Kamakura Corporation

2222 Kalakaua Avenue
Suite 1400
Honolulu, Hawaii 96815, USA
phone 1 808 791 9888
facsimile 1 808 791 9898
www.kamakuraco.com

Advanced Financial Risk Management:

Tools and Techniques for Integrated Credit Risk, Market Risk, Liquidity Risk and Interest Rate Risk Management

April 2012

Donald R. van Deventer
Kenji Imai
Mark Mesler

Part 5: Portfolio Strategy and Risk Management

Chapter 40:
Information technology considerations

In the first 39 chapters of this book, we have reviewed the tools necessary to construct a fully integrated measure of interest rate risk, market risk, liquidity risk, foreign exchange risk and credit risk. This measure, the Jarrow-Merton put option, produces an answer to the key question "what is the hedge?" Some of the tools we have discussed were focused on risk management strategy, such as usage of the Jarrow-Merton put option as a comprehensive measure of risk. Other tools have been more mechanical, such as how to value an American call option embedded in a bond when the default of the issuer is a possibility. In this chapter, we focus on another set of issues that have to be faced when implementing a comprehensive risk management system—the information technology aspects of an implementation.

The authors collectively have spent 105 years in the implementation of risk systems in major financial institutions in 23 countries around the world, and our comments below reflect the world-wide best practice in that regard.

Common Practice in Risk Management Systems: Dealing with Legacy Systems

Many of the world's financial institutions face a common situation when preparing to move forward in the risk management systems area. The move forward is usually motivated not changes in regulatory risk measures, like the Comprehensive Capital Analysis and Review 2012 of the Federal Reserve and similar stress tests introduced in Europe. Particularly in countries where there has been a recent financial crisis, managers of the most sophisticated financial institutions now clearly realize that the common situation described in the bullet points below have left the institution unprepared for a crisis. These managers are re-dedicated to the comprehensive view of risk that we have described in this book because their previous failure to take this approach nearly resulted in the failure of their institution. In a few cases, the institution did fail (as we saw in Chapter 37) but government capital injections preserved its role in the financial markets.

A common situation that senior management often faces looks like this:

- The risk management function is fragmented. Market risk, credit risk, liquidity risk, interest rate risk and performance measurement are all handled by different groups with no central coordination at either the working level or the senior management level.
- There is at best only a partially completed enterprise wide data warehouse.
- Risk management systems have been implemented on a compartmentalized basis:
    - Vendor A provides third party default probabilities.
    - Vendor B provides credit risk loss simulation software that provides cumulative losses over a long single-period but no valuation, multi-period simulation or hedging.
    - Vendor C provides interest rate risk management simulation and mark-to-market calculations, generally ignoring default risk, and using primitive 1 factor interest rate models.
    - Vendor D provides traditional value-at-risk calculations for market risk.
    - Vendor E provides performance measurement for each business unit, either on a benchmark basis or on a transfer pricing basis, depending on the nature of the financial institution.
    - Vendor F (often an in-house solution) provides capital allocation as a supplement to Vendor E's performance measurement data.
- These systems use different graphic user interfaces, different input data-base design, different financial mathematics, different reporting tools, different end-user staffs and are totally irreconcilable.

This common situation is totally wasteful, as 95% of the lines of source code in the six systems should be common. In effect, the user has bought the same system six times and then completely scrambled the picture even more with inconsistent data architecture and reporting. As wasteful is this situation is, for many years it was not the fault of software users. Legacy risk vendors as a group were incapable of producing risk systems that could analyze multiple risk silos on an accurate, consistent and transparent basis. Over the last 15 years, however, that situation has dramatically changed.

Even more than the wastefulness of this typical situation is its danger—the systems infrastructure obscures the view of the organization's risk and return and defeats the very objectives of the risk management process that we outlined in Chapter 1. Most importantly, there is no link at all between random movements in macroeconomic factors which drive defaults and valuation for every item on the balance sheet of the financial institution. How can there be an integrated view of risk when each system is looking at a different piece of the puzzle? This is perhaps the biggest reason for the failures

and near failures due to home price collapses of firms like Bank of America, Citigroup, Royal Bank of Scotland, HBOS, Northern Rock PLC, Wachovia, Lehman Brothers, Bear Stearns, Washington Mutual, Federal National Mortgage Association, Federal Home Loan Mortgage Corporation and many others in the credit crisis of 2007 to 2010. A similarly long list of failures stems from the macro factor influences on the Japanese banking and life insurance system after the burst of the Japanese bubble beginning in 1990.

