FAIR INSTITUTE

Today’s Cyber Risk Measurement Best Practices

Jack Jones
Chairman FAIR Institute

Why do we measure risk?

What single most significant cybersecurity risk your organization faces?

How much risk reduction did your organization get from its most recent major cybersecurity initiative?

You have two security-related findings…

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

A security assessment determined that the organization was unlikely to be able to identify when a cyber criminal breaches its network perimeter.

Which of them is more important to fix first?

What do the answers to those questions have in common?

Common Practice

Something that is done a lot and is considered “normal”

Macmillan Dictionary

Risk measurement models versus risk-related frameworks

Control Frameworks

ISO 2700x

CIS Center for Internet Security® Creating Confidence in the Connected World.

NIST 800-53

NIST CSF

Cybersecurity Maturity Models

C2M2 | Cybersecurity Capability Maturity Model

NIST CMMC

Risk Measurement Models

NIST 800-30

FAIR

Only these
actually provide
risk measurement

Commonly cited “top risks”…

- Insiders
- Reputation
- Phishing
- Ransomware
- Weak passwords
- Poor cyber hygiene

Most commonly cited “top risks” aren’t risks…

Insiders	Threat community
Reputation	Asset
Phishing	Method
Ransomware	Method
Weak passwords	Control deficiency
Poor cyber hygiene	Cause of deficient controls

The classic formula for risk

Risk = Likelihood x Impact

Likelihood and Impact of what?

Loss Event Scenarios

These aren't loss events

- Insiders
- Reputation
- Phishing
- Ransomware
- Weak passwords
- Poor cyber hygiene

You can only assign likelihood and impact to loss event scenarios.

How fast are they going?

Qualitatively

Challenges...

- Is your “Fast” the same as mine?

- Which car am I referring to?
  - One in particular? (Slowest? Fastest?)
  - An average for all of them?

- Which part of the track am I referring to?
  - Corners?
  - The straightaway?
  - Average over the entire track?
  - This lap, or an average for the entire race?

Without clear scoping, the odds of measuring risk accurately are much lower…

…regardless of whether you’re doing qualitative or quantitative measurement

What’s the most commonly used cybersecurity risk model?

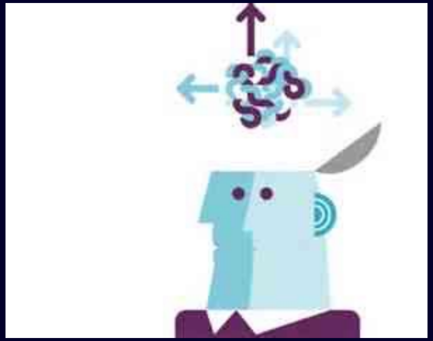

Mental models

What scope? What model? What data?

