

Risk Taxonomy: A Closer Look at Integration

The credit crisis has heightened the need for banks to analyze and manage risk in a unified and cohesive manner. To achieve risk integration, a bank must classify its different risks and then figure out how those disciplines interact and how they should be melded into an ERM framework.

BY WILLI BRAMMERTZ AND RAMI ENTIN

Historically, liquidity risk has been a primary concern of bankers. Venetian goldsmiths, for example, were already aware of the dangers associated with the paper money they helped invent, and lived in constant fear of too many creditors demanding gold for their receipts. Unsatisfied creditors often took action by breaking the goldsmith's table. Broken tables, or *banca rotta* in Italian, is a term that would not be unfamiliar to modern ears.

Credit risk was the other relevant risk for early bankers, since unpaid loans reduced their capabilities to redeem receipts. Consequently, the focus of banking regulation until recent times has been liquidity and credit risk.

More recently, with the advent of modern financial instruments such as options and futures, market risk has gained prominence. This was highlighted, in a regulatory context, in the 1996 amendment to the first Basel capital accord. Several years later, the treatment of credit risk was overhauled in the Basel II accord, which also introduced the operational risk category.

Lately, various regulatory bodies have proposed new ways to measure and manage the seemingly long-forgotten category of liquidity risk. What's more, parallel to these developments in banking supervision, the Solvency II framework for insurance regulation (which covers life and non-life insurance risk) has evolved.

Although the concept of integrating all types of risks is

not new, it received a strong boost from Basel II, which requires that the capital charge set aside to cover unexpected losses should include specific charges for market, credit and operational risks. The financial crisis fueled further calls for integration, which is now often referred to as enterprise risk management (ERM).

In this article, we discuss the when and how of risk integration, namely in which cases this makes sense and how this should be done.¹ We will proceed using a logical sequence, starting with market risk, followed by credit risk, insurance risk, operational risk and liquidity risk.

Market and Credit Risk


Market risk directly corresponds with the major asset classes of bonds, foreign exchange, equities and commodities. A financial contract is driven by sets of rules that determine which cash flows are exchanged under it, when these exchanges take place, and, in the case of contingent cash flows, how these cash flows should be determined. Market risk arises from unpredictable changes in risk factors effecting values and cash flows. Which risk factors affect which contracts depends on the terms of the particular contract. Exchange rates affect only contracts denominated in a foreign currency; commodity indices affect commodity-linked contracts, etc..

While market risk factors affect financial contracts, they are also mutually dependent. Although sophisticated mathematical descriptions of this interdependency are possible, a corre-

lation matrix is sufficient for most practical applications.


When financial contracts are well modeled, it is possible to calculate how a change in a market risk factor affects the value of any financial contract (as illustrated in Figure 1, below). What is true for a single contract also holds for a portfolio of contracts — or even for the balance sheet of a bank.

Figure 1: Market Risks and Financial Contracts



In reality, however, contracts' rules are not always adhered to, which gives rise to credit risk, as shown in Figure 2. In this diagram (see below), market risk factors are subsumed into a single category. Credit risk is determined by counterparties and their capability to repay obligations (including collaterals, ratings and probabilities of default), all of which affect expected cash flows. Market risk analysis must have precedence over credit risk, since it is not possible to model the expected loss on cash flows or a contract without first calculating these cash flows in the absence of credit risk.

Figure 2: Market and Credit Risk



Insurance Risk


Although the “natural” order of Basel II would dictate otherwise, it is insurance risk that logically follows as the next risk category. Insurance risk is generally split into life and non-life subcategories.² A life insurance policy or contract is — ignoring the probability of death during the coverage period — a

simple financial contract. Such contracts are very similar to common saving plans offered by banks. The only difference lies in the significant penalties that insurers can impose when premium payments are not met. This makes life insurance contracts practically non-callable, and, consequently, surrender rates are low.

From an analytical perspective, a life insurance policy is therefore a pure financial contract, and can be treated as such. The only real difference is the coverage of mortality during the term of the policy. Mortality is taken into account in the terms of the policy and — given that does not fluctuate in an unforeseen magnitude — it is not even a risk factor. In reality, the fluctuations are small, and mortality risk therefore occupies only a minor spot in the risk spectrum.

Non-life insurance is more difficult to handle, since neither the probability of loss events (frequency) nor their magnitude (severity) or payoff patterns (loss triangles) can be modeled to the same precision as in the life insurance case. That said, estimates for these distributions based on historical data do exist and form the input to the risk analysis (as shown in Figure 3, below).

Figure 3: Non-Life Insurance Risk



The non-saving part of life insurance can be viewed as a special case of non-life insurance where the frequency is relatively stable and well studied (mortality tables), and the severity and loss triangles are well defined (for example, payoff at death). It should be noted that these risk factors apply only to the liability part of insurance business. The investment or asset side of the business is no different than that of banks, and is exposed to the same market and credit risks.

