Why the FAIR Model Can Be So Unfair

CISOs Need Better Risk Assessment Tools

Dr. Aleksandr Yampolskiy, CEO & Co-Founder, SecurityScorecard

Jim Routh, former CISO of Aetna, MassMutual, American Express

SecurityScorecard

Cybersecurity is an ever-evolving challenge, and accurately quantifying cyber risk is a complex task. Security leaders are in constant need of specific, clear metrics to effectively measure the effectiveness of controls and manage risks. Making trade-off decisions about resource allocation can be quite difficult, even if it appears straightforward to outsiders.

The current technology landscape includes new exploitable vulnerabilities, patches, code releases, version updates, and countless other data points. Companies need automated cybersecurity with continuous monitoring if an organization wants to accurately assess its risk exposure.

> **You can't manage what you don't measure.**

"You can't manage what you don't measure." So, communicating and measuring risk in monetary terms could be an attractive concept in theory.

However, security practitioners still struggle with prioritizing the issues created by the bloat of various security tools they employ. How do they know which issues will cost the most money? And where should they focus their already limited attention?

The last decade has seen a sharp spike in cyber incidents and an increase in cyber risk governance and regulations — which require much more cybersecurity oversight at the executive and board levels. Recent enforcement actions by the SEC and other regulatory bodies have demonstrated the importance of a unified view of risk, both from a cost-savings perspective and a liability one. More to the point: giving more stakeholders a real-time, monetary value of risk will motivate them to mitigate that risk.

The cybersecurity community loves a good debate, and truthfully, there isn't one perfect method for quantifying risk. The best we can do is find one that's straightforward, scalable, easy to understand, and **trusted**.

What is the FAIR model?

The Factor Analysis of Information Risk (FAIR) has emerged as a promising framework for quantifying cybersecurity risk. It is a quantitative risk analysis model that aids organizations in assessing cyber risks unique to their environment and translating the impact of these risks into mathematical risk estimates.

FAIR analyzes scoped risk scenarios and translates them to quantify potential loss exposure. This helps organizations better understand their risk posture, where they're most likely to be impacted by a cyberattack, and potential financial loss. FAIR is an international standard for quantifying cyber risk governed by The Open Group. The non-profit FAIR Institute, while ostensibly promoting the FAIR standard, primarily seems to promote RiskLens' (now Safe Security's) FAIR calculator software.

As the saying goes, every model is flawed, but some are still useful. If you are going to trust your organization's risk assessments to any model, it's important to know its strengths and weaknesses.

Below are some of what we see as the high and low points of utilizing the FAIR model for your cyber risk program.

