

Risk Management Through Risk Concentration Analysis (RCA)

Risks are best analyzed in the context of all other risks.

All enterprises manage risk on a daily basis. For corporate executives, there is a constant balancing act that involves deciding which risks are benign and which require controls. Risks may appear in any corner of the enterprise, including, but not limited to:

- Defects in products and services
- Irregularities in the value and supply chain
- Inconsistencies with compliance
- Issues with currency, and
- Problems throughout operations.

Decades of focus on business continuity risk, particularly in financial services and large IT-based businesses, has resulted in a tendency to look at risks based on their causes. Once selected, risk managers implement controls that compensate for each isolated risk deemed significant. While this approach to risk management has successfully helped prepare for natural disasters and similar events, isolated analysis fails to recognize risks underlying other common failure scenarios.

Isolated risk analysis doesn't provide an efficient way of looking at the larger picture. It evaluates each standalone tree but ignores the much larger level of the forest. Additionally, isolated analysis lacks consistency and the ability to assign relative weights to risks, making it hard to compare risks identified in seemingly unrelated areas.

A much more effective approach is to consider all risks simultaneously, their relationships and dependencies, and their potential impact on the enterprise. This white paper provides an overview of an approach based on Risk Concentration Analysis (RCA).

Interestingly, the most important relationship isn't always a common cause. If a collection of seemingly unrelated causes can have a common result, there is a concentration of risk associated with the effect that needs to be managed. Organizations spend resources to prevent bad effects; causes only matter if their effects aren't well controlled. Once a ship is taking on water, the fact that a different cause was mitigated is of little comfort. Making a stronger hull becomes a more important task than trying to individually control each potential cause individually.

This aggregated tonnage concept of risk applies perfectly to the enterprise. Common risk dashboards classify risks (generally by cause) and light up areas with "red" to indicate a high risk exists. What's overlooked is that an area lit with many "yellow" risks might actually have a higher total risk.

Do 100 "yellow" risks combined add to something worse than a single "red" risk? They might – particularly if they are concentrated in one place. The place might be a cause, an effect, a product line, a program, a geography, etc. The good news is that a common concentration yields an ideal opportunity to determine their collective materiality and efficiently control the group of risks.

What this means is that risks are best analyzed in the context of all other risks. Executives concerned with managing risk are advised to evaluate the total existing concentration of risk. Evaluating concentrated risk provides:

- A complete view of risks as they combine and interact,
- An efficient way to control collectively unhealthy risk groups, and
- A view of how actions affect an existing risk concentration.

Examples

Counterparty Risk - A well-known financial institution had a substantial risk concentration in counterparty transactions. While each isolated transaction likely made sense (and was constantly undergoing scrutiny by the institution's risk management team), the aggregate risk level did not. When combined with another substantial risk concentration (collateralized debt) a significant, unforeseen event was likely to occur.

As real estate values decreased, the firm appeared to be less strong. This triggered the counterparties to demand more certainty. The level of counterparty exposure was greater than the institution's ability to guarantee, and the institution failed. Isolated individual transactions were managed to acceptable risk levels, but risk concentrations were not.

Manufacturing Risk - A global auto manufacturer had a defect in its electronics that could result in unintended acceleration. An unrelated cause (sticking in the throttle linkage), could potentially create a similar result. A third issue, floor mats that could jam the accelerator pedal, created a risk of the vehicle continuing to move even when one's foot was removed from the accelerator.

Each of these three risks taken in isolation was evidently determined to be low, with little risk of causing substantial problems. However, the total collection of risks from these different causes (associated with a common result), makes it clear that there is a substantial aggregated risk of unintended acceleration.

Materiality

If risks should be evaluated in context (against other risks, against themselves over time, etc.) a new framework is needed to enable this analysis. The first part of that framework is a common scoring mechanism. Scoring must be consistent and allow comparisons over time. This repeatability is extremely important.

Equally important is a way to score disparate risks so that they can be meaningfully combined across enterprise silos. Risk “materiality” is a valuable way to accomplish this. Materiality is generally considered to be the probability of an event times its impact cost (if the event should actually occur).

In reality, a single risk often has multiple impact types with different costs (financial, legal, reputational, regulatory, etc). This drives the need for multiple materiality scores per risk instance. Further deconstructing the probability into its components of exploitability (likelihood that the environment will attempt) and susceptibility (likelihood or level of success if an attempt is made) creates a highly accurate materiality score.

Once a risk has a set of impact-based materiality scores, it needs to be classified in multiple dimensions. The idea is to test for concentrations of risk in each class. Causes are the most common (aligned with exploitability), though effects can be equally enlightening. In fact, the more classes, the better: Classes can show potentially harmful concentrations of a risk materiality score in a specific product, in a business process, in the results of an action or transaction, in a geography, in a customer segment, etc.

Further, each classification improves the ability to find harmful risk concentrations and to apply fact-based tests of potential controls effectiveness.

Challenges and Solutions

Normalization – Providing a centralized view of risk materiality concentrations by class is extremely valuable, but it is challenging. Aggregating data from disparate sources is first a process problem and then a technology abstraction problem. Data from very disparate sources must be normalized for materiality to be calculated in a way that allows concentrations and comparisons.

Taxonomy – This begins with the fundamental definitions of risk components, so that terms such as "high impact" in one area is equivalent to "high impact" in another. Taxonomy is an important prerequisite for implementation – without it, the normalization effort becomes far too complex for real-world implementation.

Impact Levels – With the data aggregated and normalized, we can then score materiality for each type of impact. At this stage it is vital to consider the importance of each asset; the same risks operating on different processes, product lines, or people can have completely different impact levels. This requires understanding the role of an asset as a critical component of materiality scoring.

Data Quality – Data quality concerns are valid, but often misunderstood. When viewing risks in isolation, the accuracy (or precision) of the data demands focus. When volumes of risks are aggregated, the importance shifts from precision to the existence of bias in existing errors. Statistically, noise without bias can self-cancel over large data sets, while biased noise will reinforce an incorrect conclusion.

Visualization – A way to interactively test and present risk concentrations is needed. This suggests a visualization engine against which the user can isolate risk concentrations, find key drivers, and test the effects of potential controls or actions/transactions. It requires much more advanced visualization tools than simple static dashboards and heat maps.

Summary

Risk Concentration Analysis – the understanding of risks in the context of all other risks - can drive substantial value. Not only does its framework enable a fact-based, repeatable way to score risks at scale, but it provides an efficient way to identify potentially harmful risk concentrations. RCA also provides the ability to test and measure the actual risk-adjusted impact of a specific action/transaction, creating an ROI for action.

Executives and Boards are increasingly seeking a holistic understanding of risk. RCA, based on the materiality of risk concentrations, meets that need. For more information, please contact info@rev2.net.

Copyright © 2011, Rev2. All rights reserved. The RiskView® application is a registered trademark of Rev2. All other trademarks in this document are the properties of their respective owners.