

Unless operational risk management (ORM) re-invents itself, it is doomed to fade into the background. Now is the time for ORM to make a bold move that goes right to the heart of risk management. **Marc Leipoldt** outlines what to revisit, what to discard and what to adopt to make this happen

Winning hearts

The commemorative January issue of *OpRisk & Compliance* listed an impressive array of professionals who have shaped the operational risk management (ORM) discipline. Their views on the state of ORM are varied. Coming from regulators, banks' risk management departments, insurance agencies, academia and assorted consultants, this is to be expected. One overriding element in their comments, however, is on the need for ORM to pick itself up and move on. This article is an attempt to outline the new roads that warrant further attention, as well as the dead ends that should be abandoned.

The amount of criticism levelled at ORM in many institutions is grounded in a misunderstanding engendered by ORM itself. This misunderstanding started with the way operational risk was introduced in the capital requirements regime. Taking a cue from the now maligned value-at-risk (VAR) method, the most advanced methods of capital calculation for operational risk had to correspond to a VAR at the 99.9% confidence level interval.

This sent the ORM functions scurrying for data to compute this elusive number, using internally gathered and externally acquired data. In doing so, the core of operational risk, namely insight into the banks' processes and the underlying causes for events, as well as preparedness for the future developments in the wider environment, played second fiddle.

The remainder of the article sketches the directions ORM should take in financial institutions, which includes some common practices that must be abandoned or at least drastically reformed. The refocus can free some resources as well as redirect existing programmes for many banks. Given the ORM tasks

at hand, the upshot is, however, likely to increase the workload of ORM departments, business units, internal control departments and audit. This is a direct consequence of the fact that sound ORM is a labour intensive, bank-wide activity.

Ever since the formalisation of ORM, senior management of various banks have been wired to ask "show me the money". This often left ORM staff stammering about better processes and higher levels of prevention that could be achieved. This is of course a highly unglamorous business. As Hank Paulson is alleged to have said, you don't get much credit for averting a disaster.

Even in banks that started out promisingly, senior management quickly dismissed ORM as a compliance exercise at best and reduced their support to the bare minimum. It will be interesting to see whether banks, which have all pledged to improve risk management and all claim to be increasing the levels of accountability and control, are indeed willing to put their money, time and attention where their mouth is.

As a starting point, here is an enumeration of which elements to reconsider, what to discard and what to adopt.

1 REVISIT the remit of ORM: it must become central to the bank's risk management.

In most if not all banks, the rise of ORM has been a by-product of a Basel II programme. As a consequence, it has focused on Pillar I capital calculation rather than *managing* the risk due to "inadequate or failed internal processes, people and systems or from external events", to use the Bank for International Settlements' (BIS) own definition of op risk. In practice, the remit of ORM has been



ORM must be the heart of risk management and thus the starting place of any other risk analysis, be it credit risk, market risk, liquidity risk or any other risk

much smaller than this broad sweep and has actually shrunk over the last three years.

First of all, established risk management practices such as credit risk management and market risk management were virtually out of bounds for ORM work from the start. Instead of recognising these activities as processes, people and systems that should conform to ORM practices like any activity, it was often assumed that, because they dealt with a specific type of risk, they were exempt from the ORM programme.

stockphoto



Second, while ORM was finding its feet and establishing an appropriate governance structure, further pieces started to disappear from core ORM units. Two dominant examples are business continuity management (BCM) and information security (IS), both of which required technical specialists, who quickly set up shop with only a loose connection to risk management. Because neither BCM nor IS plays any role in Pillar I, this hardly registered with the decision-makers in the banks.

Lastly, ORM had to struggle to find a place among such disciplines as audit, internal control, quality assurance/process improvement units, business process re-engineering, documentation, accounting/finance and enterprise risk management. ORM has dependencies and overlaps with all these disciplines. The difference being that ORM was ill-defined in its programme and so found it difficult to show to the business what value it could add. In addition, it seemed that as soon as ORM did transform itself into

a more concrete set of rules, that part was prone to be lopped off and to continue as a separate function.

To rescue ORM from obscurity and eventual oblivion, it is necessary to re-establish its remit. The basic definition is just fine, if it is combined with a set of policies and procedures that allow ORM to exercise the role of a proper risk unit.

Using the BIS definition of “the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events” and ignoring the unhelpful additions, we get back to what ORM is supposed to do. It should employ a diverse set of tools to identify risks, analyse the underlying causes, gather information on events and scenarios, analyse events and scenarios, establish appropriate responses and track outstanding actions on these responses.

By focusing on internal processes, it is immediately obvious that ORM must be active at all levels. Processes at the head office in a firm are not aggre-

It will be interesting to see whether banks, which have all pledged to improve risk management, are indeed willing to put their money, time and attention where their mouth is

gates of processes at branch level: they are *different* processes. The processes at senior level are not the sum of processes at lower levels: they are *different* processes. Both have their own controls and risks.