With a typical siloed risk system infrastructure, the problems are too numerous to list them all. Here's just one critical example. A rise in interest rates will never increase default probabilities and change credit-adjusted valuation in an integrated way in a fragmented risk architecture. The default probability vendor ignores the impact of interest rates on default probabilities and does not take inputs from simulations in other systems. The interest rate risk system ignores default. The financial institution is completely blind with respect to simulations and stress tests of rising interest rates and their impact on default.

This architecture has extremely serious flaws that are of grave concern to many. If senior management has ended up in this situation prior to the developments in risk technology that we've outlined in this book, how can we escape this untenable situation?

It is to that task that we now turn.

Upgrading the Risk Infrastructure: The Request for Proposal Process

The first step in the process of reforming the risk infrastructure starts with a request for proposal process. In a perfect world, there would be a single risk management staff group that manages market risk, credit risk, interest rate risk, liquidity risk and performance measurement on a fully integrated basis. There are many fine books on the subject of the organization and execution of enterprise wide risk management, so we won't detour in that direction here. We need to remain focused on the more micro-tasks of execution. Even if the staff of these groups remains separate with a different managerial reporting hierarchy for each group, a joint project team should be organized to design and implement the request for proposal process. If the team is not a joint team, each group will fight to preserve its right to choose the vendor, rather than fighting to choose the best integrated risk vendor. Senior management intervention to fix this is often inevitable.

In order for the request for proposal process to clearly reveal the skills, strengths and liabilities of the potential vendors, it has to be both comprehensive and very detailed. This means that the project team members will have to be both relatively senior and with considerable hands on experience in each of the areas of risk we want to cover— market risk, credit risk, interest rate risk, liquidity risk and performance measurement. A typical 'best practice' request for proposal covers all four areas with equal emphasis and has 200-1,000 very specific questions on the product functionality along the lines of the approach outlined in this book.

Prior to 1995, it was uncommon to see such integrated risk request for proposals because, frankly, the technology was still in its early stages. Now it is obvious to most market participants that the technology outlined in this book exists and has been implemented by many of the world's most successful financial institutions. For that reason, the authors estimate that 80% of the requests for proposal for risk management systems reflect this integrated approach to risk. The only area which

seems to be lagging is the large American banks where the legacy of many mergers has left high walls between risk management groups of the merged banks. Even in the U.S., however, the trend toward integrated risk management is very strong and distinct. Legacy risk groups, however, often resist being "merged" to improve risk management with a tenacity that demonstrates how unfit they are for a senior role in risk management.

One of the authors learned to his chagrin when he resigned from Lehman Brothers that the only profession in the world with a lower perceived integrity than investment bankers is that of software salesman. This is a sad testimony to the industry and one that the Request For Proposal process has to be explicitly designed to deal with.

The level of detail in the question of the RFP should be high for two reasons. Firstly, it will test the comprehensiveness of both the vendors' knowledge and completeness of the vendors' solutions. Secondly, it will allow for a detailed 'audit' of the vendors' solutions after the RFP responses have been gathered. There is a third benefit to the RFP process. It is common for working level risk staff to prefer a vendor who does not have the best risk system. One reason is prior use of a legacy risk system, allowing the risk analyst to avoid getting better. Another reason, candidly, is bribery in spite of the increasing enforcement of the Foreign Corrupt Practices Act via fines and jail terms. A detailed RFP and its response make it very clear when the inferior vendor has been selected and will immediately trigger an investigation in a sophisticated firm.