A broken cybersecurity risk model

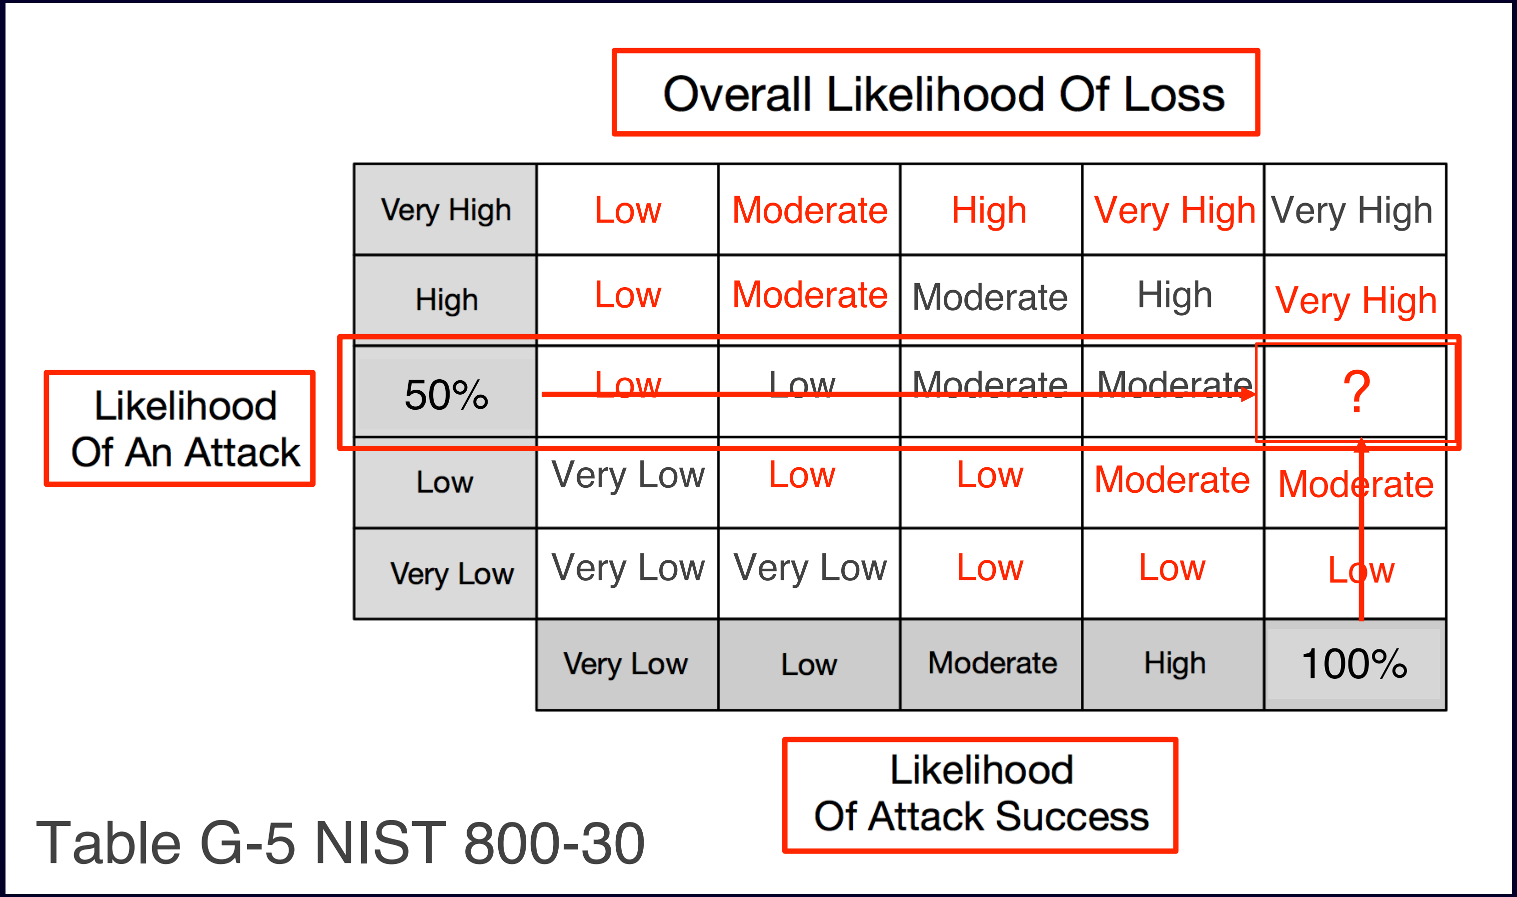

Overall Likelihood Of Loss

Likelihood Of An Attack

Very High	Low	Moderate	High	Very High	Very High
High	Low	Moderate	Moderate	High	Very High
50%	Low	Low	Moderate	Moderate	?
Low	Very Low	Low	Low	Moderate	Moderate
Very Low	Very Low	Very Low	Low	Low	Low
	Very Low	Low	Moderate	High	100%

Likelihood Of Attack Success

Table G-5 NIST 800-30

Math on colors

(🟩 Green x 🟨 Yellow) / 🟥 Red = ?

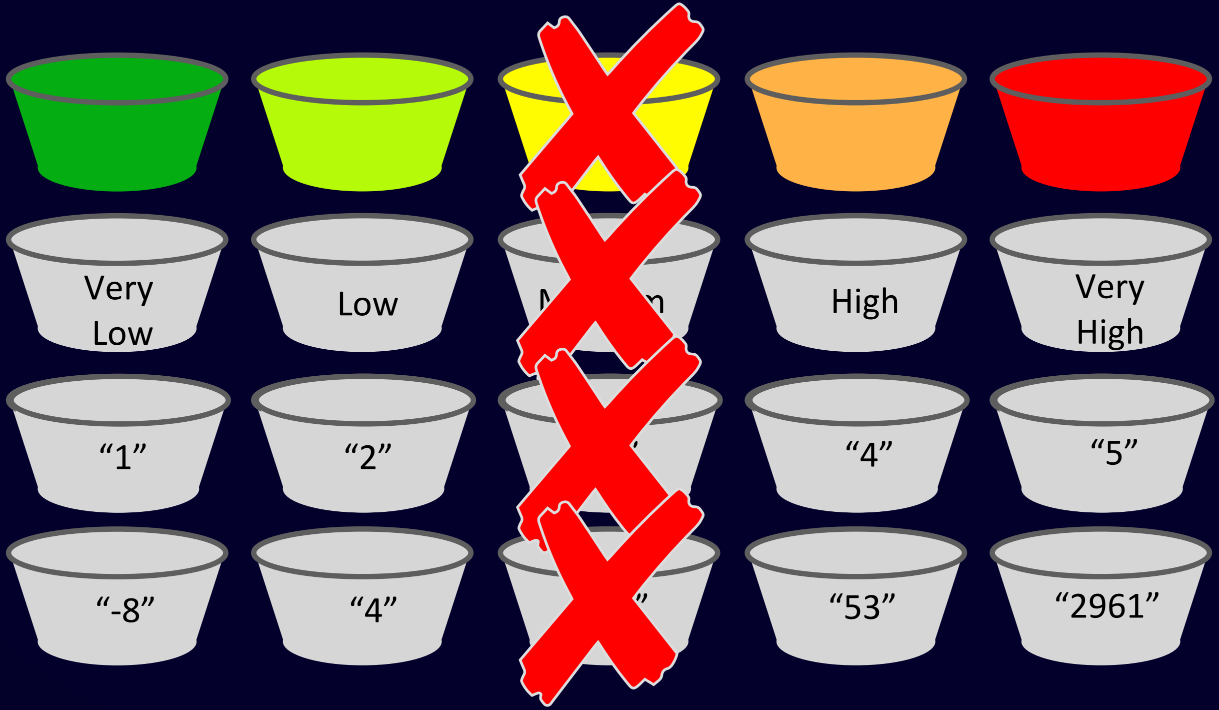

Over 70% of “high risk” findings aren’t, in fact, high risk.

No organization I’ve encountered in the past 5 years had accurately identified their top 10 cyber-related risks.

Key take-aways…

- We exist as a profession to help our organizations manage the frequency and magnitude of loss event scenarios.

- Today’s common risk measurement practices DO NOT support that objective.

Best Practice

The most effective way to do something.

Macmillan Dictionary

This is NOT a question of qualitative vs. quantitative measurement.

Three criteria for accurate risk measurement…

1. Clarity about what’s being measured
2. An accurate risk model
3. Accurate data

What are we measuring?

