UnFAIR: Simplifying and Expanding Technology Risk Quantification

John Benninghoff
Security Differently

This talk is a proposal for a new tool – and includes a working demo. While the tool is new, the concepts are based on research and observations from industry, my peer David Grimmer's work starting risk quantification at our last company, and my own analysis and research. It's also the story of my ongoing journey to better model and estimate risk. Please do ask questions at any time. I'll have a QR code at the end with links to slides from the talk including all references and speaker's notes.

SANS

usenix

SRE CON

It’s great to be back at SIRAcon! I know many you, but for those who don’t know me, here's a little of my story. Me on upper left, wife Jolene and our dog Gertie. Started in security after attending SANS Network Security 1998. 20 years later, MSc in safety science (managing risk and systems change, 2018-2021). More recently, I worked in Site Reliability Engineering, starting in 2020.

SANS: https://www.sans.org
TCD: https://psychology.tcd.ie/postgraduate/msc-riskandchange/, image: https://commons.wikimedia.org/wiki/File:Trinity_college_library.jpg
SREcon: https://www.usenix.org/srecon

Why CRQ?

What is the value of (Cyber) Risk Quantification?

What is the value of (Cyber) Risk Quantification? I see risk quantification as a tool to help inform and improve organizational decisions, primarily investments.

Blunt End and Sharp End

In safety, we talk about the blunt end and the sharp end: the sharp end of the organization is the people who do the work, and the blunt end is leadership. There's a gap in understanding; the executives at the blunt end don't and can't have as complete an understanding as those at the sharp end. (Also true for different practitioners).

Image: Figure 3 from: Cook, R., Woods, D., & Miller, C. A. (1998). A Tale of Two Stories: Contrasting Views of Patient Safety. https://www.researchgate.net/publication/245102691_A_Tale_of_Two_Stories_Contrasting_Views_of_Patient_Safety