If it is accepted that ORM must be implemented in all parts and at all levels in the bank, then two corollaries follow. First of all, ORM requires active co-operation by all. Second, it follows that ORM must be the heart of risk management and thus the starting place of any other risk analysis, be it credit risk, market risk, liquidity risk or any other risk. Part of the argument for this is grounded in the two areas to be strengthened or added to ORM, namely the role of uncertainty and the notion of causal analysis (see items 5 and 6).

This all has implications for the governance structure of ORM. Currently, the use of local ORM contacts or co-ordinators is widespread. They are typically weak powers in the organisation. Likewise, strong central ORM functions are even more rare than modest traders. ‘Strong’ would here mean capable of setting enforceable policies and standards, receiving relevant and timely data as well as being equipped to perform analyses and give advice for all business units.

2 REVISIT risk and control self-assessment (RCSA) and key risk indicators (KRI): these tools must be saved from the modellers and the capital-obsessed.

There are many misunderstandings about RCSA. Let’s start by splitting the RCSA into RSA and CSA: risk self-assessment and control self-assessment. Both are assessments but if implemented well, they are fundamentally different in their approach. In both cases, we are looking for weaknesses and vulnerabilities and, in both cases, we hope to arrive at meaningful actions that will address this. But whereas in CSA the analysis starts with the list of *known* controls, which are by definition limitative, in RSA we focus on risks, which are multifarious and in principle unlimited and uncertain.

Hence, the RSA should adopt an approach that

allows for as yet unspecified risks to be assessed, whereas CSA can work with a given set of pre-established controls. Many banks start out their RSA implementation from a blank sheet, allowing units the maximum freedom to identify risks and work from there. Because the pilot implementations and the first wave are mostly done in friendly units that are open to the idea of ORM, the outcome is successful at the unit level.

Most banks, however, cannot resist the temptation to make this cumbersome process 'more efficient'. They will, for example, present units with a fixed list of risks to assess, or they will eliminate all creativity in the identification process by merely asking them to provide numbers on impact and likelihood. When this happens, the effect of the RSA will approach zero faster than you can say rainbow chart and the programme will not survive in a meaningful way. The early enthusiasm is not carried over, and an opportunity to learn about risks is lost as the RSA rapidly devolves into box-ticking.

This sorry state combines with the desire for aggregate information by those who want to manage without truly becoming involved. This leads to an increased demand for numbers, tracking only the impact and likelihood outcomes over time and comparing them across business units or processes. This is the final blow for RSA, because it turns an open-ended discovery process into a measuring contest. The assessment has become worthless.

For CSA, box-ticking is pretty much where it starts. The nature of CSA, however, allows for a much higher frequency of polling than RSA, which does allow for a different kind of analysis. This alone would be an argument to separate RSA and CSA.

Lastly, one thing we can say for sure about assessing risk is that people are notoriously bad at it. Academic research has shown that, when asked to estimate risks, people typically overestimate the likelihood of rare events, and underestimate the likelihood of frequent events. Other peculiarities such as anchoring and adjustment, the base-rate



Marc Leijpolt

There have been suggestions that KRIs should predict losses. This is fundamentally wrong. KRIs should be an indicator of risk, not loss; the difference is important

fallacy and group think should caution us against the use of these assessment outside their context and in other than a qualitative sense.

As far as key risk indicators are concerned, they have been misunderstood in some cases. There have been suggestions that KRIs should predict losses. This is fundamentally wrong. KRIs should be an indicator of risk, not loss; the difference is important. For KRIs to be early warning indicators, they should register increases in the *likelihood* of events *unless* immediate action is taken.

Some have suggested vast lists of standard KRIs, and the jury is still out on the usefulness of that. In practice, the KRIs that work are specific, frequent and comparatively few but highly effective. They are known as operational controls and no process owner needs to be convinced of their value.

Lastly, follow up actions and monitoring thereof must be part of the programme. Failing that, the programmes can rightly be said to have no effect and could even obscure the degree of risk and control we are seeking to put in broad daylight.

3 DISCARD the loss distribution approach (LDA) for the advanced measurement approach.

The LDA is an attempt to create a value at risk (VAR) for ORM based on internal loss data, supplemented with external loss data and further augmented by scenario data. The problem with the approach is that VAR was meant for situations that behave in a statistically sensible manner, that is, there are many observations of equal validity, more observations lead to better modelling, data is time-invariant and is well classified. All of these characteristics are absent from the kind of loss events that interest op risk managers at the tail. To make matters worse, in the world of credit and market risk, VAR has been largely discredited as a meaningful predictor of losses. In fact, it has been compared to an airbag that works 99.97% of the time, except on impact.

For ORM, the very losses that drive the capital are those that are extremely unlikely to re-occur. Any chance of another Rusnak at Allied Irish? Or another Kerviel at SocGen? Also, the use of external loss data brings its own issues. The use of external data quickly takes away the incentives for ORM implementation once business owners find out that they have no influence over these external events. Lastly, the so-called qualitative adjustment is often just that, an afterthought that barely registers in the final outcome and is not sufficient to stimulate unwilling business managers.