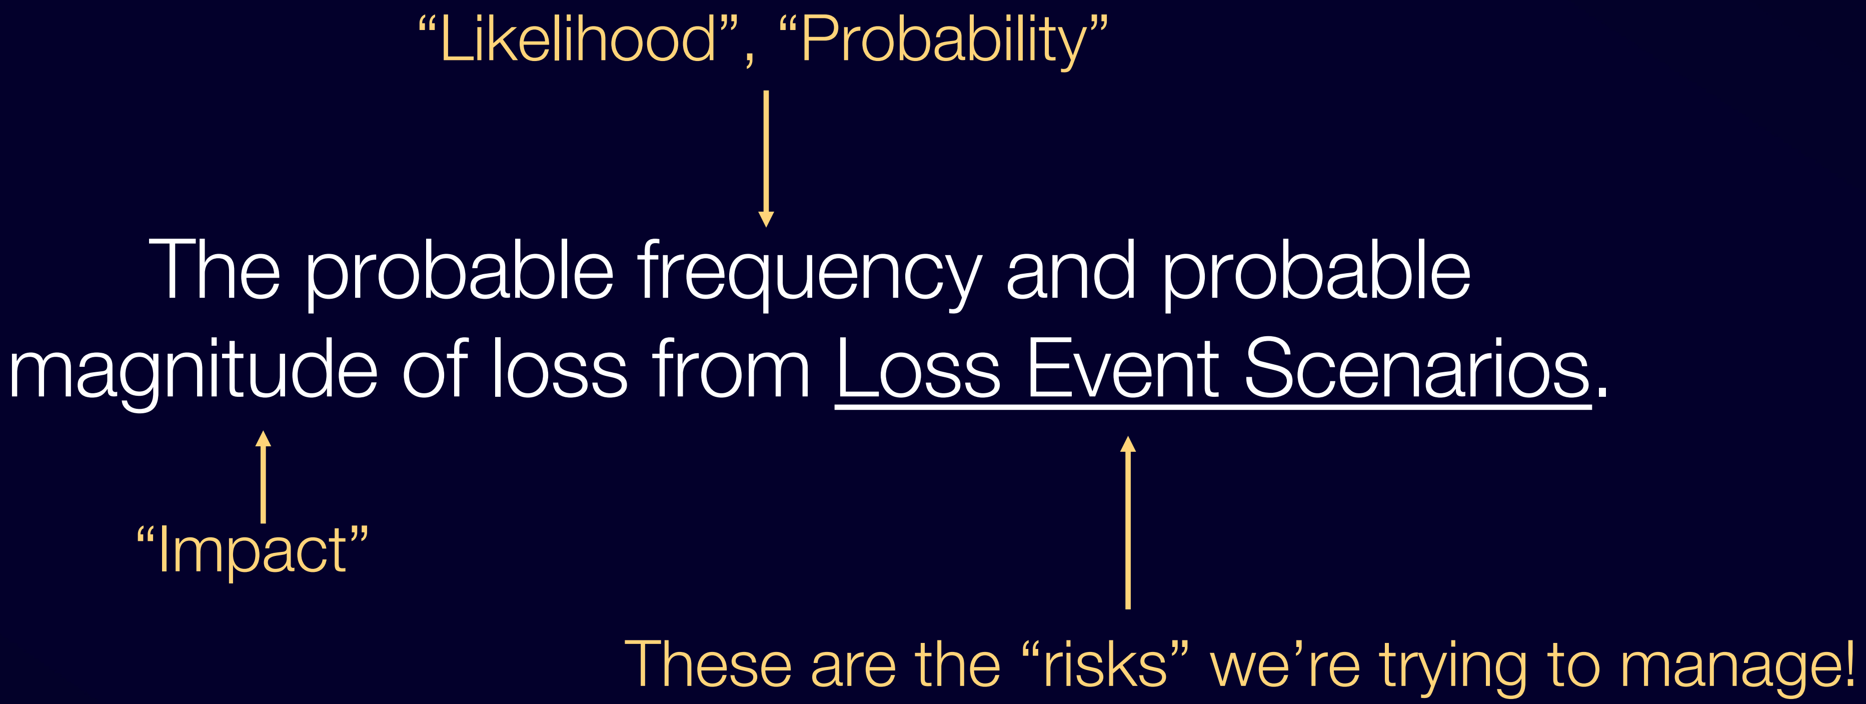

"Likelihood", "Probability"

The probable frequency and probable magnitude of loss from Loss Event Scenarios.

"Impact"

These are the "risks" we're trying to manage!

Example of a clearly scoped risk (loss event scenario)

Outage of key business systems.due to cybercriminals performing a ransomware attack via a phishing e-mail.

Asset
Threat
Effect
Vector
Method

Example of a clearly scoped risk (loss event scenario)

Disclosure of sensitive government documents.by a malicious insider who misuses their privileged access.

Asset
Threat
Effect
Method

Example of a clearly scoped risk (loss event scenario)

Corruption of customer financial information due to unintentional coding errors by software engineers in a new software release.

Asset
Threat
Effect
Method
Vector

Key take-aways…

Without clear scoping, the odds of measuring risk accurately are much lower…

…regardless of whether you're doing qualitative or quantitative measurement

Models

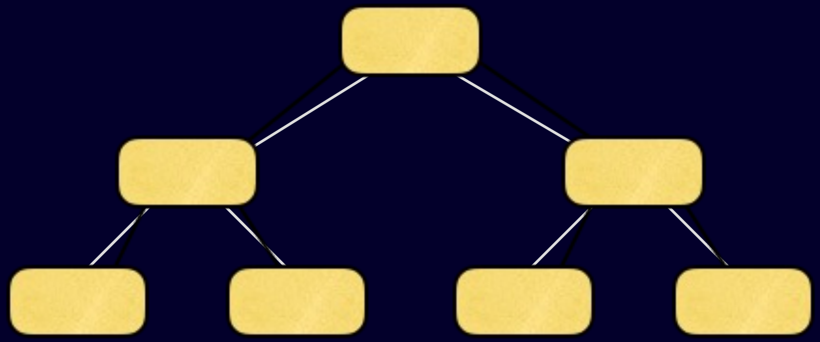

What is a model?

Models are simplified representations of a more complex reality.

Minimum risk model requirement

It must include measurement of both the probability and magnitude of loss.

Remember these?

Control Frameworks

ISO 2700x

CIS Center for Internet Security® Creating Confidence in the Connected World.

NIST CSF

NIST 800-53

Cybersecurity Maturity Models

C2M2 Cybersecurity Capability Maturity Model

NIST CMMC

Risk Measurement Models

NIST 800-30

FAIR

Only these include measurement of probability and magnitude.