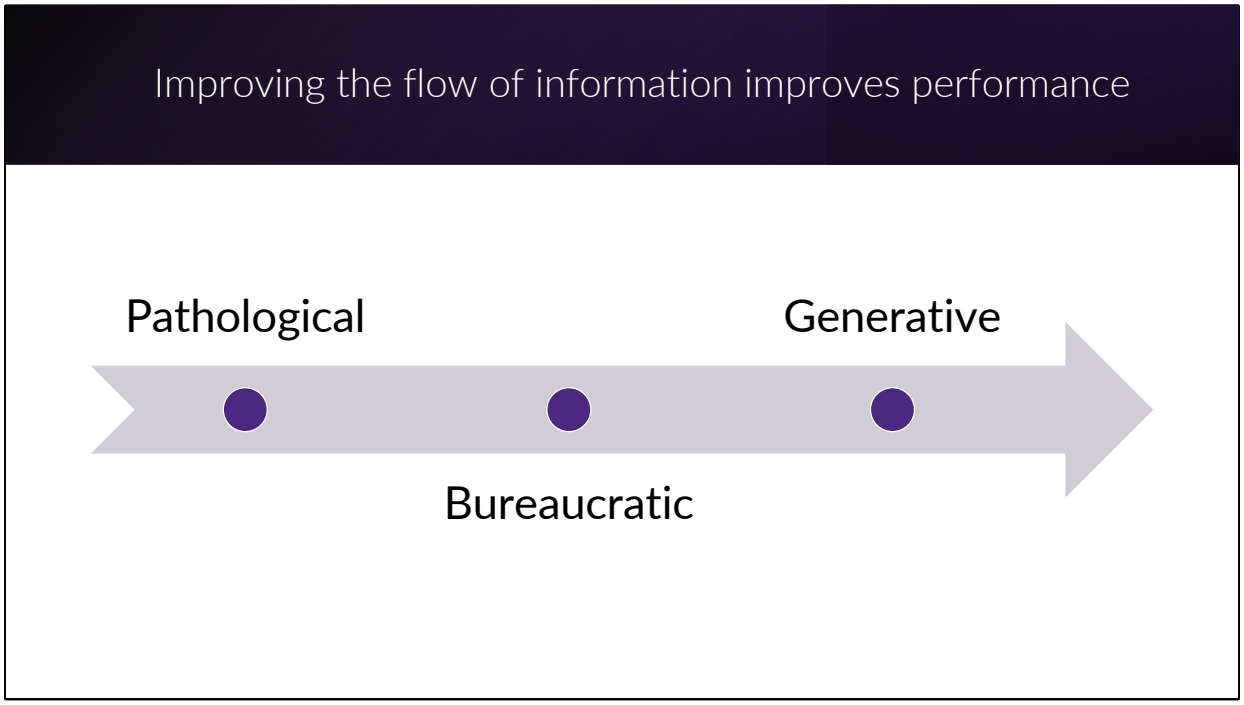

I argue that the goal for Risk Quant is to improve the flow of information and knowledge from the sharp end to the blunt end. The work of Ron Westrum and Google DORA has shown that cultures that promote the flow of information improves performance.

Westrum, R. (2014). The study of information flow: A personal journey [Article]. Safety Science, 67, 58-63. https://doi.org/10.1016/j.ssci.2014.01.009
Westrum, R. (2004). A typology of organisational cultures. Quality and Safety in Health Care, 13(suppl_2), ii22-ii27. https://doi.org/10.1136/qshc.2003.009522
Google. (2024). DORA | Capabilities: Generative organizational culture. Retrieved 2024-08-13 from https://dora.dev/capabilities/generative-organizational-culture/

Role of Risk Analyst

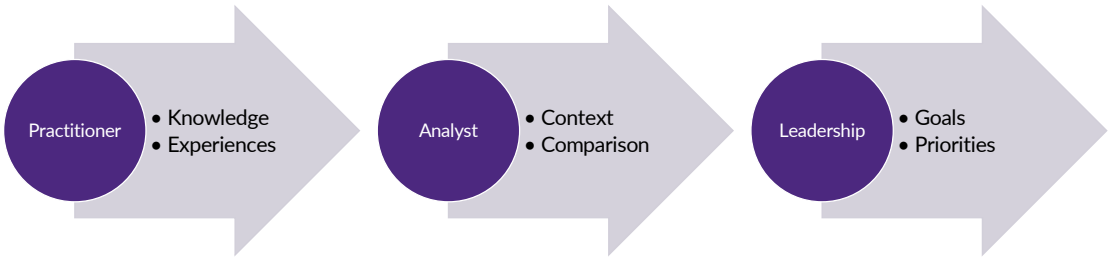

What is the role of the risk analyst? To facilitate the flow of knowledge from practitioners to leadership, to tell their story with context and allow comparison with other priorities and goals.

Limitations of FAIR

FAIR, the Factor Analysis of Information Risk, is an excellent tool, and has become the de facto standard for quantifying cybersecurity risk, but has some limitations.

https://en.wikipedia.org/wiki/Factor_analysis_of_information_risk

More Factors = More Better?

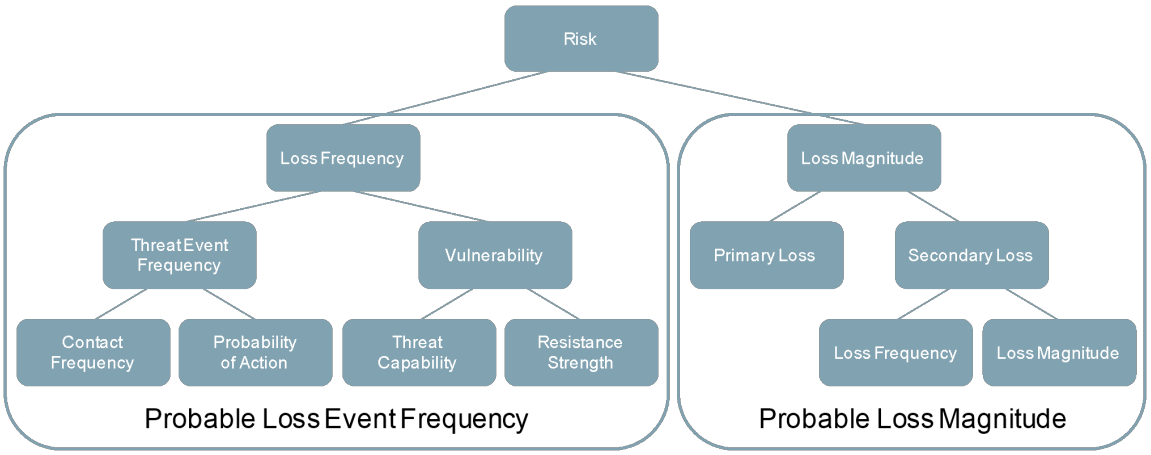

More factors isn't necessarily better; there is evidence that *fewer* factors give better estimates, and in my opinion, more factors serves the Analyst, but not the experts; it's easier for the experts to simply estimate frequency and magnitude directly, and use fast-thinking for the many factors that contribute to each.