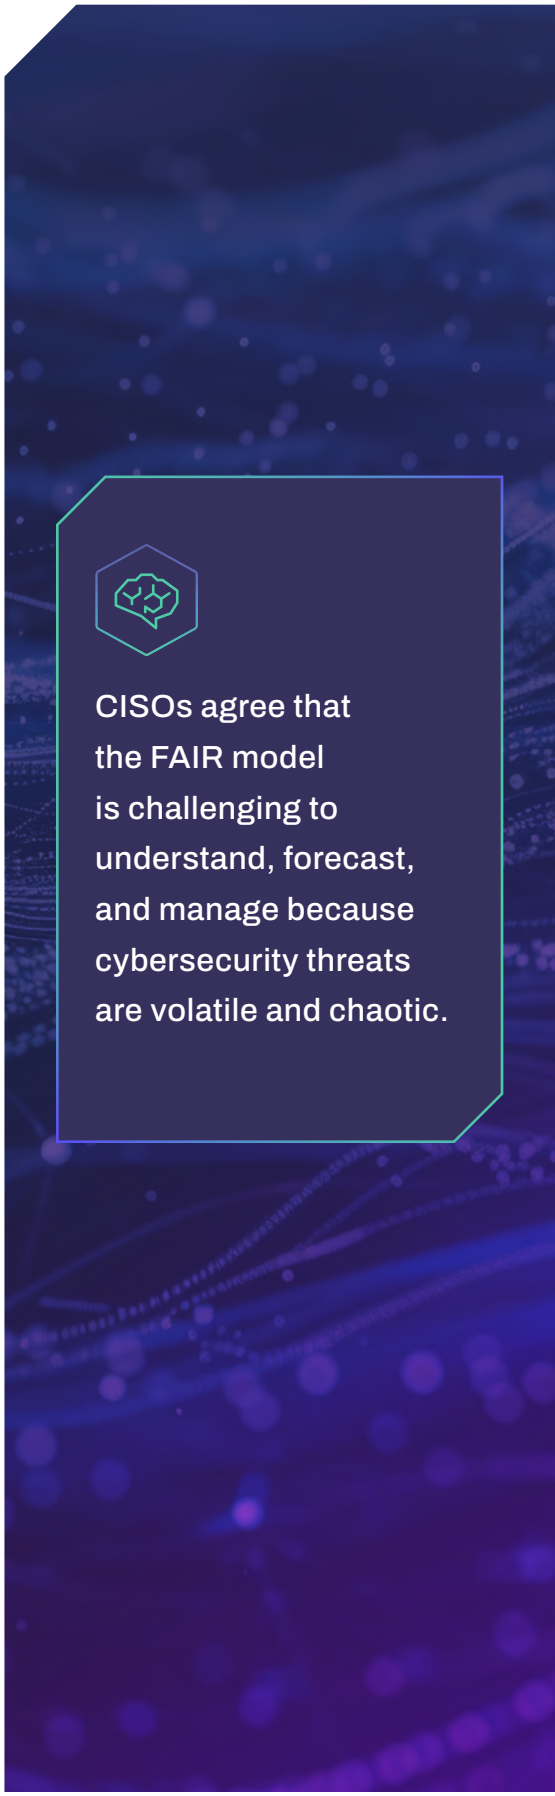

Benefits of the FAIR model

Quantitative risk assessment: The FAIR model provides a quantitative approach to risk assessment, enabling organizations to make informed decisions based on numerical data rather than subjective judgments. This can improve the objectivity and consistency of risk evaluations.

Regulatory compliance: The FAIR model may help organizations gather and present quantifiable risk metrics that align with regulatory requirements, facilitating compliance with standards such as GDPR, HIPAA, etc.

Historical data validation: The FAIR model is able to work with historical datasets – if you have them. This allows organizations to backtest their risk assessments, ensuring their models are accurate and reliable. If you don't, it also provides a process to help you create such datasets about your organization. However, without access to historical data breach datasets, we believe organizations risk having inaccurate data – resulting in vastly wrong dollar outputs.

Enhanced decision-making: The FAIR model supports more strategic decision-making by providing a quantifiable understanding of risk. Organizations can prioritize resources and investments based on various risks' potential impact and likelihood.

Limitations of FAIR

Time to Value

While it's good in principle, we do not think FAIR has crossed the chasm to become a usable and practical model (besides a few eager innovators and adopters — who may believe in its promise but not the technology itself).

However, it's important to note that the FAIR model, while promising, has its limitations. A frequent complaint about FAIR is its lengthy implementation process, which can often require a dedicated team of experts. Many organizations eventually conclude that the benefits of FAIR don't outweigh the time, effort, and investment required. Gartner predicts that by 2025, 50% of cyber-risk quantification projects will fail.

Risk = Loss x Probability Threat x Vendor Breach Likelihood

Data Quality

FAIR is not immune to the GIGO (garbage in, garbage out) problem. Unfortunately, using erroneous dollar or probability numbers can create more harm than good. It's not the model itself that's bad — but how people use it; in other words, the complexity of implementing FAIR results in security practitioners taking shortcuts, which results in less-than-desirable results.

Timeliness

The accuracy of the FAIR model is not only heavily dependent on the quality of the data, but also its timeliness. If the data is outdated or inaccurate, the risk assessments generated by the model will be less reliable. FAIR is only as good as the data it is supplied with and demands a large amount of up-to-date and accurate data from a broad spectrum of sources to perform optimally.

With FAIR and other frameworks, relying on point-in-time static data will not help to accurately determine and measure the materiality of threats across the enterprise as they evolve in real time. In heavily regulated environments, proof of material attack vectors with evidence-based data is required, and continuous proactive data is a must to stay fully compliant and reduce the likelihood of legal scrutiny in the event of a material incident.

Many organizations eventually realize that FAIR doesn't produce enough material benefit to merit the time, effort, and investment required.

Gartner predicts that by 2025, 50% of cyber-risk quantification projects will fail.