The FAIR Model

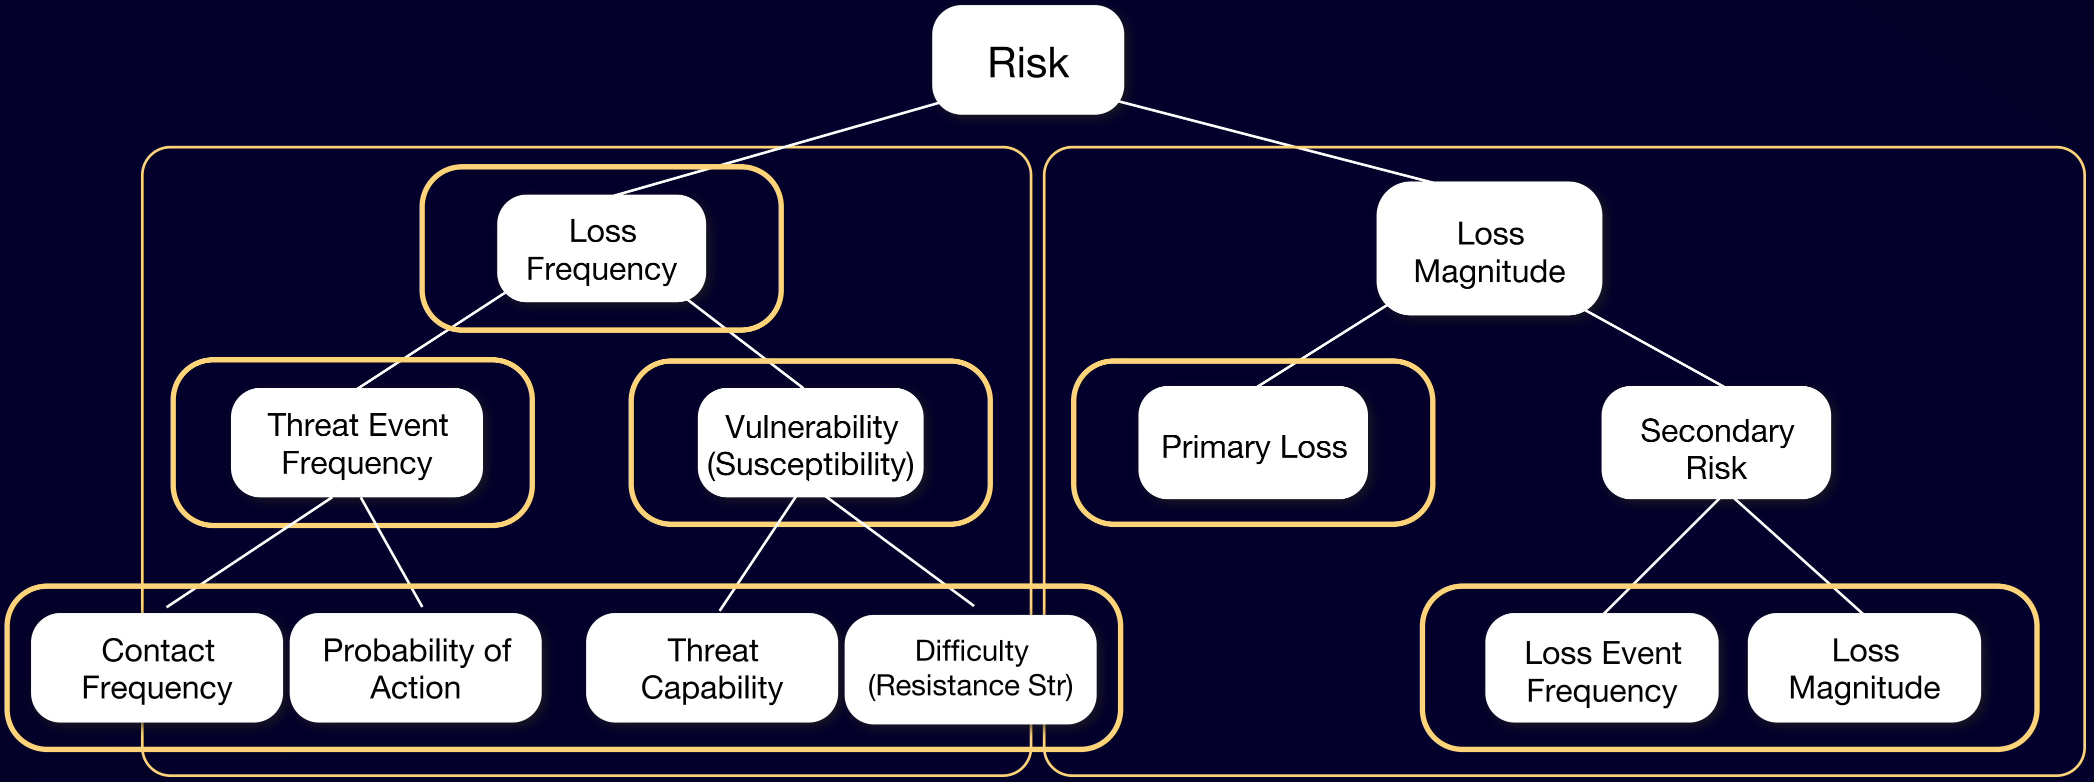

Loss Event Frequency

Loss Magnitude

Key take-aways…

Risk measurement models enable the measurement of loss event frequency and magnitude.

All risk measurement models involve assumptions.

Open models enable us to understand, challenge, and accept (or not) those assumptions.

But what about data?

What data do we need?

