Trusting and Defending Risk Measurements

Jack Jones, Chairman, FAIR Institute

FAIR22 CONFERENCE

What are the implications if cyber risk is measured poorly?

What’s the difference between trusting a risk measurement versus being able to defend a risk measurement?

How do we decide what to trust?

Using a risk measurement to influence a decision (prioritization, solution selection, etc.) implies that we believe the measurement is accurate.

That we trust it.

10,000
Hours

When you perform a risk measurement, how do you know you got it right?

Practice without feedback results in…

Firmly entrenched habits of unknown efficacy.
Unwarranted confidence.

If we automate the same risk measurement methods our profession has been using for years…

Automation will become a huge part of cyber risk measurement in the future, but automation that is fundamentally flawed simply amplifies the effects of poor risk measurement.

What’s the difference?

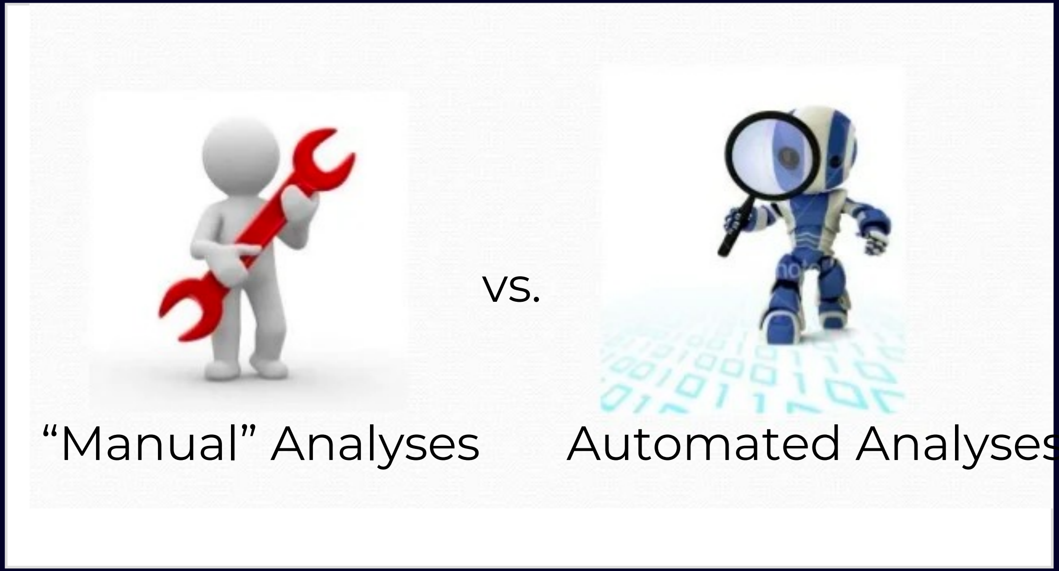

“Manual” Analyses

vs.

Automated Analyses

All risk analyses involve assumptions.

What, exactly, is being measured? (the scope)

Model parameters and relationships

What data to use, and how to use them

Example use-case…

Finding #1

An audit discovered that privileges are not consistently being updated for user accounts with access to a customer service application containing PII.

Finding #2

A security assessment determined that the organization was unlikely to be able to identify when a cyber criminal breaches its network perimeter because it has poor visibility and monitoring controls.

- Data source and form?
- Conversion to quantification?
- Which scenarios are these relevant to?
- Are other controls in place that are relevant to these scenarios?
- Recommended solutions for the findings?

Another use-case…

Using NIST CSF scores within a risk analysis

Control Subcategory	Rating
PR.PT-4 : Communications and control networks are protected	2

Is PR.PT-4 relevant to the scenario being analyzed?

How does PR.PT-4 affect risk?

What does a “2” represent?

Which other controls is PR.PT-4 dependent on, and what were their scores?

The point is…

- All risk measurements require making assumptions
- A key difference between automated and manual risk measurement is who’s making the assumptions
- If the assumptions are wrong, the measurement is almost certain to be inaccurate

What makes a risk measurement trustworthy?

FAIR22 CONFERENCE

First, some fundamentals…

- Risk measurements are either accurate, or they’re not
  - I.e., a measurement’s accuracy is a true/false question

Three estimates of my height:

- 5’6” - 6’0”
- 5’11” - 6’1”
- 5’11”

- Trustworthiness is based on the probability that a measurement is accurate
  - Higher probability of accuracy = more trustworthy
  - Lower probability of accuracy = less trustworthy
- What affects the probability of being accurate?

“Just” four key things…

- The clarity of a measurement’s scope
- Accurate and relevant input data
- A model that is logically and formulaically sound
- Results that faithfully reflect uncertainty (ranges & distributions)

Some accuracy (and trustworthiness) red flags...

