Empirical Security & Privacy, for Humans

UPenn CIS 7000-010

9/9/2025

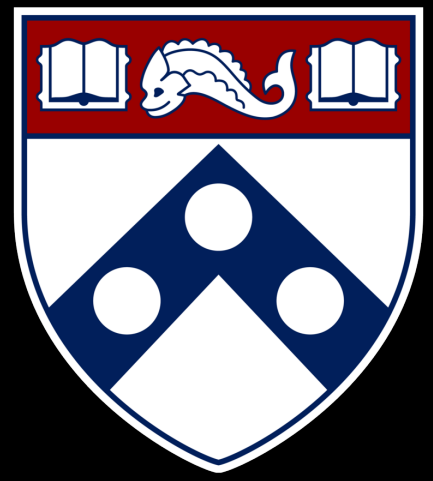

Security – Risk and uncertainty

Readings

JOURNAL OF CYBERSECURITY

Journal of Cybersecurity, 2024, tyae022
https://doi.org/10.1093/cybsec/tyae022

Article

Into the unknown: the need to reframe risk analysis

Andrew Simpson *

Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD, United Kingdom

*Corresponding author. E-mail: andrew.simpson@cs.ox.ac.uk

Received 4 December 2023; revised 31 July 2024; accepted 9 October 2024

Abstract

In recent years there have been efforts to bring a degree of quantification to the task of security risk analysis. Various arguments in favour of such developments have been offered: 'checklist'- or 'tickbox'-based security is insufficiently dynamic; risk matrices are flawed; quantitative approaches *must* (somehow) be better than qualitative ones; it makes sense to leverage advances in data science, AI, and machine learning in concert with the increasing abundance of data; there is merit in leveraging lessons from economics. While some notes of caution have been offered in the literature (with data availability and quality being prominent concerns), we argue that greater consideration and recognition of the relationship between risk and uncertainty—and, indeed, unawareness—would be of value to the community. In doing so, we look to recent critiques of the prevailing economics orthodoxy before considering potential sources of possible help.

Keywords: risk analysis; uncertainty; information security economics

Introduction

It is well understood that the notion of *risk* is at the heart of information security management.[1] In addition, it is widely recognized that recent technological developments have had a dramatic impact upon the threat landscape. To quote Wheatley *et al.* [2]: "cyber technologies have ushered in rapidly evolving cyber risks."

Such considerations are at the heart of the EU's General Data Protection Regulation (GDPR).[2] For example, Article 35 of the GDPR—

Partly in response to such developments, there have, in recent years, been a number of related and overlapping research efforts that have influenced how risk is modelled. One such effort—what might be characterized as *Data-Driven Security*—is argued for by Jacobs and Rudis in their textbook, *Data-Driven Security: Analysis, Visualization and Dashboards* [5], thus:

> "The era of the security shamen is rapidly fading, and it's time

Minerva (2017) 55:229–248
DOI 10.1007/s11024-017-9322-4

The Aviation Paradox: Why We Can 'Know' Jetliners But Not Reactors

John Downer[1]

Published online: 7 June 2017
© The Author(s) 2017. This article is an open access publication

Abstract Publics and policymakers increasingly have to contend with the risks of complex, safety-critical technologies, such as airframes and reactors. As such, 'technological risk' has become an important object of modern governance, with state regulators as core agents, and 'reliability assessment' as the most essential metric. The Science and Technology Studies (STS) literature casts doubt on whether or not we should place our faith in these assessments because predictively calculating the ultra-high reliability required of such systems poses seemingly insurmountable epistemological problems. This paper argues that these misgivings are warranted in the nuclear sphere, despite evidence from the aviation sphere suggesting that such calculations can be accurate. It explains why regulatory calculations that predict the reliability of new airframes cannot work in principle, and then it explains why those calculations work in practice. It then builds on this explanation to argue that the means by which engineers manage reliability in aviation is highly domain-specific, and to suggest how a more nuanced understanding of jetliners could inform debates about nuclear energy.