* I lost track of the reference on the benefit of fewer factors; it was from a talk Miles Edmunson gave at Secure360, where he spoke about using Monte Carlo for risk estimation (without knowledge of FAIR)
Image: https://pubs.opengroup.org/security/openfair-process-guide/#_Toc503856057

Which Distribution?

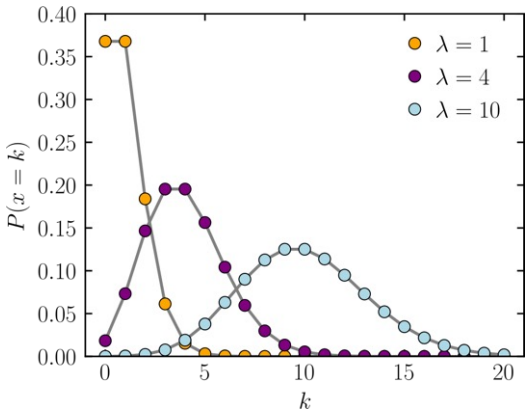

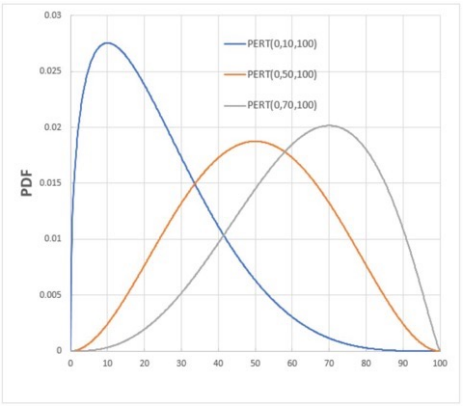

While OpenFAIR doesn't prescribe a specific distribution, historically, FAIR uses BetaPERT. There are 2 issues with this: first, frequency is better modeled with a discrete distribution, like Poisson.

The Open Group. (2021). *Risk Analysis (O-RA), Version 2.0.1*. https://publications.opengroup.org/c20a
Images: https://commons.wikimedia.org/wiki/File:PERT_pdf_examples.jpg, https://commons.wikimedia.org/wiki/File:Log-normal-pdfs.png

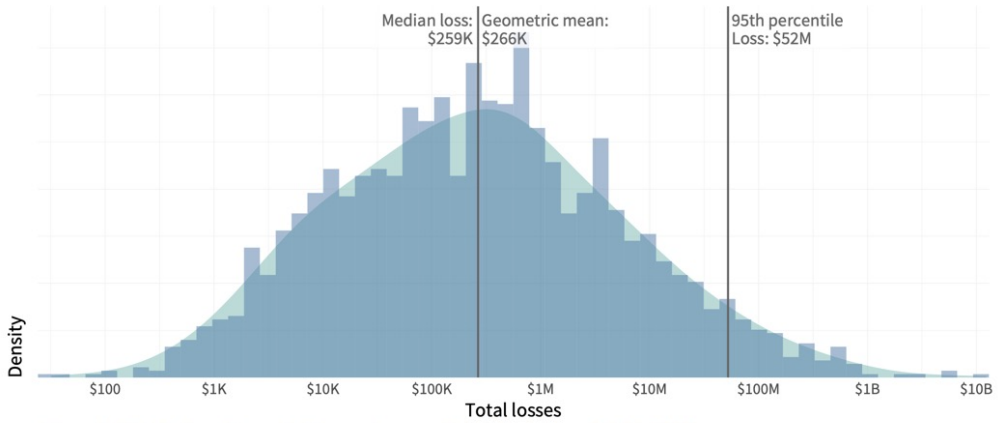

Figure 6: Distribution of reported losses for security incidents from 2012 to 2021

The Cyentia IRIS reports show that loss distribution is log-normal.

