RiskLens™

Cyber Risk Quantification

Risk Management Study (DRAFT)

How You Prioritize, Matters

Overview

This study evaluated the relative efficacy of prioritizing risk remediation activities using qualitative risk ratings, quantitative risk values, or estimates of worst-case impact.

Approach

A simulated environment was created containing 1000 virtual assets. More than 5 million simulations were run where each of the decision modes under analysis were used to drive remediation choices. To provide a baseline for comparison, simulations were also run where remediation choices were made randomly.

Results

- For preventing losses due to inaccurate or imprecise estimates, quantitative prioritization resulted in a 72% improvement over the baseline, while qualitative prioritization showed a 48% improvement over the baseline. Decisions driven by worst-case impact estimates performed 21% worse than the baseline.
- For preventing large losses, decisions based on quantitative analysis performed 21% better than the baseline, while decisions driven by worst-case impact performed 11% better than the baseline. Decisions driven by qualitative analysis performed 3% worse than the baseline.

Overview

The purpose of this study was to gain an understanding of the relative efficacy of different approaches for prioritizing cyber risk remediation activities. This is a critical concern because the complex and dynamic nature of the cyber risk landscape, combined with the inherent resource limitations within any organization, means that being able to focus on what are truly the most important issues is vital to the overall success of a risk management program.

Historically, cyber risk management decisions have been driven by qualitative assessments using ordinal scales (e.g., red, yellow, green, high, medium, low, etc.). This approach has been used in large part because the profession lacked a clear ontology and consistent nomenclature for evaluating risk. There also has been a prevailing (though inaccurate) belief that not enough data existed to do quantitative analysis.

The advantage to qualitative risk ratings is that they can be arrived at almost instantaneously. The potential downsides however, are:

- **Inaccuracy** - the absence of rigor, explicit measurement activities, or formal models increases the probability that risk ratings will not be accurate.
- **Imprecision** - using a three-level (or sometimes five-level) ordinal scale as a means to assign relative significance to cyber risk issues doesn't enable effective prioritization of issues within each level.

With the advent of Factor Analysis of Information Risk (FAIR), problems related to ontology and nomenclature have been overcome. Combined with well-established methods (e.g., calibrated estimation) for effectively leveraging even sparse data, and Monte Carlo and other stochastic methods, organizations are now able to perform true quantitative analyses of cyber risk.

This paper describes at a high level a comparison of the relative efficacy of prioritizing risk remediation activities using qualitative versus quantitative methods. Because some professionals in the industry claim that prioritization should be performed purely on potential worst-case outcomes, excluding any consideration of likelihood, the study included this as yet a third decision mode for analysis.

Analysis Approach

This study leveraged a simulated cyber risk landscape comprised of global variables and one thousand virtual assets. Each run of the analysis was ten "years" in length, with each year being comprised of 365 days. Each day included the following actions:

- A risk assessment to identify assets whose control efficacy was less than a threshold set as a global variable (established to emulate an organization policy or standard)
- Risk remediation (constrained by a global variable for budget)
- Random changes to the landscape to reflect the dynamic nature of cyber risk landscapes
- Potential loss events (the frequency of which was driven by Bernoulli trials on asset variables described below)

In addition, on a monthly basis the remediation budget was refreshed.

Global variables included:

- Decision Mode - reflecting the method (random, qualitative, quantitative, impact only) being used to select which control deficiency to remediate on any given day
- Analysis Quality - reflecting the level of calibration (or accuracy) of risk analyses performed within simulations
- Budget - the amount of resources available for remediation activities
- Large Loss Threshold - a value enabling the identification and categorization of particularly large loss events
- Control Design Efficacy - an intended level of efficacy for asset resistive controls
- Compliance Level - the probability that an asset's resistive strength will be at the intended (designed) level

Each asset was randomly assigned three values:

- Impact - the amount of loss that would materialize if the asset were to be compromised
- A control efficacy value (essentially, the percentage of time an attack would be resisted)
- A threat event frequency

These values were used to assign a Risk value to each individual asset.

Each asset was then assigned a "Perceived Risk" value based on a combination of the Risk value and the Analysis Quality global variable.

Besides simply comparing decision modes, two additional variables — remediation budget size and analysis quality — also were included to evaluate their effect on results.

In order to generate sufficient data and effectively compare the different decision modes, each 10-year analysis (described above) was performed one hundred times for each combination of decision mode, budget level (normal and 50% of normal), and analysis quality level (70% vs 90% accuracy). This resulted in 1600 separate base analyses, equating to 16,000 virtual "years" or roughly 5.8 million remediation decision-days.