Keywords Engineering · Reliability · Risk · Safety · Regulation · Technology assessment · Nuclear energy · Civil aviation · Jetliners · Reactors

If politics means anything today it must become 'the art of the impossible.'
~Lewis Mumford (1954: 7)

Readings

Minerva (2017) 55:229–248
DOI 10.1007/s11024-017-9322-4

CrossMark

The Aviation Paradox: Why We Can 'Know' Jetliners But Not Reactors

John Downer[1]

Published online: 7 June 2017
© The Author(s) 2017. This article is an open access publication

Abstract Publics and policymakers increasingly have to contend with the risks of complex, safety-critical technologies, such as airframes and reactors. As such, 'technological risk' has become an important object of modern governance, with state regulators as core agents, and 'reliability assessment' as the most essential metric. The Science and Technology Studies (STS) literature casts doubt on whether or not we should place our faith in these assessments because predictively calculating the ultra-high reliability required of such systems poses seemingly insurmountable epistemological problems. This paper argues that these misgivings are warranted in the nuclear sphere, despite evidence from the aviation sphere suggesting that such calculations can be accurate. It explains why regulatory calculations that predict the reliability of new airframes cannot work in principle, and then it explains why those calculations work in practice. It then builds on this explanation to argue that the means by which engineers manage reliability in aviation is highly domain-specific, and to suggest how a more nuanced understanding of jetliners could inform debates about nuclear energy.

Keywords Engineering · Reliability · Risk · Safety · Regulation · Technology assessment · Nuclear energy · Civil aviation · Jetliners · Reactors

If politics means anything today it must become 'the art of the impossible.'
~Lewis Mumford (1954: 7)

Ultra-high reliability assessments

- Reliability like security – a negative property.
  - Forces it to be contextual: You have to define the circumstances under which you agree the bad thing can't happen, and you have to carefully define the parameters of the bad thing
- Predictive reliability assessment goal: 1B hours of failure-free service
- Approach: Model-based analysis
  - Assess individual components **inductively**
  - Reason about their combined reliability **deductively**, using a model
- Example: Redundant engines
  - Single engine failure probability P, based on measurements
  - Total engine failure of two engines = P^2
    - Assuming engine failures are independent

Ultra-high reliability assessments, questioned

- Epistemically dubious!
  - “There is no way that experts should be able to deduce from tests and models that a yet-unrealized jetliner or reactor will be reliable to the extraordinary levels that they claim.”
  - Why? We cannot observe the tech in action long enough, or draw on other prior evidence, to extrapolate to the hoped-for conclusion.
    - 1B hours is 114,000(ish) years! Long time to run to test directly
- Nevertheless: “A stubbornly suicidal traveler would have to take a random airline flight every day for 19,000 years to stand a better-than-even chance of succumbing to a fatal crash.”

