MILLIMAN WHITE PAPER

Know your cyber blind spots

The importance of modeling cyber risk for businesses

Chris Harner, FRM
Chris Beck
Blake Fleisher, Associate of (ISC)2

Milliman

Modeling nonfinancial risks such as cyber has presented immense challenges for employers, investors, and insurers. Historically, risk managers have looked to tools based on frequency-severity (f/s), which not only provide a limited ability to answer key risk questions, but also create a false sense of security that an organization understands the risk. Often, modelers are tempted to “get to the number"; of equal importance is management buy-in of the model output. That buy-in is dependent on the credibility, transparency, and explanatory power of the model.

Concerns regarding frequency-severity modeling are amplified in the case of cyber risk, which embodies three unique attributes:

1. Cyber is an *adversarial* risk: Your opponent—whether a hacker, insider, or state actor—is trying to outthink you.
2. Cyber is a *high-velocity* risk that provides little to no warning.
3. Cyber is a *stealth* risk: You may be compromised for an extended period of time and not even realize it.

For these reasons, it is essential for any business to rethink how to best model its cyber risk with the goal of illuminating blind spots instead of missing them.

Business analytics can provide valuable clarity to a company’s business strategy and risk management decisions. Cyber models serve as a set of analytical tools that can provide such insight. Some companies neglect to model their cyber risks, whether due to other pressing business priorities or confusion on how to approach cyber risk modeling. This leaves these companies in the dark, causing them to speculate about the ramifications of a cyberattack on their business.

Some companies give high priority to the modeling of cyber risk, but are oftentimes misled into thinking that frameworks with checkboxes, heat maps, risk and control self-assessments (RCSAs), or even f/s models (i.e., Cyber VaR) provide sufficient insight. While these tools may satisfy certain questions under certain circumstances, their accuracy is highly uncertain.

It is critical that the right tools are chosen for the job. While a good model supports business strategy, a poor model can undermine it. Existing approaches leave risk mitigation insights to be desired.

Frequency-severity models

Given the popularity of value-at-risk (VaR) since the late 1990s on Wall Street to measure market and credit risk, some modelers have turned to the f/s approach by applying Monte Carlo simulations. However, such models suffer from limited or incomplete observations and unknown correlations. They fail to capture the impacts of volatility found in cyber and other operational risks. Frequency-severity models are grounded in the assumptions that the past will repeat itself and that there is evidence of future events in historical data. We will elaborate on the shortcomings of using this method to quantify cyber in this paper.

Faulty underlying assumptions

The frequency-severity approach requires modelers to assume away some of the complexity inherent in real-world conditions. Modelers fit distributions for frequency and magnitude in a siloed manner and combine these distributions using statistical methods that assume independence between the two variables. How often a risk event occurs is considered entirely separate from the magnitude of impact. This may hold true for car insurance (the likelihood of an accident does not impact the damage caused by the accident), but cyber risks are constantly evolving with new types of attacks every day.

The idea that frequency can be divorced from severity does not hold in an adversarial risk environment such as cyber. An adversary thinks strategically and develops tactics to defeat cyber controls and gain access to systems, data, and applications. When looking at these events in modeling terms, there are very good reasons that an adversary may align the frequency and severity of attacks.

A cyber risk model must account for scenarios that businesses have never faced. Lack of evidence in historical data of a novel cyberattack does not mean that such an attack is implausible or impossible. In some cases, it might mean the opposite; namely, that the novel attack a bad actor thinks will be successful has the greatest probability of success because controls have not been created to prevent such an attack.

To make up for the lack of data, cyber risk models must turn to other sources. Frequency-severity models rely heavily on curve-fitting techniques of historical data. There are two directions modelers can take: mine industry data pools or incorporate estimates from business experts.

On their own, industry data sets provide valuable insight to a business. However, if used in a model there is no guarantee that the aggregate data applies to the specific business in question. Businesses have differing policies, technologies, and levels of cyber hygiene. When a model overemphasizes such aggregate data sets, caution must be used. Too often, the end result of a frequency-severity model obscures which data inputs have the largest impact. It becomes challenging for decision makers to delineate which insights are primarily influenced by internal business data, versus which came from aggregate data sets.