The FAIR Model

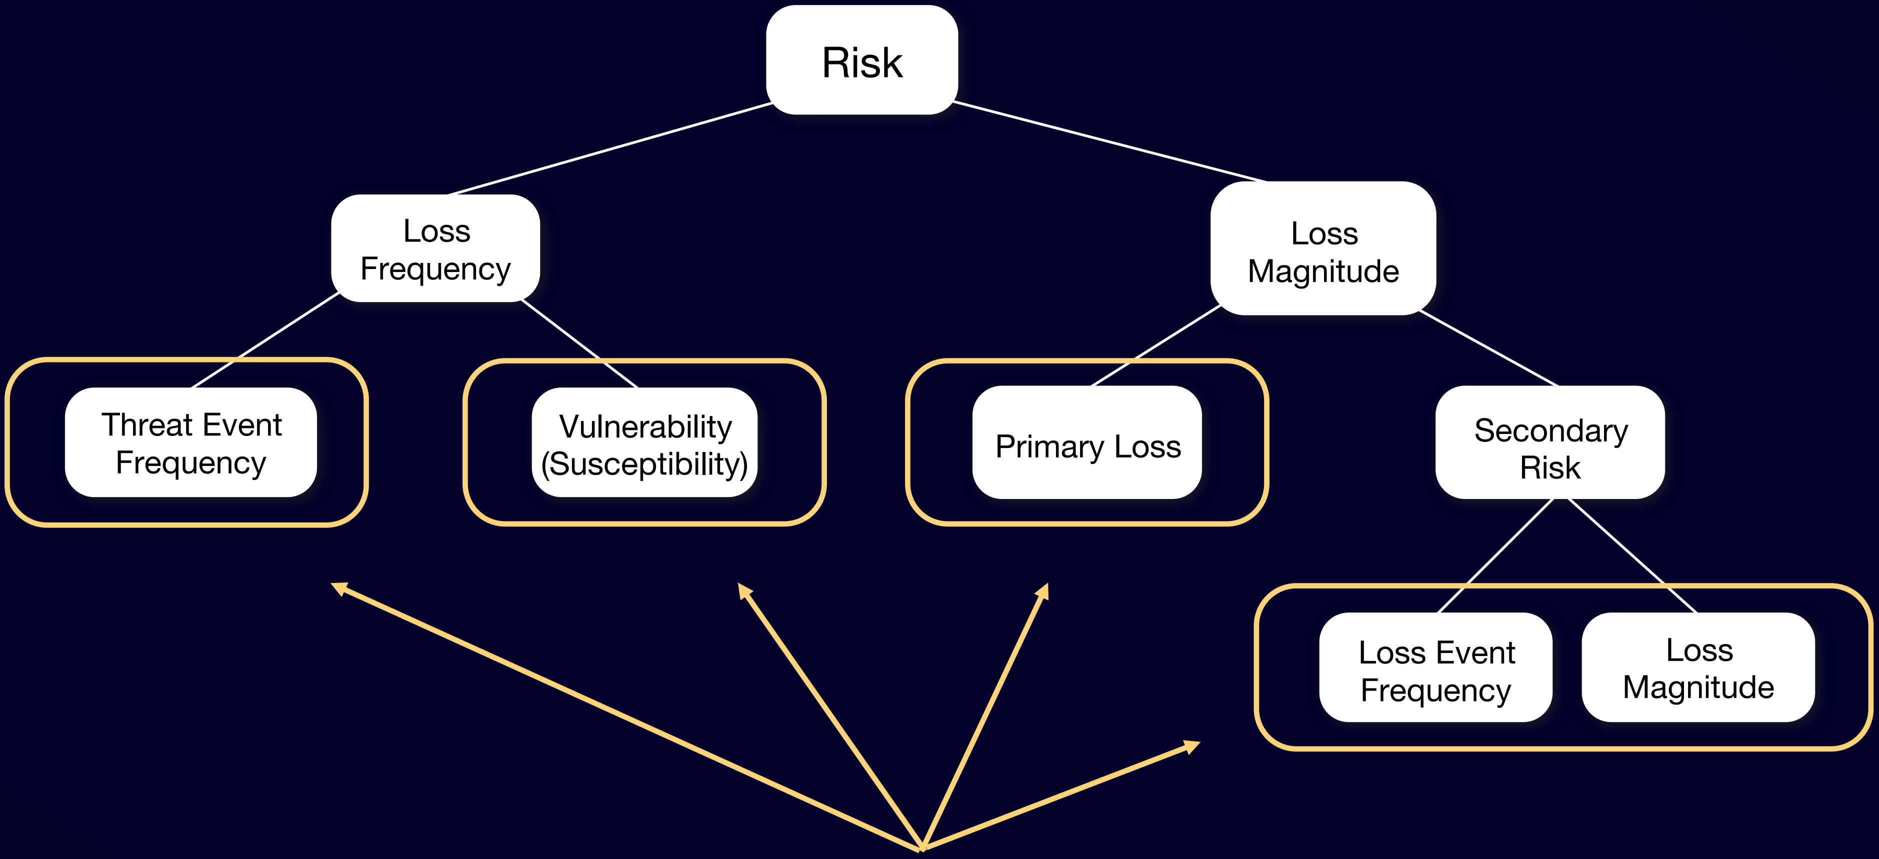

Data related to these…

“We don’t have enough data.”

- “You have more data than you think you do.”
- “You need less data than you think you do.”

Douglas Hubbard

How tall am I?

Example

What is the wingspan of a Boeing 747?

- 1 to 1000 feet?
- 50 to 500 feet?
- 100 to 300 feet?
- 125 to 250 feet?

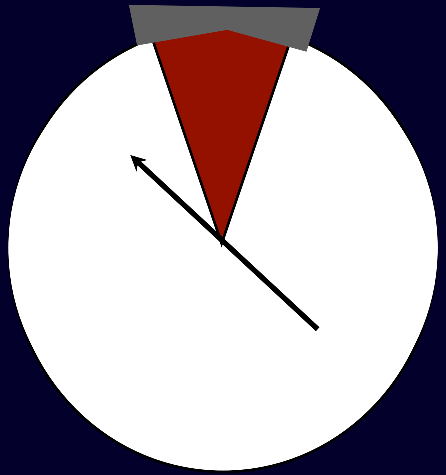

90% probability of landing in the white area.

For $1,000 place your bet…

What’s the difference between a guess and an estimate?

An estimate is something you would place a meaningful bet on. You wouldn’t place that same bet on a guess.

The problem of uncertainty…

Uncertainty is inevitable. It’s simply a matter of whether it’s accounted for in measurement inputs and outputs.

Using ranges and distributions to faithfully reflect uncertainty is crucial for accurate quantitative risk measurement.

Key take-aways…

Data scarcity is never a legitimate argument for not doing quantitative risk measurement.

There are well-established methods for dealing with sparse data

You have the data you have. You just need to faithfully represent uncertainty in your inputs and outputs.

Example Analyses

Remember these?

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

A security assessment determined that the organization was unlikely to be able to identify when a cyber criminal breaches its network perimeter.

Which one is more important to fix first?

Scoping the 1st analysis…

- What is the asset at risk? Customer information
- Who/what is the threat actor(s)? Personnel with inappropriate access
- What type of action Malicious
- What type of event is it (C, I, or A)? Confidentiality
- What is the loss event scenario? The confidentiality of customer data is maliciously compromised by personnel with inappropriate access

This is the risk →

Threat Event Frequency

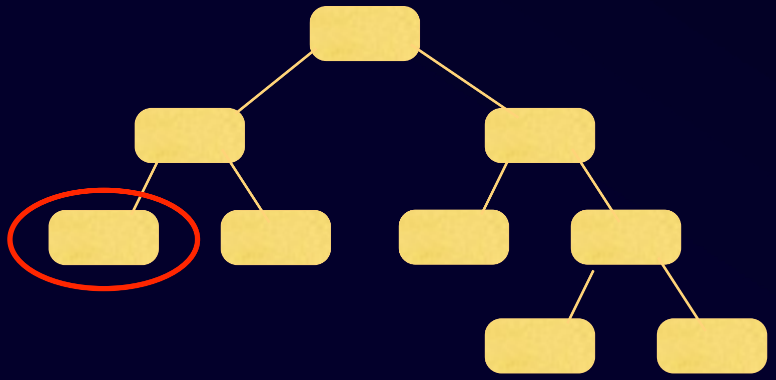

- Definition

The probable frequency, within a given timeframe, that a threat will act in a manner that may result in loss

- Estimates

Qualitative? Low

Min: .05 yr (1 in 20 yr)

Max: 5 yr (Logging!)