Data Veracity

Recent cybersecurity mandates require proof of underlying enterprise risks and vulnerability rankings. New regulatory cybersecurity requirements elevate the need for irrefutable data to maintain compliance and avoid fines.

The more complete the data, the better the results; conversely, when given less-than-ideal levels of data coverage and quality, FAIR can produce highly variable and potentially misleading results – leading to poor and inconsistent risk management decisions.

When challenged to substantiate with material data, the FAIR thesis can fall flat. Even when cloud environment data and other data are connected to the model – the lack of breadth, depth, and substantiation can leave clients with nothing more than a vanity metric to check a box while quietly yearning for more reliable solutions.

Scalability issues

While the FAIR model can be effective for an organization's internal use, scaling it across third-party or fourth-party vendors can introduce significant challenges. If users need to click 20 levels down to understand certain controls, they will never really use the model in practice, as it will create more work. While this might work for an organization's internal use, it's unrealistic and can be nearly impossible to scale for third- and fourth-party risk management.

Often, a CMO or CFO signs a contract and needs to make a snap decision for that vendor, "Do I accept risk, mitigate risk, or transfer risk?"

Given that FAIR is best used when a problem statement is enumerated in a number of scenarios, it makes it difficult for a CMO or CFO to understand the real risk if they have to look at 20 different scenarios (ransomware, DDOS, outages, etc.). The consistency and control required for accurate risk assessment are difficult to maintain when dealing with external entities.

Organizations need to have a rapid litmus test — similar to a credit score or a letter grade (A – F) to decide if they want to ask more questions or move on. There's no substitute for real-time signals, and many FAIR implementations simply don't offer that.

