GRAS Newsletter 12

July 2009

Summary
Quantifying OpRisk has proven to be an elusive sport. The various attempts, however, hold an important lesson, namely that quantifying, and by extension also *managing,* both day-to-day and extreme losses requires specific methods. This also applies to the basic governance of ORM which should acknowledge the limits to what extent line management can manage ORM by itself, and where an independent ORM function should step in.

Dear reader,
The nature of OpRisk is that it covers a broad area. It is restricted to neither specific functions, products, activities, client groups or any other specific part of the bank. The standard practical solution for the *management* of OpRisk is that it has been made part and parcel of the line managers' responsibility, coupled with a central unit to set methodology and to compile reports. It is not uncommon for banks to adopt a guiding principle for ORM, which is often a variant of the following statement: "Management at all levels is accountable and responsible for the management of operational risk in their area". Sometimes this even broadened to "ORM is everyone's task", thus making it a universal responsibility.

There are clearly areas of risk where this approach of line management responsibility for ORM works well. Identifying inefficient, problematic and risky procedures, malfunctioning systems or failing levels of product handling is quite within the capabilities of management and staff. By definition, the day-to-day processes, activities and situations are their responsibility and hence, so are the day-to-day associated risks. Where this principle brakes down is where the risks (and indeed the losses that occur when these risks materialise) are not of a day-to-day nature.

The lesson from the modelling effort
The attempts to model OpRisk have met with mixed results. The issues with completeness of data collection, clarity of definitions and demarcation of categorisations are well known. Nevertheless, we can actually see some remarkable successes in modelling OpRisk losses. For many banks, the recorded loss events display a consistent pattern *provided they are restricted to the day-to-day events*. If we exclude the large losses as 'one-off' events, then the remaining losses are statistically reasonably well-behaved and have some predictive capability. Examples of such predictable losses are credit card frauds, teller mismatches, petty theft, discretionary client compensation and even branch hold-ups.

Unfortunately, these losses are completely irrelevant for the kind of quantification that AMA capital models require. Here, the banks must make an estimation of the total yearly loss amount that is not likely to be exceeded more than once in a thousand years. And for that, these losses are as good as irrelevant.

Capital implications
It has been known for some time now that this calculation is dominated by a few single extreme events, which fully determine the capital calculation. It not uncommon to find one or two large events whose effects are several orders of magnitude higher than the next lower loss amount, especially if we include external loss events into the equation. The most important observation to come out of the analysis is of the loss distribution is then that a difference must be made between 'day-to-day' losses and extreme losses.

Extreme losses
The circumstances that give rise to 'day-to-day' losses differ radically from the circumstances that give rise to extreme losses. Extreme losses do not occur in the same pattern, either from a statistical point or view or from a causal point of view. Part of the reason why they do not appear in the same pattern, or in any pattern for that matter, is that they never fall within a bank's tolerance level. Unlike the day-to-day losses, where a trade off between the cost of preventing events and the cost of accepting the risk is an actual management choice made time and again, extreme losses (and the kinds of risks associated with such events) carry no such trade off. Hence, these risks simply do not feature in management's decisions, since there is nothing to accept. There is no appetite, and also no tolerance at all for these kinds of risks.

One solution the modellers have investigated to allow for the specific nature of extreme losses is to use methods from Extreme Value Theory, such as modelling all losses over a certain amount in a separate model[1]. So far, however, no dominant approach has emerged for modelling the extreme loss amounts. One lesson from modelling OpRisk loss data is, however, beyond dispute: OpRisk losses can not be modelled with a single distribution.

The implications for OpRisk *management*
So how does that observation help us in OpRisk *management*? In OpRisk management, our objective is to manage operational risk to levels that they lie within our tolerance levels. The comment above on modelling losses carries over to OpRisk management in the following fashion:

> The circumstances that give rise to 'day-to-day' events differ radically from the circumstances that give rise to extreme events and hence, our approach in managing these risks should reflect that difference.

No serious bank uses a single tool for ORM, but the implications run deeper than just the tools: it also affects the governance of ORM. The notion that ORM can be managed solely by line management works well for 'day-to-day' risks but does not hold for risks that are beyond the day-to-day occurrences.

[1] The idea of using PoT (Peaks over Thresholds) lead to some early success. Current thinking, however, is that modeling OpRisk loss data, especially with an eye towards prediction, is fraught with so many complexities that simple statistics and even extreme value theory does not suffice.

Governance

It is worthwhile to notice that the most extreme headline-grabbing losses (SocGen, Barings, UBS, Allied Irish, etc) are all cases where line management is either complicit or incapable. Neither situation is likely to benefit from the traditional ORM approach of relying on line management to carry out ORM.

In most banks, the governance of ORM is characterised by a Senior ORM committee, a central ORM unit with (part time) ORM officers in business units who report in the line. It is not uncommon to find a rather weak ORM unit and only a few business units that take their ORM tasks to heart. This situation is a natural consequence of the focus on 'day-to-day' risks and over-emphasising line management's responsibility for ORM.

ORM Review

One way of managing beyond the day-to-day events is to upgrade the central unit's responsibility and to create an independent ORM review role. The review role in credit risk has been well established and has proven to be an effective way of uncovering errors and omissions that crept in the credit practice. It basically allows the bank to look beyond the line-management's blinkers. That can also work for ORM.

The ORM review would also ensure a closer adherence to the letter and the intent of the existing ORM programmes across the institution. It would create a natural platform for the central ORM unit to gain deeper insight in what programmes work well and what programmes are performed as a mere compliance exercise. More importantly, however, it will allow the central unit to look beyond 'day-to-day' risks and it can overcome line-management's blind spots.

Note that the central ORMs risk review function does not diminish line management's responsibility or accountability. Far from it. In all head-line grabbing extreme events, the line management was accountable and responsible. But line management also failed spectacularly to manage its risks by itself. The reason that failed is the same reason why ORM needs an independent review function. Line management was either incompetent or complicit. Neither situation can be remedied without an independent function that has the authority and wherewithal to spot these risks before they become events.