Operational Risk

Operational risk (OR) gained in importance under Basel II, which requires banks to quantify it and set aside capital to cover unexpected losses arising from it. In parallel,

the concept of integrated risk management encompassing all risk categories has become fashionable, but with little understanding of how this should be achieved in practice. For example, there is no reliable model (let alone data) on the correlation of operational risk with other risk categories. Thus, in the absence of any better method or understanding, simple summation is generally proposed and applied.

While the “thou-shalt-integrate” instinct is fundamentally correct, integration should be done in a proper and logical way. A close examination of most OR systems shows that these systems are not risk analysis systems in the traditional sense of the word. The principal aim of OR systems is to eliminate or reduce risk, and therefore they are similar to management tools with an emphasis on workflow and control. Such tools typically define processes and identify key indicators, which measure their quality. Fire prevention, for example, requires among other things placing fire extinguishers at critical locations, making sure that they function correctly and that staff are trained to use them properly.


The classification of OR sources — keeping track of various remedial or preventive actions and measuring their effectiveness — is a core task of OR systems. It is indeed sensible to invest available managerial skills and resources in operational loss prevention, whereas modeling operational losses stochastically is meaningless.

Experience shows that managerial skills and techniques have their limits, and cost/benefit considerations therefore must also play a role. Figure 4 (above, right) shows the marginal operational loss given a certain level of resources required to mitigate it and the corresponding avoidance cost. As the resource level increases, the total loss decreases (at least ideally), but at a marginally decreasing rate. On the other hand, the marginal avoidance cost increases at an optimal point, where the two lines intersect. Investing fewer resources than the optimum leads to higher losses, while dedicating more resources is not worthwhile, since the cost of avoiding a loss is higher than the loss itself.

In Figure 4, the domain left of the optimum is managed with the workflow tools found in typical OR systems. It is commonly known as qualitative operational risk — or, as we prefer to describe it, the operational part of operational risk. To the right of the optimum is the residual loss after resources are optimized. This is known as quantitative operational risk — or, as we would also call it, the risk part of operational risk.

While the operational part has nothing to do with risk management proper, the risk part is comparable to market, credit or

Figure 4: Optimal Managerial Effort for Operational Risk



insurance risk. It is also the element that must be integrated into an overall risk framework. Basel correctly defines operational risk in this sense as the risk of loss resulting from inadequate or failed internal processes, people and systems, or from external events. This, however, should not take the focus from the main function of operational risk management: the avoidance such risks from the start.

How, then, should the risk part of operational risk be integrated into an overall ERM framework? The logical first step is to examine the loss database, where loss events (accidents, etc.) and the subsequent cash payments are recorded. This information can be used in two different ways:

1. For deriving key indicators. The number of losses and their severity is an indicator of the effectiveness of resources dedicated to operational loss prevention. Too many losses may be an indication that the resources spent on avoiding risk are left of the optimum.

2. For ERM input. A frequency (how often a loss event occurs) and severity (the cost of a loss event) distribution can be empirically derived for each loss category.

While key indicators are helpful in determining the managerial optimum, frequency and severity are the key factors that link operational risk management to integration. Frequency and severity are also used in non-life insurance for risk modeling; the only distinction is that non-life insurance additionally works with loss triangles³, which provide information about how payments are expected over time. Ideally, the loss database would also contain value and liquidity data and could be used as a basis for statistical analysis — but in practice this is not the case. Anecdotal evidence suggests that this is due to the low quality of frequency and severity information, where adding

loss triangles data would only lead to imprecision.

A firm needing to insure an operational risk faces a choice between managing it internally or externalizing it to an insurer. If the firm decides to buy insurance, the insurance company will need to analyze the risk using the data shown in Figure 3.

From its emergence on the risk radar, operational risk has been classified by many professionals as a different animal. One reason for this is the similarity of its quantitative part to non-life insurance risk, which is not a well-known concept to practitioners in the banking industry. Another reason is the statistical basis for the quantitative part compared with that of traditional market risk; operational risk is far weaker, because it is a special case at each firm and because operational loss events are few and far between. This is indeed why the operational part of operational risk management is much more important than the quantitative part.

There is, however, a more profound basis for the different nature of operational risk: At heart, analysis of market, credit and even life insurance risk is based purely on financial contracts.³ In an analytical system in which financial contracts are precisely modeled, their contributions to the overall risk profile can be calculated (for these types of risk) fairly accurately.

Operational risk, in contrast, does not have a financial contract at its center; rather it has at its core physical activity — namely, people, premises, machinery and processes. This is one of the primary reasons why the benefits of integrating operational risk are minimal. Since there is no connection between a fire drill and interest rates, for example, there is no benefit to integrating the fire extinguisher workflow into a risk management process. Although it makes every sense to integrate the output of a loss database into an ERM framework for analyzing the risk part of operational risk, one should keep in mind that it is minor compared to the operational part.