Cyentia Institute. (2022). *Information Risk Insights Study (IRIS) 2022*. https://www.cyentia.com/wp-content/uploads/IRIS-2022_Cyentia.pdf

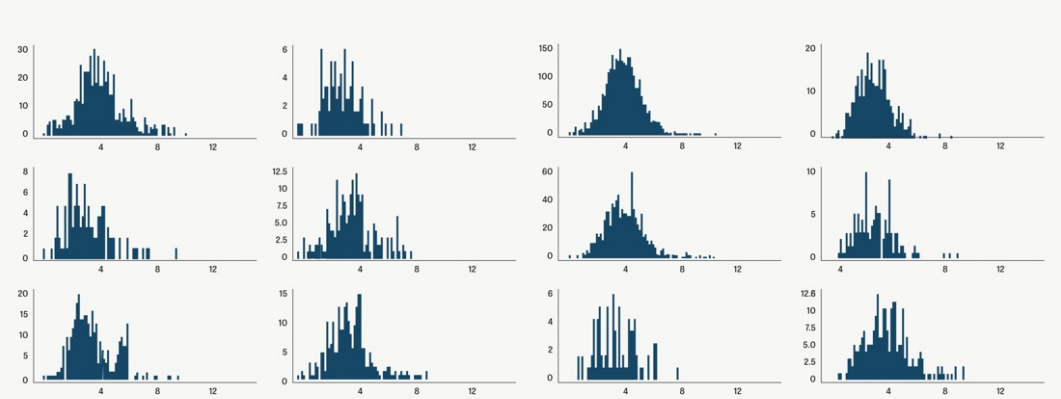

Figure 5. Distribution of log-transformed duration data from 12 companies in the VOID. Each chart represents a different company.

Technology outage duration times are also generally log-normal. In my own work, I found that outages of a specific type did fit log-normal quite well, some of these may be mixing multiple types and are multi-modal as a result.

Nash, C. (2022). *The VOID Report 2022*. https://www.thevoid.community/report

What are we missing?

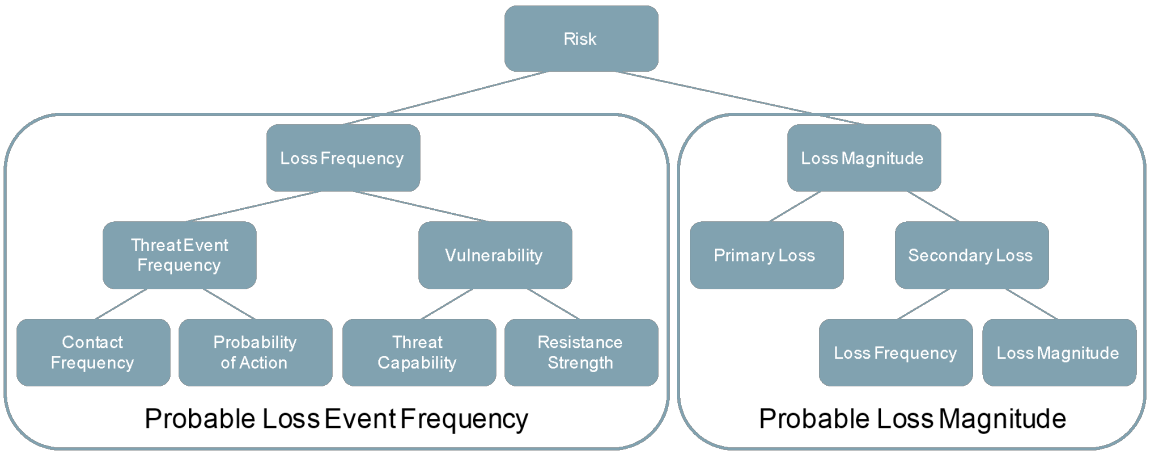

Most importantly, by focusing only on cybersecurity risk, FAIR excludes non-security risks that may be larger. By asking "What are we missing?" the analyst is mining for knowledge of hidden risks.

Why do we need *another* RQ tool?