ML: .1 yr (1 in 10 yr)

- Data/Rationale
  - 30 user accounts (out of 200) with inappropriate access levels (15%)
  - HR records show 2 events of misuse in the past 3 yrs ("snooping")
  - Snooping was performed by personnel <u>with appropriate access</u>
  - No history of malicious misuse

Vulnerability

- Definition

The probability that a threat event will become a loss event

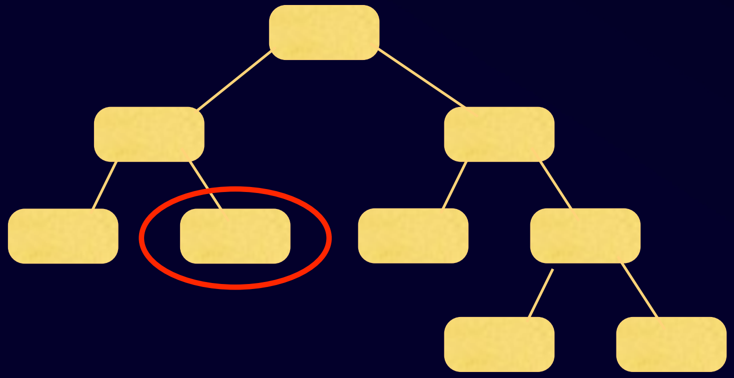

- Estimates

Qualitative? High

100%

- Data/Rationale
  - These are privileged insiders who don't have to overcome controls in order to execute the illicit action

Primary Loss Magnitude

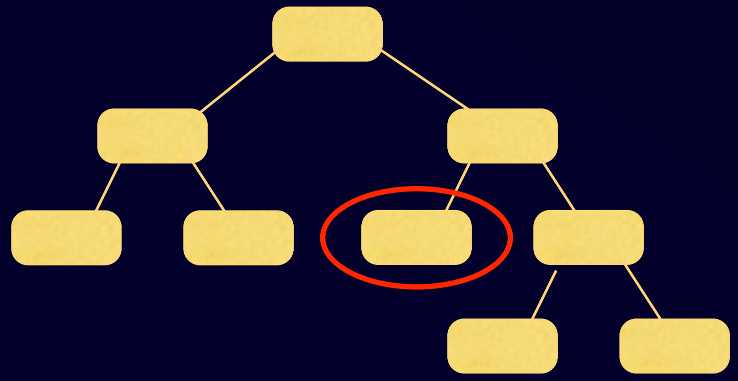

- Definition

  Loss that occurs directly as a result of the threat act against the asset.

- Estimates

  Qualitative? Moderate

  Min: $ 25k

  Max: $ 150k

  ML: $ 40k

- Data/Rationale
  - Forensic/investigative costs
  - Costs associated with replacing the malicious employee

Secondary Loss Event Frequency

- Definition

The probability of secondary loss (fallout)

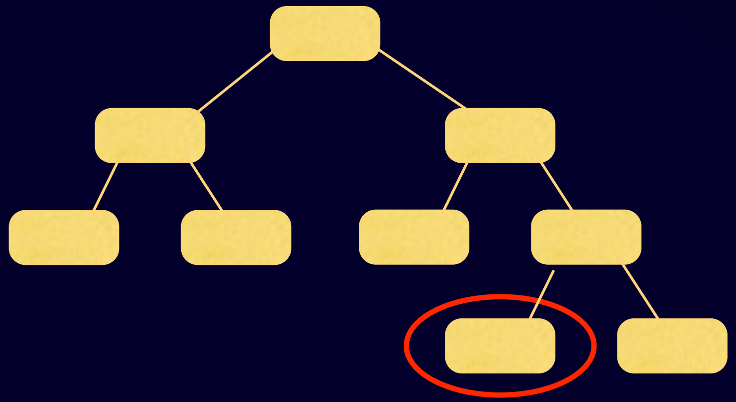

- Estimates

Qualitative? High

100%

- Data/Rationale
  - Assumes that any compromise of customer information would require notification and other secondary costs

Secondary Loss Magnitude

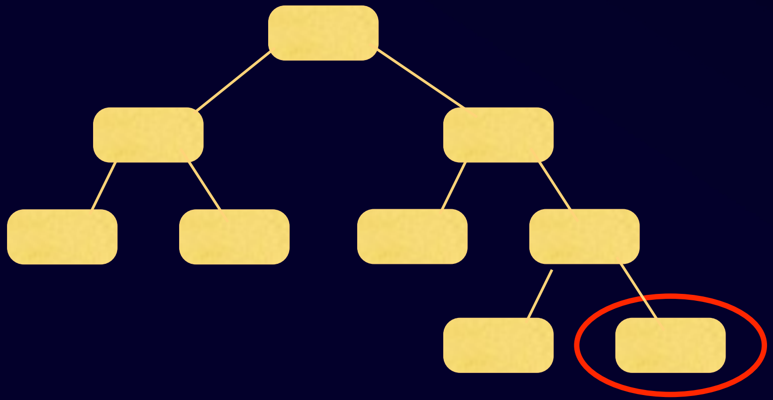

- Definition

The probable loss magnitude resulting from fallout

- Estimates

Qualitative?	Moderate
Min:	$ 100
Max:	$ 500k
ML:	$ 17k

- Data/Rationale
  - Minimum of 1 customer record
  - Most Likely 20 customer records
  - Max 100 records (only accessible one at a time)
  - Includes notification costs, credit monitoring, legal defense, and customer churn

