, ultimately, with corporate finance. In the end, this is all about value-based management, where shareholder value is a function of a firm's risk and reward. Now an organization can make those value-based decisions at the micro level so that they will aggregate up to be consistent with its strategic view of its business.

The most challenging aspect of all of this is often in the implementation. Most banks have not one but several legacy credit systems. It becomes more appropriate to implement major change and replace than graft on improved elements. And that implies significant investment and lead time. The first step down that road is a deep understanding of what is now possible; the second is a keen sense of why it matters.