# Advances in wholesale energy credit risk management

**Donald C. Mumma of Axiom Software Laboratories considers some advances now available to the industry that can improve competitive advantage**

The Enron debacle was the last wake-up call for the credit risk issues facing the energy industry. Earlier events, such as the 1998 Cinergy spike and California in 2001, were also wake-up calls. So did the industry wake up? Yes, and the general response has been to take less credit risk, or transfer it to others, rather than managing it better. Keeping up with advances in managing credit risk can improve a firm's competitive advantage.

### What is energy credit risk management really about?

Simply, it is about measuring the uncertainties and variability of cashflows caused by customers (buyers) and suppliers (sellers) who fail to honour their obligations. Then it is about capitalising this activity and, using strategic and tactical management techniques, to optimise the return on this risk.

So how do you get there? To manage return on risk, you first need a numerator and a denominator. The value of the numerator is, in many respects, a management accounting measure that is not as market-sensitive as the denominator, which is more variable. As deals go in-the-money, credit exposure increases. It decreases as deals go out-of-the-money. This article offers some new insights and shows how Axiom SL's RiskMonitor implements credit risk management.

### What are the drivers of credit risk?

The three elements used to measure credit risk capital are exposure, defaults and recoveries. All are unknowns. There are, however, some things we can know or estimate in our measurement model. Let's take one at a time.

**Exposure:**

This is not just the amounts that are owed to you. It also includes the cost of carrying positions where you owe others. The task is to measure both the current exposure and to estimate the potential future exposure profile over some time horizon. There are several components to measuring exposure. The first is contract replacement value, also known as pre-settlement exposure. The second is settlement exposure. This is when you send something out the door (physical product or money) before you have received the counterparty's side of the bargain. This is usually considered to be receivables, but it also includes unbilled deliveries, margin payments, eastern global time zone deliveries against western global time zone receipts, etc. The determinants of future amounts may be known or unknown. One commonly overlooked known variable is the time to a 'default discovery'. In a sale, you send out product over some billing interval. Then you send more product out the door while you wait for receipt under normal payment terms, and then you send more product out the door while you chase down the late payment. In a purchase, the time to discover a default is much shorter, and you do not have the same settlement exposure.

The unknown variables that drive future credit exposure are contract price structure versus market prices and, perhaps, volume. When a contract price is fixed, settlement exposure is not so sensitive to market price changes, but the pre-settlement exposure is. Conversely, when a contract price is floating, the pre-settlement exposure is not so sensitive to market price changes, but the settlement exposure is.

The process for estimating settlement and pre-settlement exposure is fairly well established. Axiom SL's RiskMonitor performs a

**Figure 1. Potential future exposure**

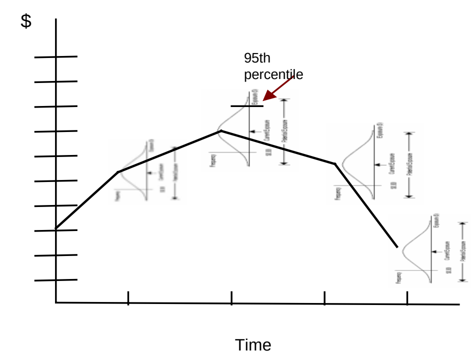

Monte Carlo simulation at user-defined intervals (points) over an overall holding period, also user-defined. Thus, at each future point there is a distribution of exposures, as illustrated in figure 1, at a deal level and for each aggregation criteria, which is used for the set of returned potential future exposure values.

The average potential future exposure (APFE) is the time-weighted mean value of each interval's exposure distribution. The maximum potential exposure (MPFE) is the highest of the averages. The confidence level potential future exposure (CLPFE) is the highest of the nth (user-defined) percentile observed value of all the intervals' distributions. The average and maximum PFEs are often used in the downstream calculation of credit risk capital, while the CLPFE is most often used for credit limits. The system also keeps track of the MPFEs and CLPFEs at each simulation point, so that they can be applied against RiskMonitor's time-laddered limit structure. For example, if the highest CLPFE occurs in year three, but the six-month point CLPFE is half the amount, it is indicative of the worst-case exposure of that deal if default occurred in the sixth month. Typically, credit exposure appetite is higher for shorter terms, because of the lower uncertainty of default estimates in the near term. The exposure profile also gives effect to deal run off in the aggregation of portfolios.