Results

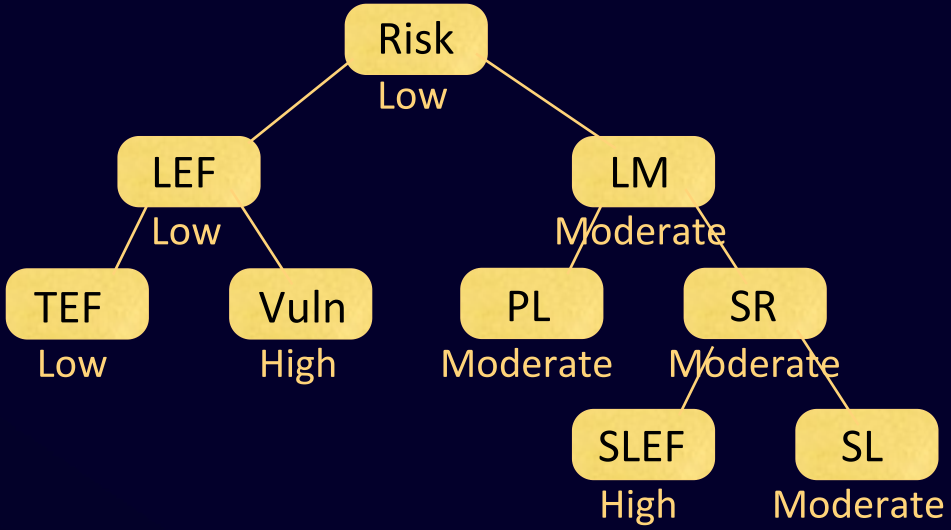

Risk
Low

LEF
Low

LM
Moderate

TEF
Low

Vuln
High

PL
Moderate

SR
Moderate

SLEF
High

SL
Moderate

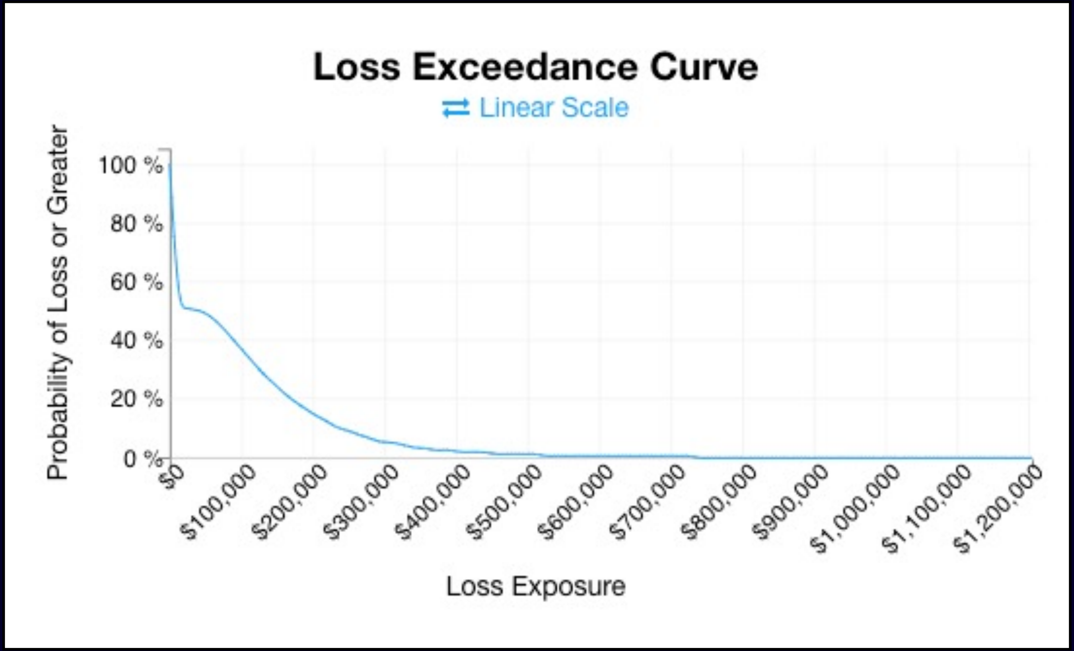

Loss Exceedance Curve

Scoping the 2nd analysis…

- What is the asset at risk? Customer information
- Who/what is the threat actor(s)? Cyber criminals
- What type of action Malicious
- What type of event is it (C, I, or A)? Confidentiality
- What is the loss event scenario? The confidentiality of customer data is maliciously compromised by cyber criminals who are able to breach the perimeter.

Threat Event Frequency

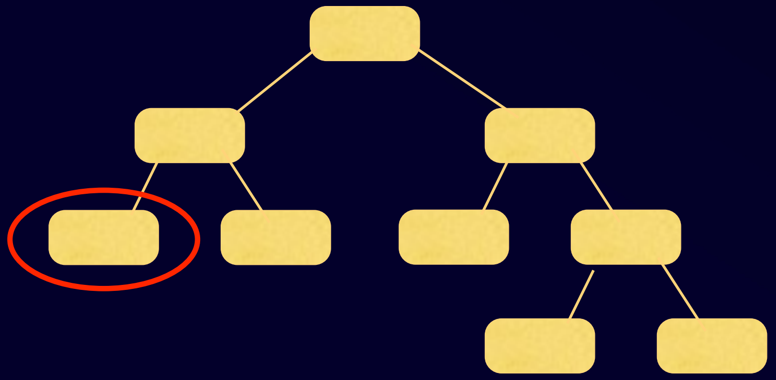

- Definition

The probable frequency, within a given timeframe, that a threat will act in a manner that may result in loss

- Estimates

Min: .1 yr (1 in 10 yr)

Max: 5 yr

ML: .5 yr (every other year)

- Data/Rationale

Based on SME estimates as well as on data from compromised systems at the perimeter that have evidence of attempts to move deeper and laterally within the network.

Vulnerability

- Definition

The probability that a threat event will become a loss event

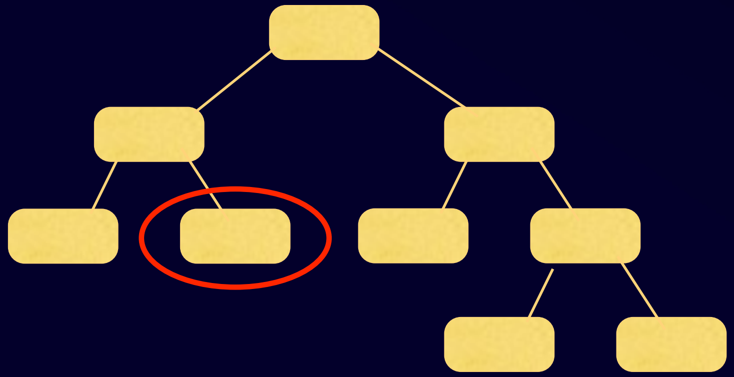

- Estimates

Min: 75%

Max: 99%

ML: 95%

- Data/Rationale

Breaching the perimeter typically involves gaining (or positions the threat actor to gain) access to legitimate accounts, which makes it much more likely that internal resistive controls will be ineffective.

Primary Loss Magnitude

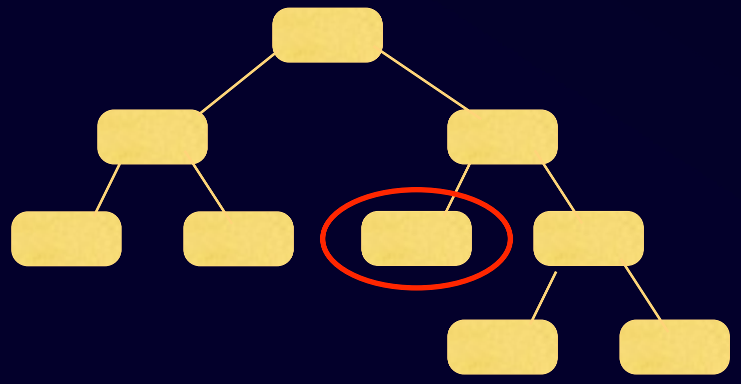

- Definition

  Loss that occurs directly as a result of the threat act against the asset.