The authors believe that 'best practice' in the RFP process is for short-listed vendors (typically two or three firms) to do more than just demonstrate the product. Even a half-day demonstration can be scripted by a vendor in such a way that deficiencies in the product will remain out of view of the potential client. In addition to the demonstration, the potential client should go through every response to every question provided by the vendor and insure that the responses provided by the vendor are accurate. Much to the authors' surprise, one very large international bank which followed this process (i.e. warning the vendors in advance), revealed that only one vendor had been truthful in their responses to the RFP.

By following this process, risk managers will both insure that they have a good understanding of the competing products and will help to root out unethical practices and unethical people from the software industry.

Most importantly, no product is perfect and no product has all features that all potential clients might want. Working with a key vendor for many years on the best future development path for the product is the very best way for a financial institution to achieve 'best practice' in integrated risk management.

Paid Pilots as Final Proof of Concept

The authors believe that once a vendor has been tentatively selected, best practice leads to a paid pilot with that vendor's system while the number two vendor waits on the sidelines. This is essentially a 'mini-installation' and insures, in as complete a fashion as possible, that the vendor's solution works as advertised. While a few vendors resist the pilot concept, the authors would not recommend purchasing any risk management system without one if there is any doubt about the vendor selected. Paid pilots reveal the strengths and weaknesses of both the system and the management team of the vendor in short order. Since everyone in the world would be interested in a free pilot, potential clients should

realize that vendors have to ration their scarce resources by charging for the pilot, with a portion of the pilot fee often credited against the ultimate purchase price of the pilot. We strongly recommend that the key team members on the pilot from the user's perspective lead the user's installation team.
In the last few years, mergers and acquisitions have reduced the number of significant risk vendors to a very small single digit number. By "risk vendor," we mean a vendor who can analyze all 700 million assets and liabilities on the balance sheet of one of the largest banks in the world, not just the 2,000 traded products with a CUSIP or an ISIN number. Because the number of vendors is so small, it is very easy to do detailed due diligence on each vendor by calling their existing clients on an unsolicited basis. This will reveal, in the case of a few vendors, very high installation costs, high failures rates, a lack of innovation, and antiquated product architecture. By making these phone calls, one may be able to skip the costs of a paid pilot and proceed immediately with a full installation of the right vendor.

Keys to Success in Software Installation

The success rate on software installations in general is surprisingly low. The word 'surprising' is the key word here. There should be only one reason that an installation fails and that is the sad situation where the software doesn't work—this should never happen if the process above is carefully followed. As a result, the failure of an installation of software that works is truly a failure of the management process and can be avoided 100% of the time if management is diligent at both the user and the vendor level.

The key to success from the user perspective is to have a dedicated project team 100% committed to the installation as their full time job. The installation team should have a mix of business experts who represent the ultimate end-user internally and IT experts who are experts at getting the necessary input data and changing its format using standard industry tools. These tools include both data mapping and reporting of data base (both input files and output files) contents.

Turnover on this project team is the greatest risk to the success of the project, because it is quite difficult for the vendor to replace the institution-specific knowledge that is lost when someone on the user's project team resigns. This knowledge normally can come only from within the user's institution itself.

The vendor's installation team normally has three types of experts. The first, the project manager, is a generalist whose expertise involves both clear and organized planning and reporting. The second type of expert is the excellent business user of the software from the vendor's perspective. The third type of expert is the IT expert who works closely with his peers on the user's project team.

The vendor's project team members typically come from three sources: from the vendor's staff itself, from a formal distributor of the software, or from a third party systems integrator. The authors believe that the success rate is highest if the project team on the vendor side comes either from the vendor itself or from a long-term distributor of the project. If the only relationship with the vendor is on a project basis, a third party systems integrator is in a complex position. The short-term revenue from the project is maximized by dragging out the installation, and the systems integrator suffers relatively little from the long-term reputational damage to the vendor from projects where installation was too slow, too expensive, or failed. From an end-user's perspective, the user should want an installation team very focused on how important a successful installation is to the reputation of the vendor.

When this long-term reputational perspective is properly in place, a project should never fail unless the end user for some reason doesn't staff the user side of the installation correctly.

Vendor Size: Larger Vendor or Small Vendor?