There are two other factors that will affect the aggregation of credit exposures. The first is netting counterparty deals and the application of margin. The second is the portfolio effects within or across counterparties. RiskMonitor provides both a gross and net calculation for each of its current exposure and PFE measures. Figure 2 shows the PFE time profile for different types of deals for a single counterparty with netting. If the deals were for different counterparties, gross and net totals would be the same, because a negative valued deal has a credit exposure profile with minimum value of zero. If margin is required to be given to a counterparty, this amount should be added to the exposure.

**Figure 2. Total potential future exposure**

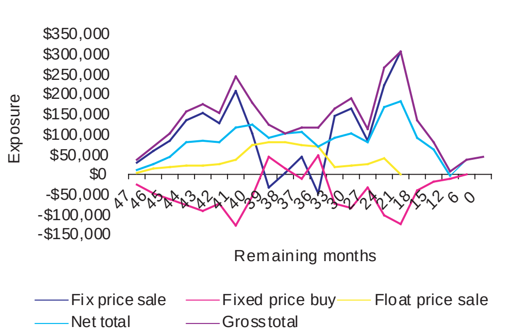


**Defaults**:
While exposure is a continuing measure, default is a one-time event for a counterparty. There are only two possible states: in compliance or in default. The challenge in calculating credit risk capital is to first decide the default rate input for a counterparty, and then how to use it in the portfolio calculation.

## The challenge in calculating credit risk capital is to first decide the default rate input for a counterparty, and then how to use it in the portfolio calculation

Default frequency estimates are expressed as a probability over some time horizon. There are many different ways to estimate the default probability of a counterparty, but the two most frequently used methods are the Merton model approach, and the second is a financial analysis-based assignment of a credit rating, and using the historical default frequency statistics applicable to the rating. Many use the rating agency statistics.

Each method has its strengths and weaknesses. The most frequently mentioned weaknesses of the rating statistics are a) they are based on a historical (sometimes rating agency) universe, which may not be valid for the energy company's portfolio today, and b) external default data is for debt defaults, not energy contract defaults. This writer is not aware of any research analysis of historical differences between energy default frequencies and debt default frequencies, but there is plenty of anecdotal evidence of companies defaulting on their energy contracts before they default on their debt. The weaknesses frequently mentioned in using the Merton model approach are that a) it is a point in time measure, rather than 'through the cycle', and b) the model does not adequately capture the volatility of liabilities. We recommend analysts use the same market price-based distribution of credit exposures to enhance liability inputs to either of the default probability estimation approaches mentioned above.

Another unknown, but possibly significant measure is the degree of correlation of default probability among other counterparties, and/or to market price factors that influence exposure. The market offers reasonable default correlations among counterparties with the same business characteristics.

**Recoveries**:
This driver of credit risk is also known as loss given default, which is simply one minus the recovery rate. It is usually expressed as a percentage of the exposure. The levels of recovery estimates are typically 'bucketed' by structure features with variability sometimes 'fixed' and/or correlated with the variability of the exposure. For example, in the case of deals 'secured' by margin and netting, which are in effect recovery enhancement mechanisms, these values are often de-

ducted off the top of exposures. If the variability of recoveries and exposures are extreme enough, the recovery rate of secured exposures can be significantly impaired by the time to default 'discovery', contract termination and liquidation of deals and margin. Long-Term Capital Management is a sad example of this possibility and is a reason why many firms also have gross exposure limits.

**Credit risk capital**:
The calculation of credit risk capital can be performed by a combination of simulations of these risk factors (exposures, defaults and recoveries) and/or closed-form calculations. An additional user-defined parameter is the confidence level for the credit losses. The resulting measure is sometimes referred to as CVAR. There is not yet a consensus of opinion whether the measure should be the nth confidence level of the distribution of credit losses or the difference between the nth confidence level and the mean loss of the distribution. Although most firms that perform this measure use the latter. The first-order calculation is the enterprise portfolio calculation across all counterparties. Because of portfolio effects, the total enterprise credit risk capital will not be the sum of individual counterparties' credit risk capital. It will also not be equal to the sum of each counterparty's (exposure X default probability X [1 – recovery rate]).