- Unclear or undefined scope
- SME estimates are not calibrated
- Input values are ordinal data (1 thru 5, etc.) vs. ratio data
- Uses weighted values
- Outputs are discrete values vs. distributions
- Doesn't account for threat event frequency
- Doesn't account for control dependencies and relationships*

Defending Risk Measurements

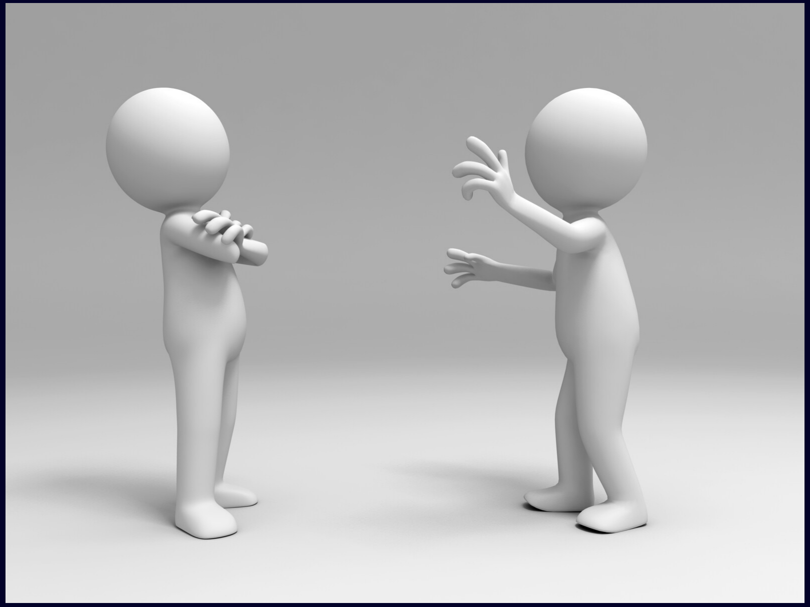

Especially at first, CRQ tends to bring out the curiosity/skepticism in stakeholders

As it should!

Two things you can base a defense on...

- You did the measurement yourself
  - You know and can explain the scope, data, and model that were used

- You trust the source of the measurement (and can explain why you trust the source)
  - You know how the measurement was done (scope, data, and model), or
  - The source has been independently validated

Making your own risk measurements defensible

- Clearly scope your measurements.
  - Assets at risk
  - Threat community (Cybercriminals, Nation-state actors, Insiders, Mother Nature, etc.)
  - Type of threat (Malicious, intentional but non-malicious, accidental, etc.)
  - Method, vector, etc.
  - Type of event (outage, data breach, fraud, etc.)
  - Relevant controls
- Choose an established model, or clearly define your own.
  - If it’s not your own model, understand its provenance.

Making your own risk measurements defensible

- Understand your data
  - Know where the data came from, as well as any concerns about data quality.
    - Industry data versus internal SME?
    - NOTE: Your data will never be perfect, and there are diminishing returns in digging for “enough” data.
- Faithfully represent uncertainty in inputs and outputs.
  - Do NOT express risk as a discrete value.
  - Use proven methods.
    - Calibration for SME estimates
    - Stochastic methods to account for uncertainty
    - Use ranges and distributions versus discrete values

Defending someone else’s measurement

- Be able to explain why your stakeholders should trust it
- How well do you understand how the results were arrived at?
- Has the model/method been independently vetted?
  - ~~Historical performance~~
  - Backtesting
  - Scenario testing
  - Methodological review

Understand the testing scope, parameters, and methods!

NOTE: Patents have absolutely nothing to do with validity!

https://www.investopedia.com/terms/b/backtesting.asp

When presenting results…

- Don’t shy away from your assumptions.
  - Understand the assumptions you’re making, and be prepared to explain them.
- Welcome skepticism and challenges to your measurements.
  - If you’ve done your homework, then it’s an opportunity to gain credibility.
  - Regardless, avoid being vain about your measurements.
- Feedback and discussion is how we improve.
- Remember that the goal is accurately-informed decisions, regardless of how that occurs.
- The process of getting to results is often just as valuable as the results themselves.

Wrapping up...

As a profession we have presumed as fact that the measurement methods we’ve been using for years are effective.

We can’t rationally defend that position.

Our problem space is highly complex.

Overly simplistic methods and models do not help, and are in fact, damaging.

Where is Darwin when you need him…

Without a feedback loop, there are no consequences <u>to the measurer</u> for being bad at risk measurement.

It’s the decision-makers and other stakeholders that pay the bill.

As the cybersecurity world leans into CRQ, measurement defensibility will become a bigger deal.

The bottom line...

If you can defend a risk measurement,
then it’s reasonable to trust it.

If you can’t defend a risk measurement,
it’s unreasonable to trust it.

Questions?