Rather than tinker with the LDA, it would be preferable to focus on the results of causal analysis and scenario analysis to quantify the amount of op risk and hence capital requirements for a given business.

4 DISCARD the notion of risk appetite for operational risk.

Risk appetite, defined as addressing the policy question of how much risk exposure the bank is prepared to carry, has almost no meaning in op risk. In principle, op risk carries no upside, and risk appetite is inextricably coupled with the risk-reward

trade-off. In op risk, this trade-off does not exist in a material sense. There is no real trade-off between, for example, tax-avoidance, forgoing segregation of duties, failing to install a proper fire wall, not training staff, misinforming management, churning accounts, mis-selling or kiting, on the one hand, and the rewards that these risks might bring. Without that trade-off, appetite in the classical sense is meaningless.

A far more useful concept to use is that of 'tolerance'. This could be introduced by assessing the status quo, and determining whether the situation is tolerable, that is, whether it needs imminent improvement or not. Here again, the inability of people (and even experts) to assess risks will continue to plague a scientific approach. The assessment will probably not be a number or a formula, rather it is the outcome of causal and scenario analysis on a process or business level coupled with the owner's limit of what the unit can bear.

5 **ADOPT** a taxonomy for uncertainty into the ORM toolkit.

One of the blind spots in risk management is accounting for uncertainty. This might seem anomalous, but the untrammelled faith in the models used, ranging from the context in which the models operate and the inputs in the models to the relations of variables in the models and the interpretation of the outcome of the models to the externalities assumed, carries a degree of uncertainty that is not addressed.

One of the reasons this uncertainty is not in focus in the risk world is the lack of an accepted language with which to discuss it. One area of research where this problem surfaced early, and which has developed an initial taxonomy for uncertainty, is policy analysis. This discipline typically is engaged in evaluating ill-defined, complex, long-term projects with major social and economic consequences. It could be the question of how to protect a delta against floods, which type of medical research to sponsor or how to reduce income inequality while raising general welfare.

One of the blind spots in risk management is accounting for uncertainty. The untrammelled faith in the models used carries a degree of uncertainty that is not addressed

An attempt to unify the many frameworks used for uncertainty lead to three dimensions: object, level and variability. This has proved to be a useful way to communicate about uncertainty. The 'object' can be a model, an input, a relation between variables, a parameter or an external circumstance. For each of the objects in an analysis, a level of uncertainty is evaluated, which might of course vary up or down over time.

Statistical uncertainty is one possible 'level' of uncertainty for any object we care to introduce, but so is scenario uncertainty or acknowledged ignorance. Highlighting the level of uncertainty (which is a continuous range from completely deterministic to total ignorance) is only one aspect. In terms of uncertainty, there is also some stickiness associated with the level of uncertainty, called its 'variability'. This variability (or stickiness) of the level tells us whether the level of uncertainty associated with the object is fixed (such as the statistical level associated with the roulette wheel) or has potential to change over time (such as the credit history of a new client, which might move from scenario uncertainty to statistical uncertainty with enough effort of data gathering).

This language enriches the risk field for banks enormously and forces risk managers and decision-makers to include more than just the statistical analysis.

6 **ADOPT** causal analysis and scenario analysis as the basis for ORM.

The BIS's focus of classifying losses by event type alone has had unfortunate consequences. It suggests that these event types are the starting

point for analysis, and regrettably they are not. It is the causes underlying the events that matter, not whether we spend hours debating whether a squirrel chewing through the fence leading to a power outage, followed by a client using the confusion to misappropriate a printer, is an external event or an external fraud. Or whether you reclassify it as internal fraud once you find out that the client is actually an employee of another branch.

What matters is that, to lower the risk of re-occurrence, a causal analysis needs to take place, which requires an investigatory approach. The language of causes is not that of financial accounting, which is all too often all that happens when a loss is reported. Many banks do require extensive reporting for their more extreme losses. That is in itself a good development, which needs to be extended to other events. Not all events require external auditors and extensive swat teams, but an event that is not analysed is an event the bank learns nothing from. A specific ORM-driven tool should facilitate the causal analysis, which must be mandatory part of the loss data policy.

Conclusion

Operational risk management has every chance to reassert itself in the current crisis. To do that, it must assert its remit as core to risk management, reaffirm the role of the classic ORM tools, shed some of the traditions that have not delivered sufficient value (such as LDA and risk appetite), and must give causal and scenario analysis a more central role in its approach. If it does that, there is much to gain for banks in boosting ORM. ■

Marc Leipoldt is managing director of Global Risk Advisory Services. Since July 2008, as managing director of his firm, he has been dedicated to operational risk governance, monitoring, analysis, review and implementation consulting. Previously, he spent five years developing internal ORM and an advanced measurement approach at Dutch bank ABN Amro. Email: marc.leipoldt@globalras.com