Four key observations about calculating Pillar 2 operational risk capital

Executive summary

For many buy-side firms, the challenges of operational risk exceed their analytic bandwidth. Operational risk is inherently complex and experienced executives know this. Financial firms spend more effort on risks within their core competencies, especially market and credit risk. However, the reality is that, for many buy-side firms, operational risk is their biggest class of prudential capital. Failing properly to address analytically operational risk simply leaves money (in the form of unnecessary prudential capital allocation) on the table. It also raises the risk of supervisors addng to a firm's capital requirement via individual capital guidance.

Using an approach of expert estimation of plausible scenarios for operational loss provides both analytically meaningful and managerially useful information for financial firms. This requires estimation of (i) probability of event occurrence; (ii) expected loss levels and (iii) worst-case loss levels; as well as (iv) the interaction between risks, i.e. correlations. With this minimum information, using an appropriate analytic tool, firms can:

- improve their understanding of operational risk analytically
- eliminate unnecessary prudential capital provisions for operational risk
- derive a meaningful and defensible number for prudential capital requirements for operational risk that will satisfy regulators
- provide important and useful management information to support decision-making around risk transfer and investment in other operational risk management strategies.

By using expert estimates supported by whatever data firms may hold (e.g. partial loss history, incident data, risk and control self-assessments) and a cost-effective, purpose-built scenario analysis tool such as MC+, buy-side managers can improve the validity and utility of their operational risk analysis, eliminate unnecessary capital allocation for operational risk and reduce the complexity of analysis required for Pillar 2 operational risk capital.

Observation 1

Probability of an event occurring and size of loss given the event occurring are logically and statistically separable (as for credit risk)

Most operational risk applications generate a probability number – just one number – but that doesn't work for operational risk. Operational risk is analogous to credit risk in that the probability of occurrence (or default (PD), in credit terms) and size or value of loss (loss given default or LGD) once the event occurs are logically separable. Also, as with credit risk, there are different probability structures or distribution functions for each element.

In operational risk, probability of an event occurring is best fitted using a Poisson distribution. The probability of loss curve is essentially empirically derived. Experience shows that a long-tail distribution such as lognormal fits most data. For each of these distributions, minimal data points are needed to describe a distribution function that can be used for simulation (see Figure 1, overleaf)

Observation 1, cont.

Unlike using a single operational risk distribution, using separate distribution functions for event occurrence and loss given occurrence is logically and analytically robust. Combining stochastically (using a powerful random number generator) a Poisson distribution for probability of event occurrence and a lognormal function for loss given event provides a robust distribution of operational risk outcomes.

Because it is logically and analytically robust, it focuses managerial attention more appropriately than less robust methods: that is, the key elements in defining risk capital are probability of event, expected loss values (which management can reduce) and worst-case scenarios which focus managers' attention on exposure and defending against very-low-probability events with potentially high resultant losses. These properly-specified loss scenarios drive managerial attention without over-stating capital impact.

FIG. 1

Separate distributions for likelihood of occurrence and loss given occurrence

... combined through MC simulation to give modelled losses

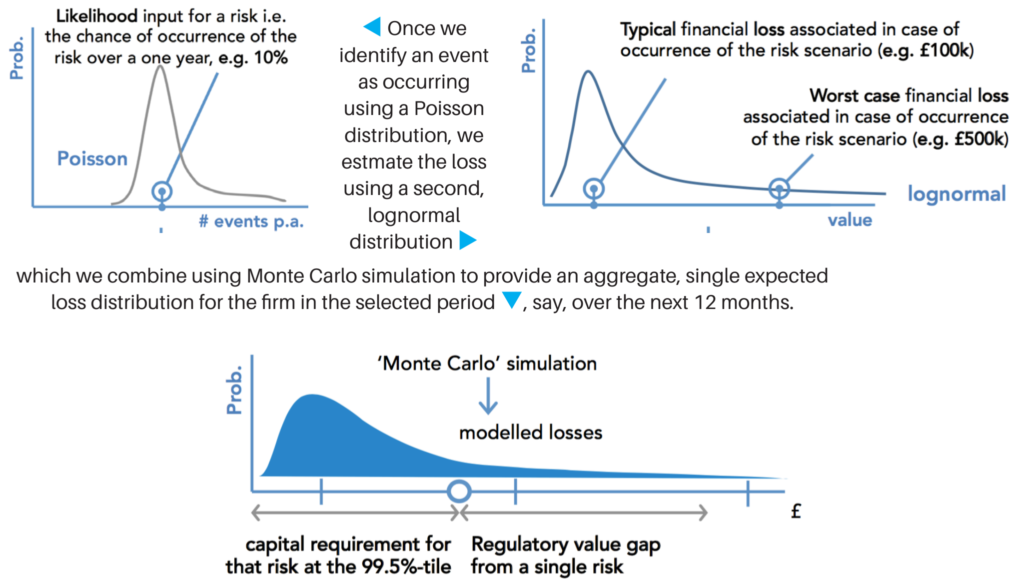

Observation 2

By failing to address the real structure of operational risk exposure and the interaction between different risks, most firms 'leave money on the table' in the form of unnecessary allocation of prudential capital for operational risk.