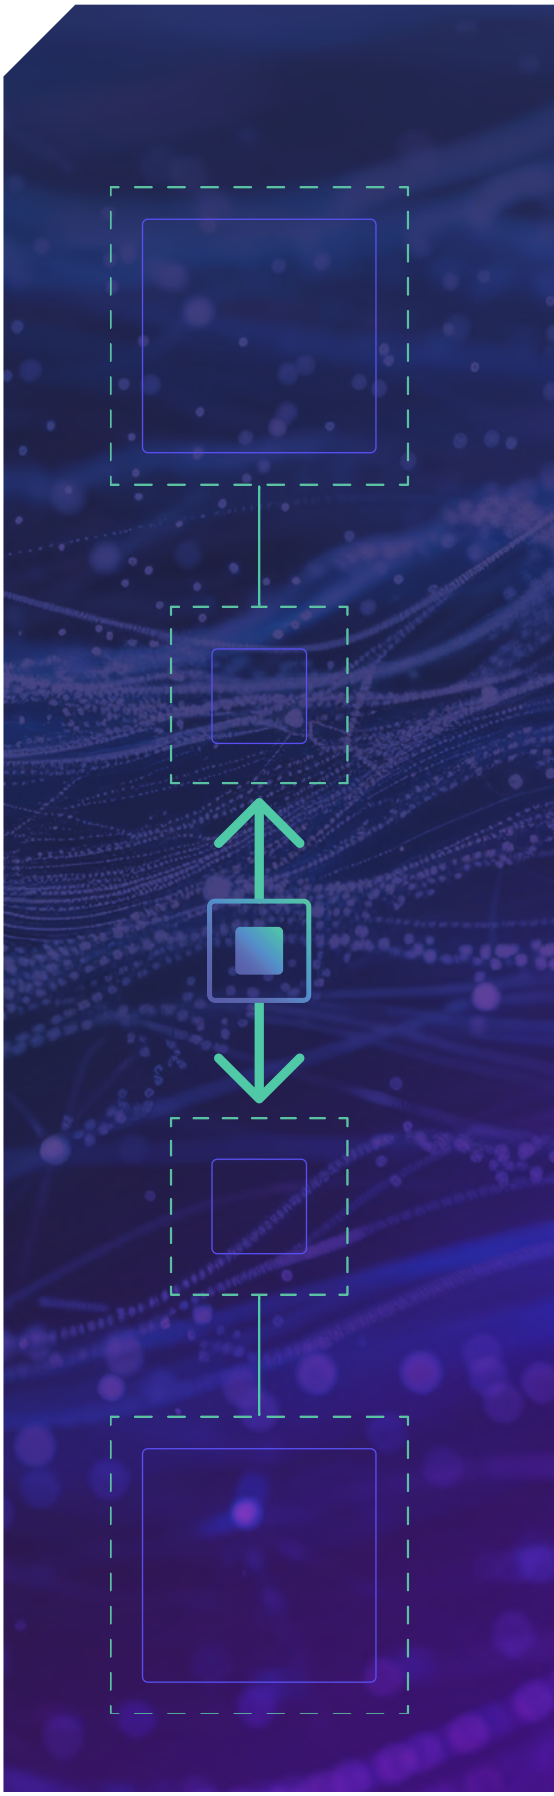

Limited inputs

The predominant inputs into FAIR calculations are subject matter expert-calibrated range-based estimates. Often, these are augmented by surface scans, complex technical integrations, and lengthy questionnaires, which may not provide a comprehensive view of the cybersecurity landscape. For instance, relying on an SOC report from six months ago might not capture recent changes in the threat environment. The weights of these inputs relative to the nature of the entity may vary drastically, again demonstrating that one size does not necessarily fit all when it comes to cyber risk modeling.

Missing Controls and Vulnerabilities

FAIR requires a control taxonomy. While this isn't an inherent weakness, we believe the solution (FAIR-CAM) is as overly complicated as it is universally unwanted. Most organizations are burdened with the extra effort of mapping back to NIST ISO and other more established models.

There are no specific variables for loss magnitude controls. Instead, organizations run simulations twice — once without the controls they are modeling and once with them.

FAIR also does not natively prioritize vulnerabilities. This is an important step for enterprises that need actionable steps to improve their own cybersecurity posture and that of their third- and fourth-party vendors.

Although FAIR allows for validated data, many practitioners leverage SME-estimated data ranges to derive the upper-level loss variables (LEF and LM). This complexity can lead to uncertainty, mainly when the model is applied to third-party risk assessments, where data quality and availability vary widely. This faulty approach can create a false sense of safety for users while potentially creating more dangerous conditions and outcomes and widening the gap within the risk mitigation program.

Significant onboarding leading to turnover

While the SME modeling approaches mentioned above are statistically valid, they are difficult for many to understand and require extensive training, both on the data collection processes and on the interpretation of results. You may also have to work to convince executives that the results are valid.

As a result, operationalizing and scaling FAIR across the organization can be highly challenging. FAIR should also be used sparingly because it can require many experts' time — time they could spend on other high-value projects.

Difficult to automate

Valid data and pre-defined scenarios are needed to automate FAIR. Although FAIR allows you to analyze any scenario and any level of granularity, the training leads users to think that the only way is to use very specific and unique scenarios.

Cybersecurity leaders know that many of the cybersecurity tools they purchase are left unused due to a lack of internal resources. Therefore, keeping massive amounts of data up-to-date manually is not a viable option due to resource constraints, such as talent gaps, burnout, cybersecurity budgets, and more.

Difficult to operationalize

Admittedly, some of these drawbacks are less about the FAIR model itself and more about how practitioners apply it. Overall, FAIR is more or less a bespoke version of the universally used loss distribution approach (LDA). In other words, the model may be valid, but the failures are in how many people and teams choose to operationalize it.

Next steps: What does the future hold?

We believe the FAIR model was a decent starting point but it also demonstrated how limitations and assumptions in computational models can be mis-sold and misused in corporate governance.

Next-generation, dependable risk computation will require continuous quantification based on real-time intelligence — by monitoring settings changes as well as the entire supply chain. It's no longer about just your attack surface for an organization - but about an attack surface of **yourself + your third-party suppliers**, backed by many years of historical outside in + inside out data.

Because FAIR generally requires manual analysis and expert knowledge, organizations that rely on it will never truly have an up-to-date understanding of their attack surfaces. Considerations for the nature of how business is done within a vertical or between companies should also have to be factored into the weights of the inputs.

Sometimes, the medicine is worse than the illness itself, and simplicity always wins over complexity. Navigating the landscape of cyber risk management is a complex journey, but understanding the benefits and drawbacks of the FAIR model is a necessary first step. There is no perfect risk model, and there are a variety of CRQ (cyber risk quantification) models in the market, so every organization can choose one that best fits its needs.

SecurityScorecard.com
info@securityscorecard.com