Figure 3 is a RiskMonitor display of exposures, exposure time profile, capital and exposure sensitivities to directional price changes.

**Figure 3. RiskMonitor**

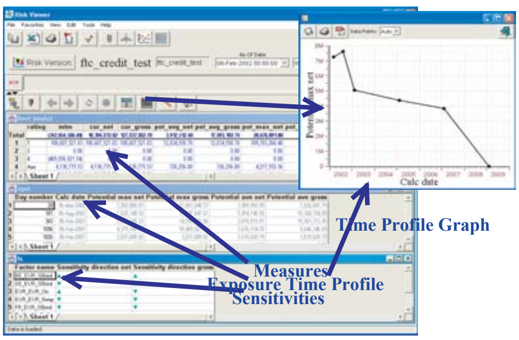


There are many techniques for apportioning credit risk capital to a single deal. The most commonly used method is sometimes called incremental capital. For an existing portfolio, a deal's incremental capital is the difference between the amount with the deal in the portfolio and the amount without the deal. This measure can be easily performed when the portfolio's credit risk capital is first calculated. For a new deal, the calculation process is somewhat more challenging. The computation time needed to do an enterprise simulation calculation each time a new deal is being considered is not practical. So performing the measure at the counterparty level or a closed form method is preferable because of the faster calculation speed. Many firms already calculate credit exposures for pre-deal pro-forma limit compliance, so adding the credit risk capital measure with the new deal at a counterparty level, during this process, does not take much longer to compute.

What is often not considered is that certain deals with certain counterparties may actually lower credit risk capital. One way to see whether a deal may reduce capital is by the directional price sensitivity of the counterparty's portfolio to increased credit exposure. For example, if a portfolio increases credit exposure when a risk factor price rises (the firm is a net fixed price purchaser or floating price seller with a counterparty), then an offsetting deal (one that increases credit exposure when the same risk factor price declines) will affect the potential future exposure of the portfolio more favourably than if the deal is done with a counterparty that is credit exposure sensitive to a decline in the risk factor price. If deals can be netted, potential future exposure will be reduced as an increase in exposure on one side of the portfolio will be offset by a reduction of exposure on the other side. If the deal cannot be netted, at least the potential future exposure will, perhaps, be no greater than it is now, because both deals cannot increase exposure at the same time. Of course the size and sensitivity (delta) of the deal relative to the portfolio will determine the magnitude of its effect on credit risk capital.

**Return on risk capital**:
So now we have a denominator. The next step in managing credit risk is to add a numerator (return) to the measure. Few do this now. It may take incremental steps at a time. But the ultimate objective is to be able to perform this measure quickly enough so that transaction professionals and management are able to see, pro-forma, the impact of a deal on not just profits, but return on risk capital.

Allocation of profits (and credit carrying costs, such as margin)to a counterparty in a trading or wholesale business is no easy task. It takes considerable internal discussion and, hopefully, consensus. But it should not be abandoned because it is too hard. Usually there is a benchmark bid-and-offer price where the mid-point spread to the seller's price can be allocated to seller profit and a mid-point spread to a buyer's price can be allocated to a buyer. If the horizon of the capital measure is multi-year or a fractional, the return on credit risk capital should be annualised. Measures can be absolute or in relation to a specific hurdle rate. If credit risk capital and profits cannot be allocated to counterparties, then at least the first step is to measure it at a business unit level and make it a performance criteria for management. After that, managements will become more interested in breaking it down into smaller bits that they are able to manage.

Best practice is emerging where business units are charged for their cost of credit risk capital. The tools to perform these measures are no longer out of reach financially, physically or intellectually. Those players doing it now have a competitive advantage. ■

*Donald C. Mumma, Managing director, Axiom Software Laboratories.*
*E-mail: don@axiomsl.com*