Why it works for (most) civil aviation

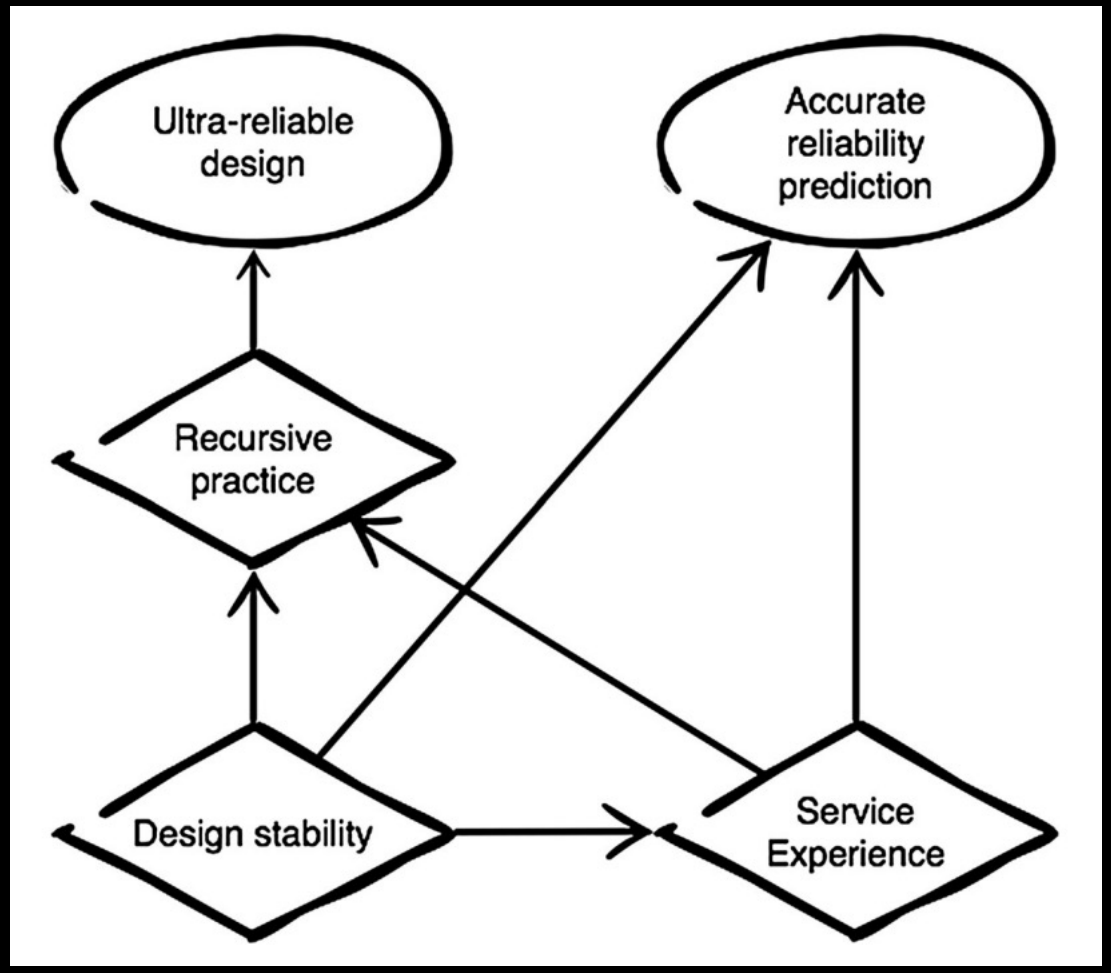

But not the Concorde

But not the Concorde

1 catastrophic failure, July 25, 2000

Very different design

- 1354 mph cruising speed
- ‘Double delta’ wing shape
- Unorthodox landing gear and a nose that moved
- Special heat-resistant alloys