The efficacy of each decision mode (including random decisions used as a baseline for comparison) was evaluated on two dimensions:

- Their ability to help an organization avoid "surprises" — i.e., their ability to properly identify and remediate higher risk issues and thus avoid losses from risk issues that had been mis-prioritized and not remediated.
- Their ability to reduce the potential for large loss events

Results

With Budget at normal levels and Analysis Quality at 90%, quantitatively driven choices out-performed random decision-making by 72% in terms of avoiding surprises — i.e., losses that could have been avoided with proper prioritization. In comparison, qualitatively driven choices out-performed random decision-making by 48%, and impact-only choices actually performed 21% worse than the randomly derived baseline.

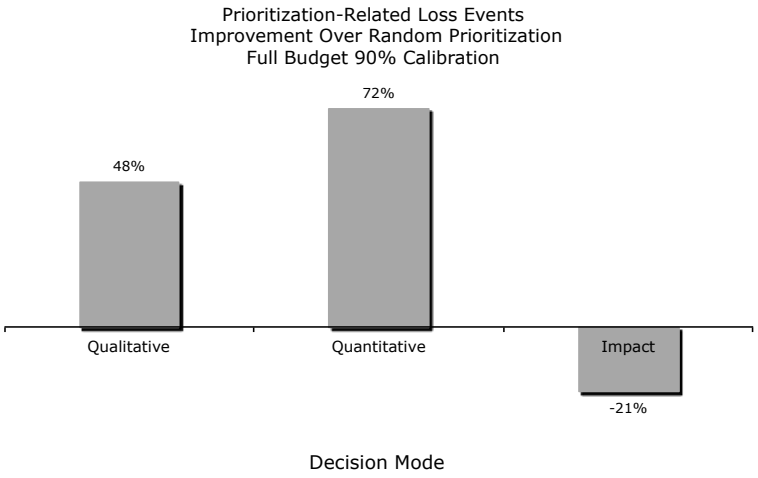

When the remediation budget was cut in half, the efficacy of all three decision modes improved relative to the baseline — qualitative by an additional 6%, quantitative by 3%, and impact-only by 16% (which still left impact-only decisions 5% worse than the baseline).

Reducing the accuracy of analyses within simulations from 90% to 70% resulted in a 22% reduction in the efficacy of qualitatively driven remediation choices, a 42% reduction for quantitatively driven choices, and had no measurable effect on impact-only driven choices.

With Budget at normal levels and Analysis Quality at 90%, quantitatively driven choices out-performed random decision-making by 21% in terms of avoiding large losses. In comparison, impact-only choices out-performed random decision-making by 11%, and qualitatively driven choices performed 3% worse than the randomly derived baseline.

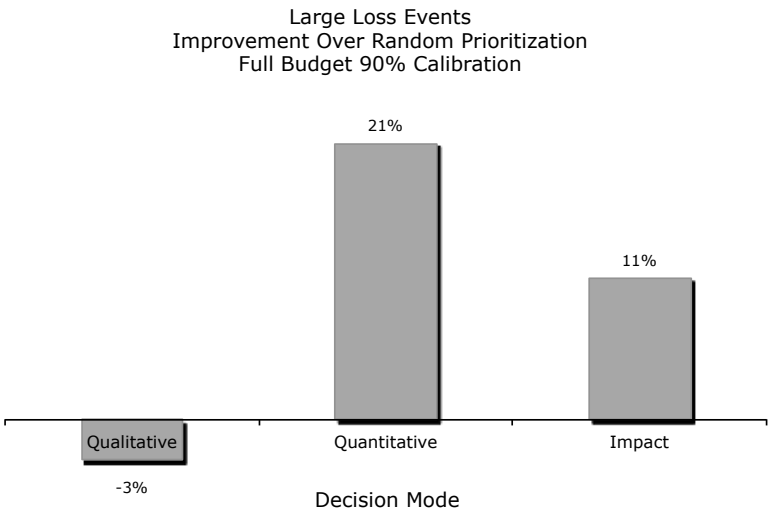

Cutting the remediation budget in half had no measurable effect on the efficacy of qualitatively driven choices in reducing large loss events. It did, however, reduce the efficacy of quantitatively driven choices by 10% and drove impact-only driven decisions to essentially nil over the baseline.

Reducing the accuracy of analyses within simulations from 90% to 70% resulted in roughly a 2% reduction in efficacy for all decision modes in their ability to limit large losses.

SUMMARY

Based on this study it is clear that, all other variables being equal, the ability to prioritize risk remediation activities using quantitative methods represents a significant improvement in an organization's ability to manage risk over both qualitative and impact-only driven prioritization.

The results also suggest that more accurate analysis (via calibrated analysts and improved models and data) is more important for quantitative analyses than it is for qualitative analyses, and is of little value for impact-only driven decisions.

For highly resource constrained organizations, the benefits of quantitative prioritization over the other decision modes, although reduced, was still substantial.

Contact RiskLens for more information about this study or to learn about how your organization can begin to realize the benefits cyber risk quantification provides.