Commercial
- Safe Security
- RiskRecon
- Ostrich
- Riskonnect
- Derive
- Onspring
- Axio
- SecurityScorecard
- Black Kite
- ProcessUnity
- Kovrr
- Erambra
- ThreatConnect
- CyberSaint
- Alfahive
- Monaco Risk
- Vivo Security

Limited Use
- FAIR-U
- Open FAIR™

Free!
- tidyrisk

So many tools… but most are commercial products or limited-use. tidyrisk is great! But I wanted to create something simpler, and not based on FAIR.

PaulS, Aaron Arutunian, SIRA Slack, #tools-apps-software

Design Goals for “UnFAIR”

"Risk Quant on Rails" - an opinionated approach that is easy to use and for the analyst to understand, focusing on the flow of knowledge

Image: https://commons.wikimedia.org/wiki/File:Ruby_On_Rails_Logo.svg

Design Goals for “UnFAIR”

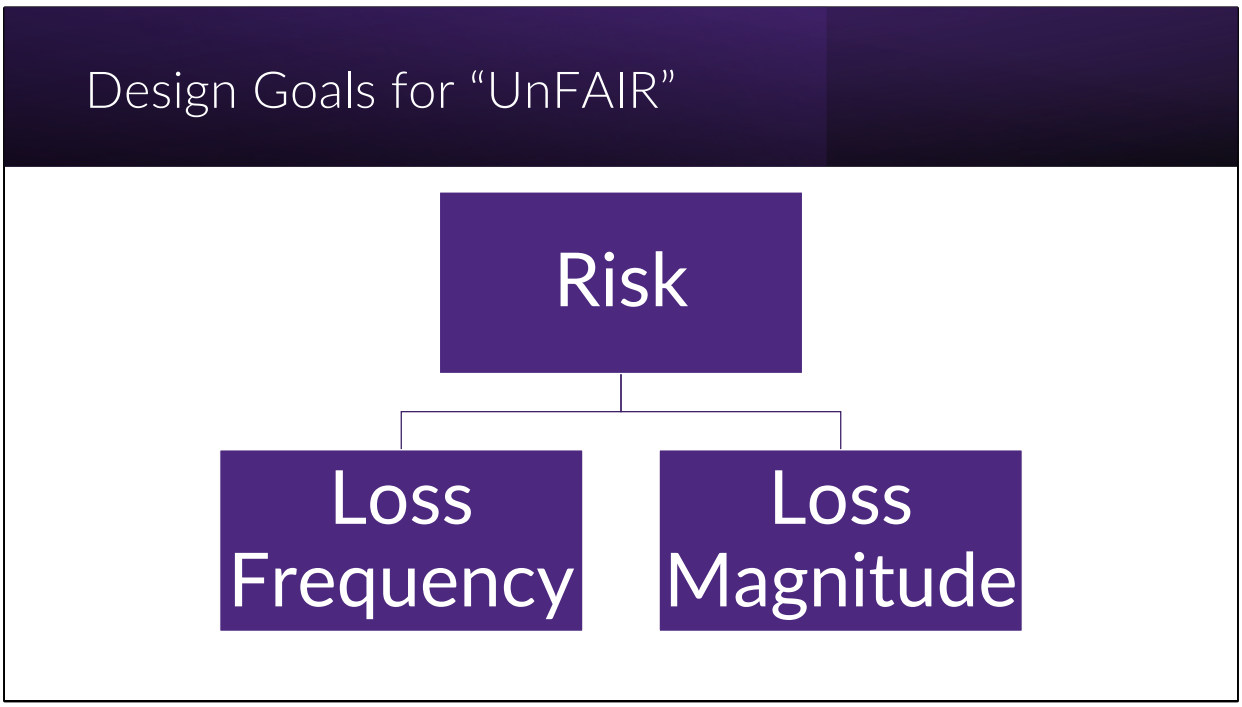

An alternative approach that follows the core of FAIR and Hubbard's methodology, estimating only two factors: loss frequency and loss magnitude.