Many firms simply add scenarios of losses from operational risk to arrive at a number for prudential capital required to support operational risk. But that implies that an event of every loss type will occur within its specified probabilistic repeat period. Because both distributions of probability of occurrence and loss given occurrence are skewed, that is simply not the case. Adding risks also fails properly to reflect the interdependencies between different operational risk scenarios or events. By building in realistic interdependencies between risks, many of which will be independent of each other, the prudential capital required to support operational risks assumed by the firm reduces. This diversification effect is an essential element of any realistic analytic framework for operational risk.

See Figure 2, overleaf.

FIG. 2

Example of a correlation matrix using qualitative scores

	Risk 1	Risk 2	Risk 3	Risk 4	Risk 5
Risk 1	1				
Risk 2	L	1			
Risk 3	L	H	1		
Risk 4	L	M	M	1	
Risk 5	L	VH	H	M	1

EXAMPLE CORRELATION VALUES

No correlation	0.00
Low	0.10
Medium	0.25
High	0.50
Very high	0.75

Observation 2, cont.

The key difficulties in modeling operational risk are estimation of maximum foreseeable loss exposure (i.e. worst-case losses) and adequacy of empirical information around interdependencies or correlations, both of which can be estimated by the firm's subject matter experts or advisors. Given the lack of applicable data for most firms, banded qualitative estimates of correlations (L, M, H, VH) suffice. In any case, simply by encouraging managers and analysts to focus on co-dependencies and relationships between risk scenarios, consideration of correlations improves understanding of the firm’s risk environment.

Observation 3

Complexity is inherent in operational risk. But that does not mean robust analysis of operational risk needs to be complicated. The right place for the complexity is to embed it in a suitable analytic tool.

Because the operating environment of firms is inherently complex, so is properly-framed analysis of operational risk. However, that does not mean it needs to be complicated. By focusing on the key data elements required to define the distributions for each risk type, and simulating using event and loss distribution functions, a well-structured risk tool can embed that complexity while requiring only a simple specification structure from the user. This makes sophisticated modeling of operational risk through definition of plausible worst-case loss scenarios both efficient and meaningful. The few feasible alternatives are slow, usually analytically questionable and difficult to operationalize.

Observation 4

Because most current operational risk analysis focuses on what we already know, much operational risk-related effort is wasted. For operational risk analysis to be useful, it must be prospective and readily re-specified to support real-time decision-making.

That requires the right data . . .

Since the inclusion of operational risk requirements in BCBS rules and the implementation of operational risk (often GRC-type event capture tools) systems in many financial firms, attention has focused on what we know: observed operational events and losses experienced. These are usually supplemented by well-reported losses in other firms. However, the utility of other firms’ experience depends crucially on understanding the scalability of their losses to the subject firm; that is, how relevant is their experience to your firm? Since the early 2000s, for large, international banks, such losses have been captured systematically but differing underlying processes and business models mean their applicability to buy-side firms can be limited and values irrelevant.

Observation 4, cont.

By using known losses, whether from event-capture tools or value-driven recording in specific general ledger codes, firms focus only on expected loss. This is critical knowledge and an essential modeling input, but does not relate to potential worst-case losses; they can only be related distributionally. However, incorporating estimates of worst-case losses without the separate event probability / loss probability distributions distorts results; it simply fails to reflect reality which, intuitively, most users of such information recognise.

As stated above, the answer is separate probability distributions for occurrence or incidence and impact or value of loss experienced (which, for some classes of risk may, meaningfully, be negative – i.e. unexpected gain). This approach uses both expected losses from captured or known data and subject-matter experts' estimates of feasible worst-case losses. However, without a suitable tool for such analysis, it rapidly becomes unwieldy and unfeasible. A purpose-built tool for handling robust scenario analysis using distinct probability functions for occurrence and loss with minimal data specification requirements provides a feasible and operable solution. MC+ is such a tool.

. . . and rapid analysis . . .

To be useful, calculation based on specified inputs must be immediate. Often, complex simulations can take half an hour to run. However, even with up to 25 scenarios specified, MC+ recalculates within 30 seconds. This enables it to be used to support live decision-making. With immediate recalculation, investment decisions in methods and technologies to reduce incidence or impact are feasibly supported in real time. This means managers can understand the economic impact of their decisions to invest in, for example, risk transfer contracts for cyber insurance based on a realistic assessment of the costs and benefits of doing so; also, the actuarially fair price is far clearer – that is, the benefits that have to be secured for the quoted costs of insurance.

. . . to pass the supervisor's use test.

By supporting management decision-making in real time, operational risk modeling based on plausible scenarios becomes useful. Although, of course, it is up to implementing managers to determine whether or not this analytic utility is applied, supervisors will be able to see its use in real decisions, validating utility and meeting the supervisor's use test.

As a stand-alone system, scenario-based operational risk modeling can supplement any existing data capture or analysis approach. With low entry costs, the approach offers exceptional value for money by quantifying more robustly prudential capital requirements for operational risk.

Contacts

To find out more about technical solutions to operational risk or using scenarios and Monte Carlo simulation for calculation of capital requirements for operational risk, visit us at

www.montecarloplus.com

Mustafa Çavuş PhD
Founder & Director

e. mustafa.cavus@montecarloplus.com
m. +44 7432 633 984

Peter Bonisch
Director

e. peter.bonisch@montecarloplus.com
m. +44 7974 168 559