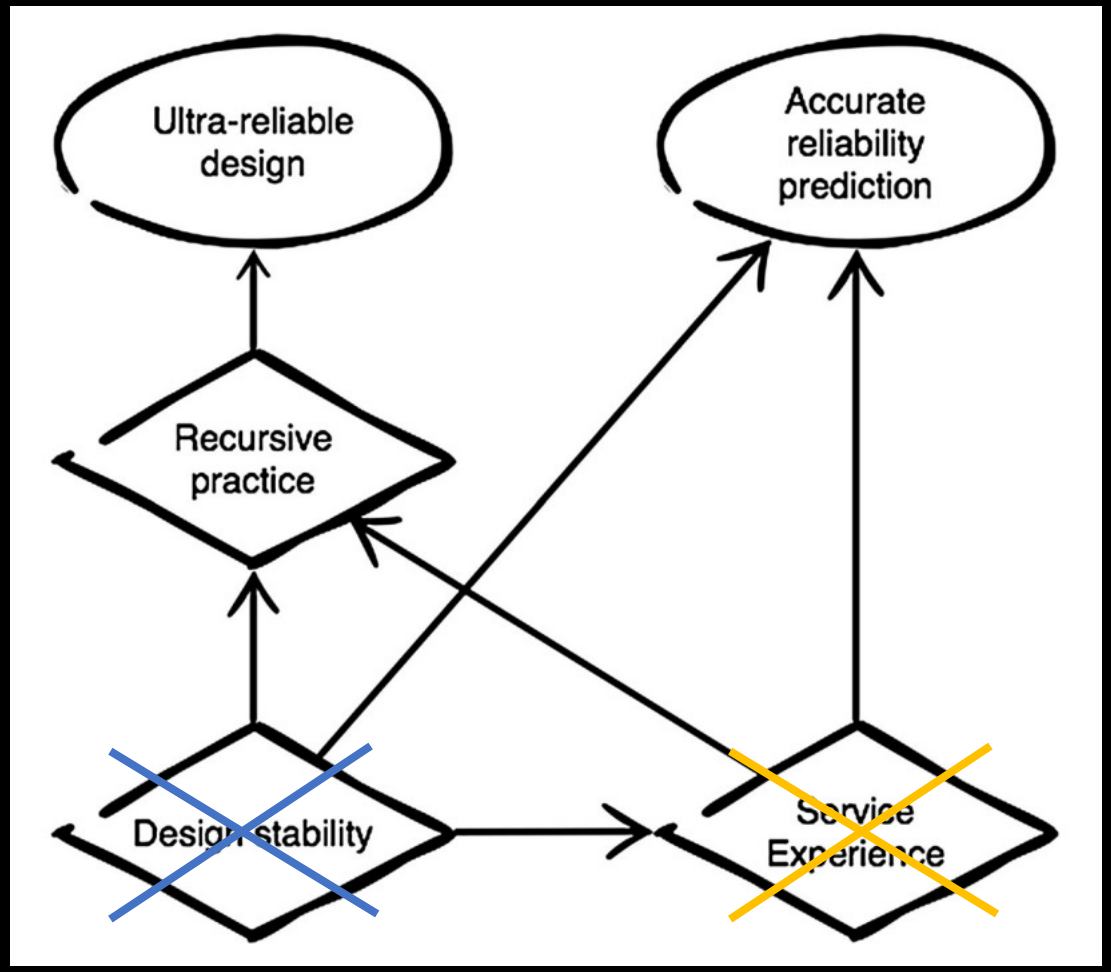

Only 14 planes in operation totaling 50,000 flights over 27 years

And not Nuclear

Lessons are learned and improvements made, but do not generalize

Reactor designs broadly dissimilar

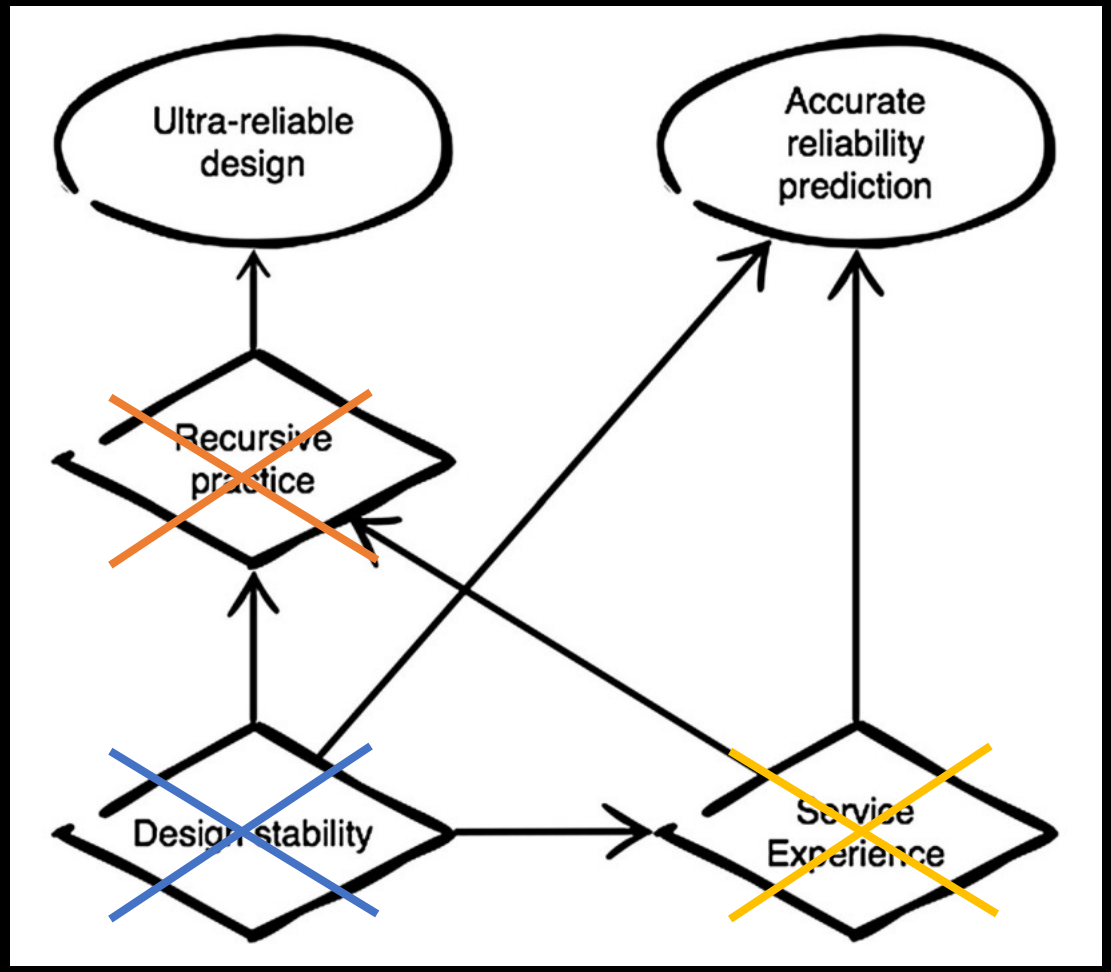

Not enough reactors, not enough service time for each

Themes, discussion points

- Situation on the ground since this paper was published (2017)
- Role and kinds of uncertainty in modeling (next paper!)
- Comparing aviation to cybersecurity
- Consequences of applying the aviation process, given all this

Thoughts?

Readings

JOURNAL OF CYBERSECURITY

Journal of Cybersecurity, 2024, tyae022
https://doi.org/10.1093/cybsec/tyae022

Article

Into the unknown: the need to reframe risk analysis

Andrew Simpson *

Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD, United Kingdom

*Corresponding author. E-mail: andrew.simpson@cs.ox.ac.uk

Received 4 December 2023; revised 31 July 2024; accepted 9 October 2024

Abstract

In recent years there have been efforts to bring a degree of quantification to the task of security risk analysis. Various arguments in favour of such developments have been offered: 'checklist'- or 'tickbox'-based security is insufficiently dynamic; risk matrices are flawed; quantitative approaches *must* (somehow) be better than qualitative ones; it makes sense to leverage advances in data science, AI, and machine learning in concert with the increasing abundance of data; there is merit in leveraging lessons from economics. While some notes of caution have been offered in the literature (with data availability and quality being prominent concerns), we argue that greater consideration and recognition of the relationship between risk and uncertainty—and, indeed, unawareness—would be of value to the community. In doing so, we look to recent critiques of the prevailing economics orthodoxy before considering potential sources of possible help.

Keywords: risk analysis; uncertainty; information security economics

Introduction

It is well understood that the notion of *risk* is at the heart of information security management.[1] In addition, it is widely recognized that recent technological developments have had a dramatic impact upon the threat landscape. To quote Wheatley *et al.* [2]: "cyber technologies have ushered in rapidly evolving cyber risks."