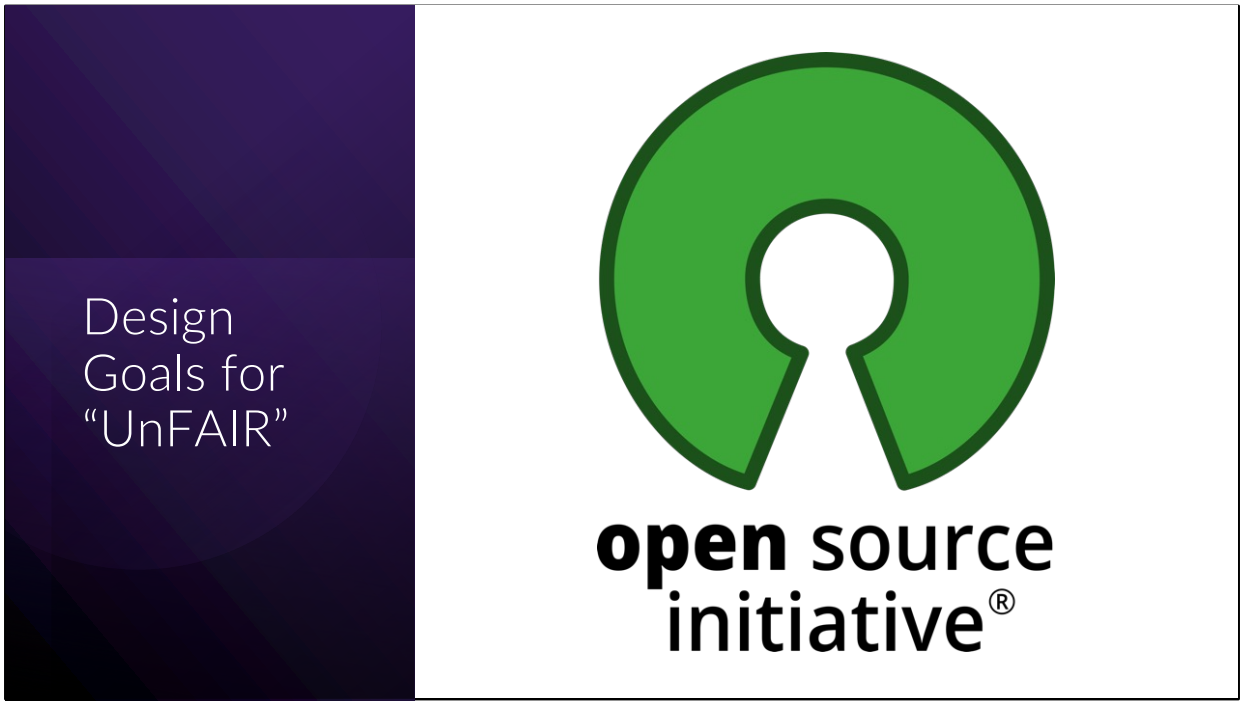

Use only free, open source, and readily available tools (it’s hard to escape use of Excel or its equivalent).

Image: https://commons.wikimedia.org/wiki/File:Open_Source_Initiative.svg

Design Goals for “UnFAIR”

Limited Time

No Budget

Ultimately, the demo is for people with limited time and essentially no budget; the one person trying to start Risk Quant who understands enough math to explain a basic model.