To construct a frequency-severity model when no loss data is available, modelers must transform expert opinion into loss curves. Thus, modelers turn to program evaluation and review technique (PERT) distributions. PERT distributions are used when little data is known and expert analysis is limited. This technique generates a curve based off of only three data points: an estimated minimum, mode, and maximum. This approach attempts to turn expert opinion into curves that mimic the shape generated by historical loss data.

PERT distributions are useful in obtaining basic insight into a distribution when there is a limited amount of data. However, frequency-severity models treat these roughly estimated distributions as statistical fact. Modelers run simulations on this data as if it is historical record without regard for the lack of precision. Current industry applications rely upon PERT distributions to form the basis of a cyber model.[1]

Decision makers who use the model cannot determine to what extent these estimates contribute to the end result. Furthermore, with this method, experts are asked to guess absolute probabilities about the likelihood of a loss event. They must estimate the minimum, mode, and maximum with limited context. A PERT distribution amplifies any error in these guesses and may break down altogether as it is not suitable for all situations.[2]

Example

STATE-BACKED ADVERSARIAL ATTACK SCENARIO

The resources of cyber groups and attacks supported by state actors give them some of the most sophisticated cyber threat capabilities. These cyber threats tend to be highly disciplined, well-funded, and secure due to their military or intelligence origins. Their motivation is often driven by the goals of the nation-state. As such, the frequency and/or severity of their actions are byproducts of sophisticated strategies and inextricably linked. For instance, a useful strategy for a state actor could be to create a series of high-frequency attacks, albeit of minor severity, followed by a limited highly sophisticated attack. In this case the state actor would both gather information and create a false sense of security while planning a more significant effort. This makes the frequency and severity of attack variables purposely linked. They cannot be modeled as independent. Furthermore, a company using frequency-severity modeling would be looking to historical loss data to inform its model. This data would only include methods that have been previously employed by bad actors. A new method for a cyberattack, even if it is one that experts widely agree is a concern, would not be included in the modeling.

Takeaway: In this scenario the likelihood of the event decreases as we consider increasing severities. An effective model would explicitly show that the identity of the adversary may inversely affect frequency and severity. Causal modeling can accomplish this by taking into account different types of threat profiles and capabilities, and incorporating them into the context of interconnected risks. Furthermore, a causal model can account for the attack method that is plausible, but has yet to be used.

[1] Freund, J. & Jones, J. (2015). Measuring and Managing Information Risk: A FAIR Approach. Oxford, UK: Butterworth-Heinemann. p. 99-103

[2] Antkiewicz, M. (December 13, 2011). Some thoughts about PERT and other distributions - part 1. Society of Information Risk Analysts. Retrieved March 19, 2021, from https://www.societyinforisk.org/Blog-Posts/4355797. The author notes: “PERT doesn’t return useful results when the minimum or maximum are very large multiples of the most likely value.”

Causal modeling

Unlike frequency-severity models, causal models combine principles from complexity theory, network science, and Bayesian statistics. Causal models incorporate insight from experts without forcing the data to represent a loss curve directly. For instance, causal models use Bayesian networks to express, according to experts, how one event interacts with another to produce outcomes. Using conditional probabilities, the Bayesian network is able to build up a view of how likely different pathways are. When modelers construct such a model, they are able to ask experts simpler questions about how the factors interact, rather than asking them to directly estimate the ultimate likelihood. Questions are asked to organically uncover the paths a risk will take as it manifests. By situating unprecedented events in context, experts provide better estimates (even of cataclysmal events, such as a ransomware attack or a breach of proprietary secrets). Building up a logical explanation of how outcomes can be derived enables experts to reflect chains of events that have never been seen in past data, but are plausible, and for which reasonable parameters can be estimated.

Causal models add situational contexts and ask experts to address the kinds of information that they easily have on hand. Experts within a company likely understand the relationships between drivers of one loss event and another. In the actual model, each input is "weighted" to determine the overall likelihood of a risk manifesting. This allows causal models to more intuitively and accurately represent a company's cyber risk landscape, as well as to easily identify paths that lead to cascading failures.