Such considerations are at the heart of the EU's General Data Protection Regulation (GDPR).[2] For example, Article 35 of the GDPR—titled *Data Protection Impact Assessment*—prescribes the following:

Partly in response to such developments, there have, in recent years, been a number of related and overlapping research efforts that have influenced how risk is modelled. One such effort—what might be characterized as *Data-Driven Security*—is argued for by Jacobs and Rudis in their textbook, *Data-Driven Security: Analysis, Visualization and Dashboards* [5], thus:

> "The era of the security shamen is rapidly fading, and it's time to adopt the proven tools and techniques being used in other dis-

Security is risk, but risks are uncertain

- “It is well understood that the notion of *risk* is at the heart of information security management.”
- “in this paper we argue that
  - (a) greater caution is needed when embracing quantitative approaches and
  - (b) greater consideration should be given to the relationship that exists between risk and uncertainty”

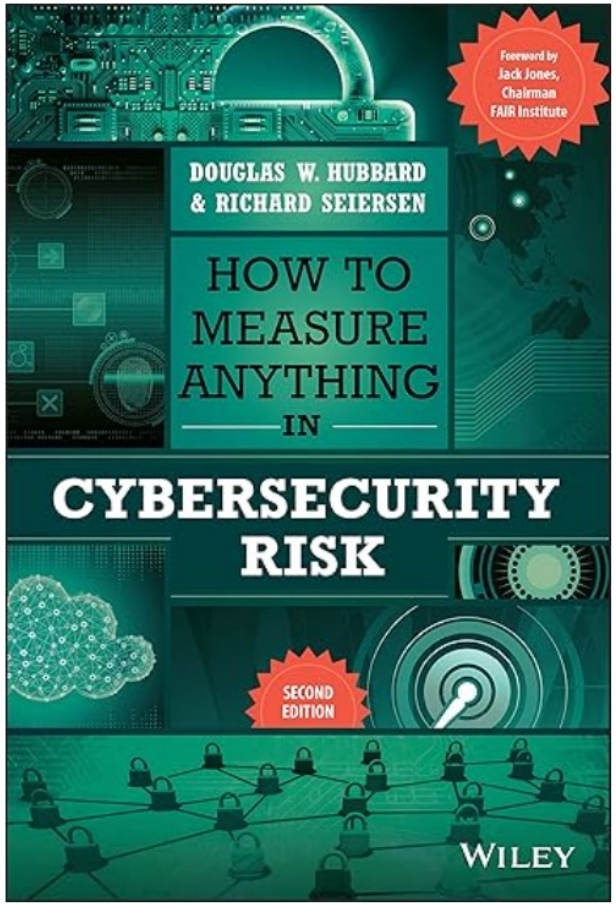

How to Measure Anything in Cybersecurity Risk 2nd Edition

by Douglas W. Hubbard (Author), Richard Seiersen (Author)

4.5 ★★★★½ 88 ratings

See all formats and editions

A start-to-finish guide for realistically measuring cybersecurity risk

In the newly revised *How to Measure Anything in Cybersecurity Risk, Second Edition*, a pioneering information security professional and a leader in quantitative analysis methods delivers yet another eye-opening text applying the quantitative language of risk analysis to cybersecurity. In the book, the authors demonstrate how to quantify uncertainty and shed light on how to measure seemingly intangible goals. It's a practical guide to improving risk assessment with a straightforward and simple framework.

Advanced methods and detailed advice for a variety of use cases round out the book, which also includes:

- A new "Rapid Risk Audit" for a first quick quantitative risk assessment.
- New research on the real impact of reputation damage

with other portfolios. The aggregation process is typically some form of invented mathematics unfamiliar to actuaries, statisticians, and mathematicians.

Just over 50% of respondents plot risks on a two-dimensional matrix. In this approach, "likelihood" and "impact" will be rated subjectively, perhaps on a 1 to 5 scale, and those two values will be used to plot a particular risk on a matrix (variously called a "risk matrix," "heat map," "risk map," etc.). The matrix—similar to the one shown in Figure 1.1—is then often further divided into sections of low, medium, and high risk. Events with high likelihood and high impact would be in the upper-right "high risk" corner, while those with low likelihood and low impact would be in the opposite "low risk" corner. The idea is that the higher the score, the more important something is and the sooner you should address it. You may intuitively think such an approach is reasonable, and if you thought so, you would be in good company.