From a conceptual point of view, every head of information technology wants a software vendor with a ten year default probabilitiy of 0.03%, a 50-year history of innovation, and a project team of 25-year veterans that works for minimum wage. Unfortunately, the nature of the risk management software industry makes this dream highly unlikely to come about. As we look at the industry, a number of trends are quite clear:

- The innovation in the risk management business has consistently come from new, small firms.
- Many of these small firms have been able to accomplish only one major innovation and end-users are forced to switch vendors in order to purchase the next innovation.
- As a result, small firms which can do multiple innovations and which allow changes in risk management models in the software with just a mouse-click by the user are very highly prized.
- This increases the size of the system that must be constructed by a new entrant, 'raising the bar' and discouraging new entrants to the business.
- Large software firms generally add to their risk management products by acquisition of other firms.
- After an acquisition by a large firm, culture clashes between the large firm-small firm mentality often lead to high turnover among the most valuable staff of the acquired firm.
- In response, the large firm ceases innovation on the product and puts it into 'cash cow' mode until it dies.
- When a risk product does die because of lack of innovation, the large software firm simply buys another vendor.

How does a sophisticated potential buyer of integrated risk management software deal with these issues? The most intelligent IT experts at major financial institutions look at the stability and ability of senior management of the vendor and of the core development team. This measure of 'vendor risk' is almost flawless in indicating the chances of a successful installation and a long-term relationship, because it says a lot about the skill of the management of the vendor at all levels. It is a much better indicator of the probability of success than the size of the vendor, because firms of all sizes can score well on this measure if the firm is well-run. At a large firm, compensation schemes designed to preserve the intellectual property of a small firm that is acquired are simple to structure and execute and yet relatively few large firms have been good at it.

A couple of incidents from the recent past illustrate how risky a large software vendor can be from an end-user perspective if the low turnover strategy is not pursued:

- A major U.K. software house bought an innovative risk software firm with a 15 person development team based in Los Angeles. Shortly after the acquisition, each member of the Los Angeles development team was ordered to move to London. As attractive as London is as a place to live, 14 of the 15 team members resigned rather than relocating. This essentially destroyed the product.

- A major New York-listed software company purchased an innovative risk management firm with substantial revenue and more than 300 employees. Within a year, the CEO of the acquired firm had left and revenue in the product line fell by more than 50%.
- Another New York-listed firm acquired a company in the credit risk area with revenues of more than $50 million a year and staff of almost 200 people. Within two years, twelve of the thirteen most senior people in the acquired company had resigned.
- A New York-listed firm in the credit risk area acquired a small but highly innovative software firm with expertise in interest rate risk management. Within one year, the CEO of the acquired firm had resigned and within two years, the acquirer notified software clients of the acquired firm that the software product of the acquired firm would no longer be supported.

From the end user perspective, assessing the maturity, longevity and commitment of the management team of the vendor is the single best way to assess the long-term risk of the vendor. As the examples above show, large size can increase risk if the managerial compensation systems are not handled skillfully.

Being a Best Practice User

Once an installation is successfully accomplished, the vendor-user relationship is just beginning. There are a number of best practices that we feel will enable end users to get the maximum benefit from an enterprise-wide software installation:

- Build mutual knowledge with the vendor through a regular exchange of visits and information. These relationships should span many levels of management, not just be concentrated at the working level.
- Encourage continued innovation by the vendor by adopting new modules as they emerge and by being an active participant in the design of expanded functionality.
- Be an active participant in the user group.
- 'Invest' in the vendor by sponsoring new functionality with financial support to speed the degree of innovation.
- Establish a test database for regular testing of new versions from the vendor and make this test database available to the vendor.
- Regularly send staff to the vendor for advanced training.

In addition to these steps, we see an increasing trend toward long-term consulting contracts between the senior management team of the vendor and the most senior end-users at major financial institutions. This kind of relationship insures that, above and beyond the day-to-day working level relationship, senior management of the financial institution is fully exploiting an enterprise wide risk management solution to achieve the objectives that we have outlined in this book.

We turn to that discussion in the next and final chapter in this book.