This sort of model is more appropriate for a complex risk like cyber because the model is built from the ground up to handle expert insight with context. Causal models allow the outcome to emerge from real-world complexity, by piecing together expert estimates and data through probabilistic modeling.

Example

PHISHING DEFENSE SCENARIO

When looking at a company's ability to prevent a phishing attack, how well the staff is trained, how often the staff turns over, and the ability of the spam filter to block phishing emails all are part of the overall prevention picture. Using standard other methods, it is difficult to combine these three components. By using causal modeling we can not only measure each component with data that naturally describes the node (e.g., percentage of staff that passed phishing training), but also combine the conditions for an outcome at the *phishing readiness node* (e.g., a high level of success with staff training, combined with low turnover, is more optimal than the same with high turnover). Additionally, a robust *spam filter* may be as valuable as a well-trained staff. The fewer potentially harmful emails that users receive, the lower the risk of a successful phishing attack. Moreover, because we can observe the staff nodes (*Training* and *Turnover*), the *Spam Filter* node, and the overall *Phishing Readiness* node, we can continue to refine the probabilities that define the nodes to make the model more accurate.

FIGURE 1: CAUSAL MODEL

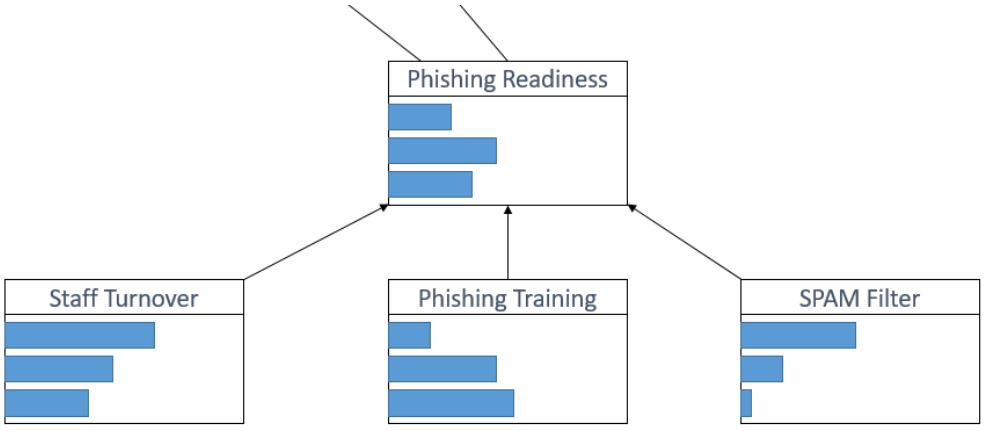

Conclusion

Operational risk is inherent to business. The challenges with cyber are compounded. Not only is there the potential for a control to fail or for human negligence to cause an issue, but there is also a bad actor looking to exploit these vulnerabilities. Cyber risk can be greatly reduced by modeling a business's cyber landscape, revealing the most critical paths that amplify risk. This allows a business to understand the interconnectedness of various people, processes, controls, and risks that lead to loss events. With this capability, management can measure which mitigating actions will reduce firms' risk profiles most effectively.

A good cyber risk model should go beyond the fundamentals of satisfying regulators or providing management with a number on how much capital to hold. Rather, the model should be complex enough to reflect the risk but simple enough to understand, providing management with real insights for risk mitigation strategy.

The key to modeling cyber is to have understanding of the business problems it can solve. Understanding cyber risk is not just about getting a number that describes the exposure. It is about weighting the value of risk mitigation decisions and controls implementation against risk reduction and business needs. There is no set of controls that can eliminate cyber risk. Cyber risk cannot be mitigated away. Firms need to model their overall cyber risk exposures in ways that allow them to see the impact of their risk decisions so that they can make the best business decisions.

Technology is complex and interconnected. It makes sense that the risk involved is likewise. When money is at stake, decision makers need a model that can be trusted to handle the complexities of risk. While no model is perfect, causal models create a far closer depiction of reality in cyber risk than current approaches.

Milliman

CONTACT

Chris Beck
chris.beck@milliman.com

Blake Fleisher
blake.fleisher@milliman.com