Various versions of scores and risk maps are endorsed and promoted by several major organizations, standards, and frameworks such as the National Institute of Standards and Technology (NIST), the International Standards Organization (ISO), MITRE.org, and the Open Web Application Security Project (OWASP). Most organizations with a cybersecurity function claim at least one of these as part of their framework for assessing risk. In fact, most major software organizations such as Oracle, Microsoft, and Adobe rate their vulnerabilities using a NIST-supported scoring system called the "Common Vulnerability Scoring System" (CVSS). Many security solutions also include CVSS ratings, be it for vulnerability and/or attack related. While the control recommendations made by many of these frameworks are good,

			Impact				
			Negligible	Minor	Moderate	Critical	Catastrophic
			1	2	3	4	5
Likelihood	Frequent	5	Medium	Medium	High	High	High
	Likely	4	Medium	Medium	Medium	High	High
	Occasional	3	Low	Medium	Medium	Medium	High
	Seldom	2	Low	Low	Medium	Medium	Medium
	Improbable	1	Low	Low	Low	Medium	Medium

FIGURE 1.1 The Familiar Risk Matrix (aka Heat Map or Risk Map)

it's how we are guided to prioritize risk management on an enterprise scale that is amplifying risk.

Literally hundreds of security vendors and even standards bodies have come to adopt some form of scoring system including the risk matrix. Indeed, scoring approaches and risk matrices are at the core of the security industry's risk management approaches.

In all cases, they are based on the idea that such methods are beneficial to some degree. That is, they are assumed to be at least an improvement over not using such a method. As one of the standards organizations has put it, rating risk this way is adequate:

> *Once the tester has identified a potential risk and wants to figure out how serious it is, the first step is to estimate the* likelihood. *At the highest level, this is a rough measure of how likely this particular vulnerability is to be uncovered and exploited by an attacker. It is not necessary to be over-precise in this estimate.* Generally, identifying whether the likelihood is low, medium, or high is sufficient. *(emphasis added)*
>
> —OWASP[20]

Does this last phrase, stating "low, medium, or high is sufficient," need to be taken on faith? Considering the critical nature of the decisions such methods will guide, we argue that it should not. This is a testable hypothesis, and it actually *has been* tested in many different ways. The growing trends of cybersecurity attacks alone indicate it might be high time to try something else.

So, let's be clear about our position on current methods: *They are a failure. They do not work.* A thorough investigation of the research on these methods and decision-making methods in general indicates the following (all of this will be discussed in detail in later chapters):

- There is no evidence that the types of scoring and risk matrix methods widely used in cybersecurity improve judgment.
- On the contrary, there is evidence these methods add noise and error to the judgment process. One researcher we will discuss more—Tony Cox—goes as far as to say they can be "worse than random."
- Any appearance of "working" is probably a type of "analysis placebo." That is, a method may make you feel better even though the activity provides no measurable improvement in estimating risks (or even adds error).
- There is overwhelming evidence in published research that quantitative, probabilistic methods are an improvement over unaided expert intuition.

Quantitative risk assessment questioned

- Qualitative ("tickbox exercise") vs. quantitative risk assessment
- Problems: Do we have good estimates or data to work from? Does it make epistemic sense, given an active adversary?
- Idea: Embrace "Radical uncertainty". Why?
  - (a) uncertainty is at the heart of risk management
  - (b) precision can be spurious,
  - (c) probabilities can hide uncertainty [what does this mean?!]
  - (d) while it is undeniably the case that it is essential to plan, it can be dangerous to pretend to know

Themes

- Small world vs. large world for modeling
- Are probabilities representing knowable outcomes, or uncertainties?
- Unawareness (unknown unknowns) vs. uncertainty (known unknowns)
- “It is possible to conclude that there is a need, as a community, to reflect upon (a) what lessons we might *meaningfully* take from economics, and (b) the limits of quantification and modeling.”

Thoughts?

END