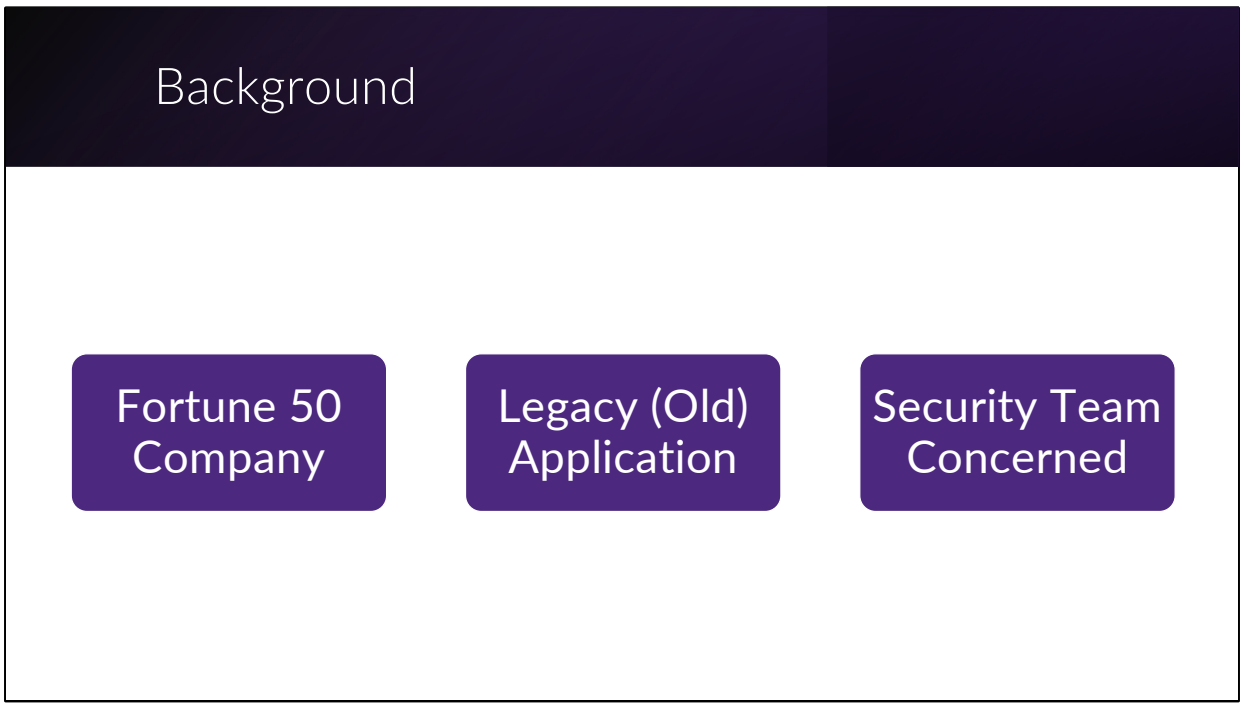

Some background for the Demo – the story so far – the security team at a Fortune 50 size company is concerned about the security risk of the widget management system in use at the company. The widget system is critical for their business and for reselling the company's B2B services to their clients. It was installed over 30 years ago, so the technology is significantly out of date. The security team is worried that it could be a source of a breach and starts a Risk Quant analysis. They soon learn that the infrastructure team is also concerned about frequent outages and near-outages. In the initial analysis, the team realizes that they don't have a good understanding of the cost of either a breach or an outage and connect with SMEs on the business side. When asked "what else are we missing?" the business team shares their concern that they are losing both current and prospective customers due to the functional obsolescence of the widget system; one customer has already left, and more are expected due to increasing competition in the widget space. The RQ team interviews several experts on these three risks.

Demo

Show me, don't tell me!

The demo is meant to be interactive – please ask questions! Spreadsheet > Report > Code Walkthrough.

https://jabenninghoff.github.io/security/analysis/rq-demo.html

Practical Advice

SCOPING AND SCHEDULING

OUTLIER EXPERTS

MODELING AND COMMUNICATION

Some common themes I've heard from those who have done risk quantification:
1. Scoping and Scheduling: the main challenges are scoping the risks, finding the experts, and scheduling time for the interviews. The time spent and running the models is comparatively easy.
2. Outlier Experts: sometimes an expert is far different from the rest. Methods of weighting expert opinion don't improve the estimate, and typically one expert won't change the story much.
3. Modeling and Communication: in practice, the model is less important – it's primary value is in facilitating discussion, discovery, and bringing knowledge from front-line workers to management.

Slides, Connect & Resources

Connect:
linkedin.com/in/jbenninghoff/

Website:
security-differently.com

Resources:
cyentia.com
thevoid.community

Thank you!

Appendix

Slides that got cut but are potentially useful.

Westrum's Organizational Typology

Pathological	Bureaucratic	Generative
Power oriented	Rule oriented	Performance oriented
Low cooperation	Modest cooperation	High cooperation
Messengers "shot"	Messengers neglected	Messengers trained
Responsibilities shirked	Narrow responsibilities	Risks are shared
Bridging discouraged	Bridging tolerated	Bridging encouraged
Failure leads to scapegoating	Failure leads to justice	Failure leads to inquiry
Novelty crushed	Novelty leads to problems	Novelty implemented

In pathological organizations, the focus is on the leader's power and advancement
Bureaucratic organizations focus on departmental goals
Generative environments focus on the organization's mission above all else

From the DORA summary at https://dora.dev/capabilities/generative-organizational-culture/