Liquidity Risk


Liquidity risk was the preoccupation of bankers up to the early 20th century, and was finally tamed with the emergence of central banks — the lenders of last resort. By the end of the 20th century, it was considered to be an extinct species, only to lash back at the beginning of this century. Analyzing how this risk fits in a unified risk analysis methodology therefore merits a discussion.

Liquidity risk comes in two forms: idiosyncratic (also known as funding liquidity risk) and market. Idiosyncratic liquidity risk is the classical risk of a bank run, observed so often in the 19th century. This happens when depositors lose confidence

in the ability of a bank to redeem deposits or, in general, to honor its liabilities. This risk is idiosyncratic because an individual bank is responsible for its actions and their effect on depositors' trust.

Idiosyncratic liquidity risk must be managed using classical asset and liability techniques. The liquidity gap (shown in Figure 5, below) is one example of such a technique. This analysis aggregates projected cash inflows and outflows over specific time intervals, where the vertical bars represent the net cash inflow or outflow at each interval.

Figure 5: Marginal and Cumulative Liquidity Gap



The solid line in Figure 5 shows the accumulated net cash flows. As long as their sum remains in positive territory, the bank is capable of surviving on its own, assuming no new business. When this sum becomes negative, the bank requires external help — for example, a bailout from the central bank or another government agency. Regulators demand a minimum survival time in order to have the necessary time to consider a course of action in such cases.

Beyond the simple gap analysis, there are additional factors to consider, such as behavioral assumptions for saving accounts or the quality of liquid assets held for such crises. The stressing and testing of such assumptions should be an integral part of liquidity risk analysis.

Even more than operational risk, liquidity risk is a different beast. Alas, the same urge to integrate appears to be at work here, making it appear that liquidity risk should be treated as just another risk category (as shown in Figure 6A). This, however, is a misconception. Although it is common to think of risk factors in terms of their value effect, idiosyncratic liquidity is a parallel analytical concept on the same footing as

value, and is affected by the same risk factors. It is therefore not a risk source comparable to market, credit, insurance or operational risk, but a risk effect (see Figure 6B, below).

Figure 6A: Liquidity Risk as an Independent Risk Category





Figure 6B: Liquidity Risk as a Parallel Analytical Output



Market liquidity risk, which arises when credit markets freeze due to a loss of confidence or widespread fear, is different. While a bank is responsible for its own idiosyncratic liquidity risk, the same does not apply to market liquidity risk, which affects every factor, regardless of its financial strength. Loss of confidence is manifested through increasing credit risk spreads, and it is this phenomenon that reveals the true nature of market liquidity risk. Indeed, market liquidity risk is a special case of credit risk where there is a general distrust of any counterparty partaking in the market.

To summarize, liquidity risk is not an independent risk category. Idiosyncratic liquidity risk is a risk effect, and arises mainly from reputational, or more broadly, operational risk. In particular, it is an analytical output on an equal footing with value, and should be derived in parallel with it. Market liquidity, at least from an analytical perspective, is a special

case of credit risk, and should therefore be managed and measured via credit spreads.

Final Thoughts

Market risk and credit risk analysis can be consistently performed using a calculation methodology centered on a financial contract. The same is true for life insurance analysis, and can be extended in a straightforward manner to non-life insurance by adding severity and loss triangles information.

Operational risk is less straightforward than it is usually presented in literature. Managing this risk is primarily done using classical managerial tools, and has little to do with risk analysis.

Liquidity risk is the most difficult to understand. This discipline covers two distinct types of risk: idiosyncratic (or structural) liquidity risk and market liquidity risk. Idiosyncratic liquidity risk is not a risk factor at all; rather, it is an analytical output similar to value, and since Idiosyncratic liquidity risk and value are affected by the same risk factors, they should be analyzed in parallel. Market liquidity risk, on the other hand, is an indication of widespread distrust toward all counterparties, and thus represents a special case of credit risk.

FOOTNOTES

1. For a more detailed explanation of how to build an ERM framework, see *United Financial Analysis: The Missing Links of Finance*, by I. Akkizidis, W. Brammertz, W. Breyman, R. Entin and M. Rustmann (Wiley, 2009).
2. Health insurance is often considered as another subcategory, but conceptually there is no difference between health and non-life insurance.
3. Loss triangles determine when the loss results in payment. In case of theft or fraud, the payment happened (involuntarily) when the act was committed and is therefore equivalent to a single cash flow. In case of fire, payments can stretch over a few years while settling a product liability lawsuit may take decades.
4. With the exception of the mortality risk of life insurance policies.

Willi Brammertz (PhD) co-founded Iris Integrated Risk Management, where he also worked as the CTO. In 2008, he sold Iris to FRSGlobal, where he currently serves as a senior risk advisor. He can be reached at willi@brammertz-consulting.ch.

Rami Entin (PhD) is a senior consultant at FRSGlobal, where he works on integrating financial risk analysis with regulatory reporting. He has worked extensively on Riskpro, a unified financial analysis application that was originally developed by Iris. He can be reached at rami.entin@frsglobal.com.