HEADLINE: From Enterprise Risk Management to "Quantitative Enterprise Management"

AUTHOR: Joe Pimbley, Managing Director of Quantitative Risk Management, American Capital Access.

Enterprise risk management (ERM) is a philosophy and set of procedures with which financial firms aspire to monitor and manage risk. Effective ERM remains "next-generation" in the sense that few, if any, firms have fully implemented its methods. In fact, the discussion of what constitutes ERM continues.

In this short article I will present a concrete view of ERM and what it means to a financial institution. The purpose is not to emphasize this "next-generation" ERM but, rather, to look past ERM to its eventual successor – an approach that could be termed "quantitative enterprise management" or QEM.

As a way of setting the stage, let's consider how ERM works by looking at a small, self-contained modelling problem from the world of structured credit. Investors in collateralised debt obligations (CDOs) purchase tranches at different points in the capital structure (equity to senior). Our hypothetical CDO contains 100 bonds of roughly the same face amount. How should one model the risk and return for each tranche?

The conventional approach is to specify a (possibly time-varying) default rate that applies uniformly to each bond in the portfolio and then determine the extent of losses to each tranche. But such a rudimentary method (which is still quite challenging to implement) cannot answer basic questions, such as how differing bond face amounts or credit ratings affect the risks to each tranche.

The state of the art in CDO modelling uses the Monte Carlo methodology, which permits the calculation of stochastic (probabilistic) defaults for each bond. This best-practice modelling considers each "atomic" risk element (in this case the potential default of each bond) along with estimated probabilities and correlations among all elements.

So, an investor who owns one tranche in this CDO, must analyse the investment with a good CDO model. Imagine now that this investor purchases a second CDO tranche of a different transaction. There are 20 bonds in the two CDO portfolios that "overlap". Clearly, then, separate analyses of the two CDO investments will be incomplete. The investor must model both CDO positions within one larger portfolio such that the default of an "overlap bond" will harm both investments simultaneously.

It gets harder. Following this line of thought, every investment – whether an exotic structured finance position or a stand-alone bond – has to be considered as a constituent of the investor's portfolio. The most challenging problem in finance is to determine rigorously whether a new investment is advantageous. The lesson from the CDO world is that this determination must take into account the investor's existing portfolio.

This observation is the cornerstone for the necessity of ERM. That is, the risk of an enterprise (or portfolio) is not a simple sum of the individual risks of the "atomic" elements. ERM requires a concerted modelling effort of the entire portfolio of risks. The modelling, in turn, demands copious quantities of accurate data. For this exercise to succeed, the modelling and data for all positions must exist on the same software and hardware platforms.

Expressed in those terms, it's easy to see why no firm has yet fully implemented the tenets of ERM. It is much, much more difficult than it may sound to build firm-wide software with state of the art models that robustly incorporates all the necessary data feeds and data specific to the risk positions.

To take ERM to the next level, we must recognize that a corporate firm is itself a CDO. A firm's potential loss in its risk portfolio is limited by its capital. The creditors of the firm bear all the risk beyond the firm's capital. A firm-wide management system must recognize this inherent loss limitation for determining the investment quality of a new risk position.

Quantitative enterprise management (QEM) will go beyond the ERM assessment of firm-wide risk and return of the current risk portfolio to generate stochastic information for firm performance: QEM models, software, data, and IT will include ERM as a subset but will also express (probabilistically) future changes in market size and market share. Upside surprises such as new product introduction will enter the model as will negative surprises, like operational errors, the closure of a foreign market, and so on.

QEM output will likely include stochastic projections of future financial information – like net income, free cash flow, and retained earnings – and even probability distributions for future equity and debt values of the firm.

ERM arguably sprang from the modelling rigor of structured finance and is very much a "bottom-up" methodology. That is, get everything rigorously correct at the lowest transaction level and then build upward. QEM is more of a high-level concept.

But the huge effort involved in pursuing firm-wide risk management invites scepticism. Is it worthwhile?

Even if years of toil provide a firm with high-quality ERM (or, in the next generation, QEM), there's still a problem. The inputs to the models are not well known. When we specify a default probability for an underlying bond, for example, how much confidence do we have in this probability? How do we test it?

If we're honest, we have no idea how well we've estimated these default probabilities, correlations, or volatilities of various market variables. The sceptic would then exclaim that unknown uncertainty in the input must lead to unknown uncertainty in the output. And (galling as it is to admit it) the sceptic is right.

On our side, though, we have two strong arguments in rebuttal. First, if application of deep quantitative analysis cannot succeed, then what is the alternative? Shall we all rely only on subjective intuition to decide whether investment in a new triple-B CDO tranche at a yield of LIBOR + 200 basis points per annum is a good, firm-wide investment? Or should we do as much objective, quantitative analysis as possible and then interpret the results of this analysis with some "fuzziness" to recognize the "unknown uncertainty".

Our second of the two arguments rebutting the sceptic is simply that modelling is a tremendously healthy exercise. Even when we know that the model results have "unknown uncertainty", the modelling process itself teaches us a great deal about the nature of risk and return in the financial world. Modelling teaches us "how to think".

Rigorous modelling also exercises the IT platform and serves as the best assessment of data quality. There are gigabytes of data for every business. Data quality is of paramount importance, but it's extremely difficult to create algorithms that test the data quality. In fact, it's possible that the models of ERM and QEM are the best "data test algorithms" conceivable.

The point here is that the exercises of ERM and QEM will pay ancillary dividends beyond output results that say something like "adding a mortgage warehouse business now will shift the mean and variance of the future firm equity value" in a particular way.

In summary, while ERM is tomorrow's paradigm for Chief Risk Officers (CROs) of financial firms, QEM is the likely succeeding paradigm. In this next world, risk management will encroach so fundamentally on the role of the CFO that, in the spirit of QEM, we may see the CRO and CFO coalesce into one role. If so, we will, in a way, be coming full circle: it was the growing sophistication in risk management that separated the two roles in the first place.