- Estimates

  Min: $ 50k

  Max: $ 500k

  ML: $ 100k

- Data/Rationale

  - Internal personnel response efforts
  - Outsourced forensic/investigative costs

Secondary Loss Event Frequency

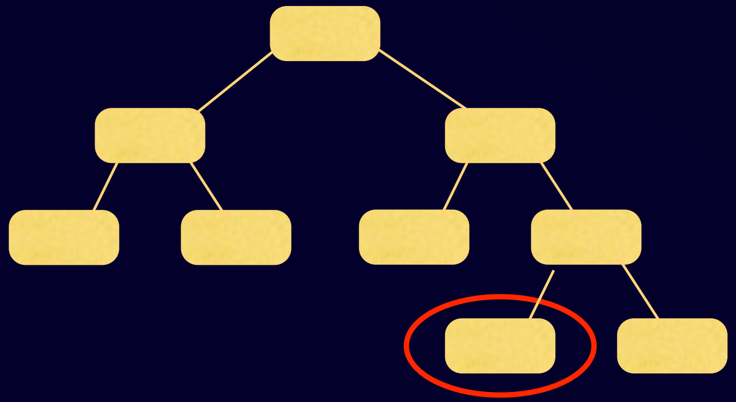

- Definition

  The probability of secondary loss (fallout)

- Estimates

  100%

- Data/Rationale

  Assumes that without early detection the threat actor will eventually compromise some amount of customer information, which would require notification and other secondary costs.

Secondary Loss Magnitude

- Definition

  The probable loss magnitude resulting from fallout

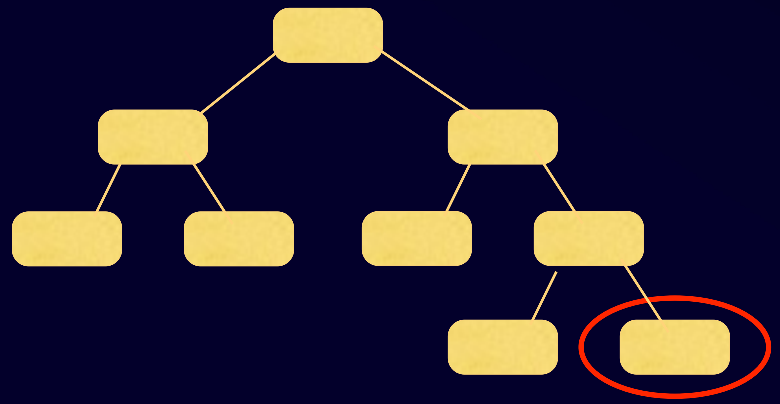

- Estimates

  Min: $ 5k

  Max: $ 100M

  ML: $ 1M

- Data/Rationale
  - Minimum of 1 customer record
  - Most Likely 1M customer records
  - Max all customer records
  - Includes notification costs, credit monitoring, legal defense, and customer churn

Result

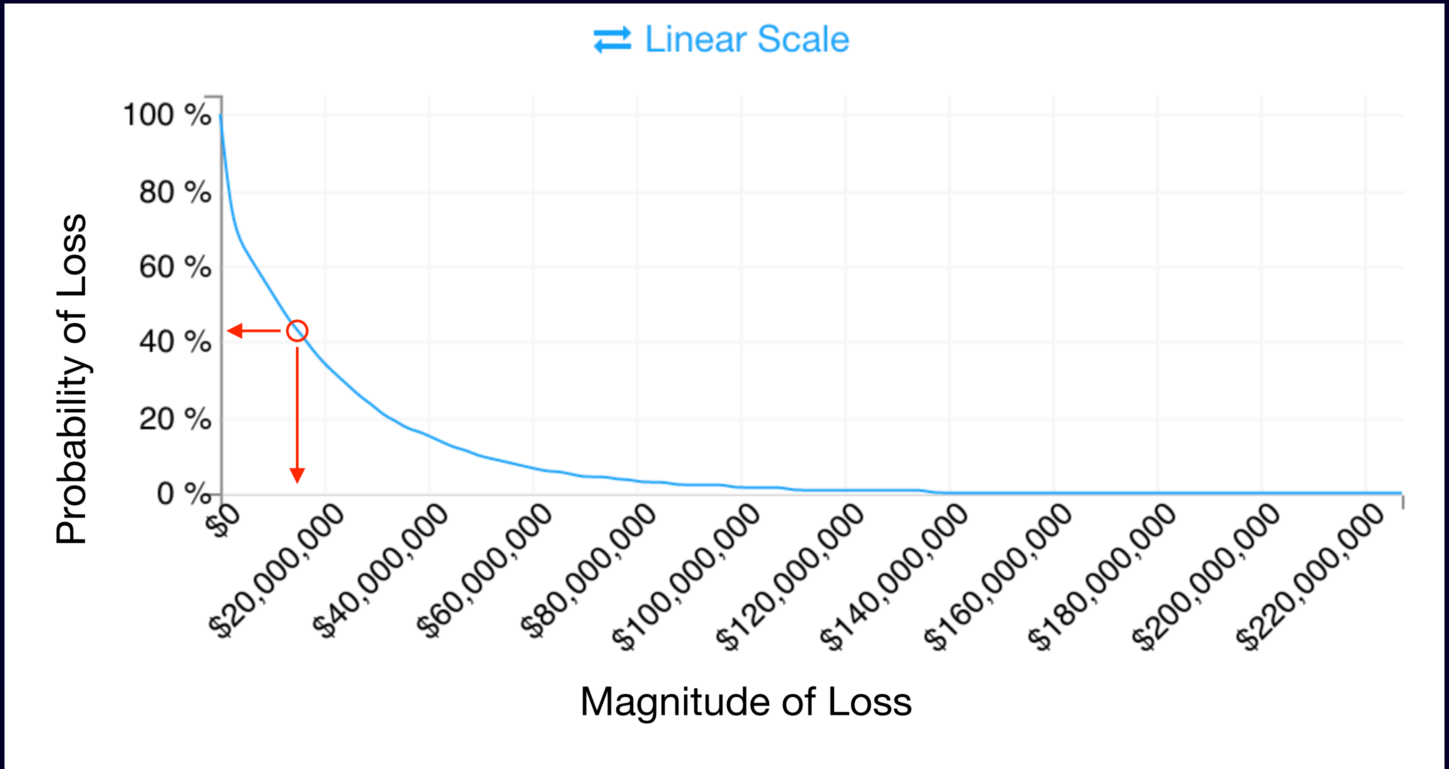
⇄ Linear Scale
Probability of Loss
100 %
80 %
60 %
40 %
20 %
0 %
$0
$20,000,000
$40,000,000
$60,000,000
$80,000,000
$100,000,000
$120,000,000
$140,000,000
$160,000,000
$180,000,000
$200,000,000
$220,000,000
Magnitude of Loss

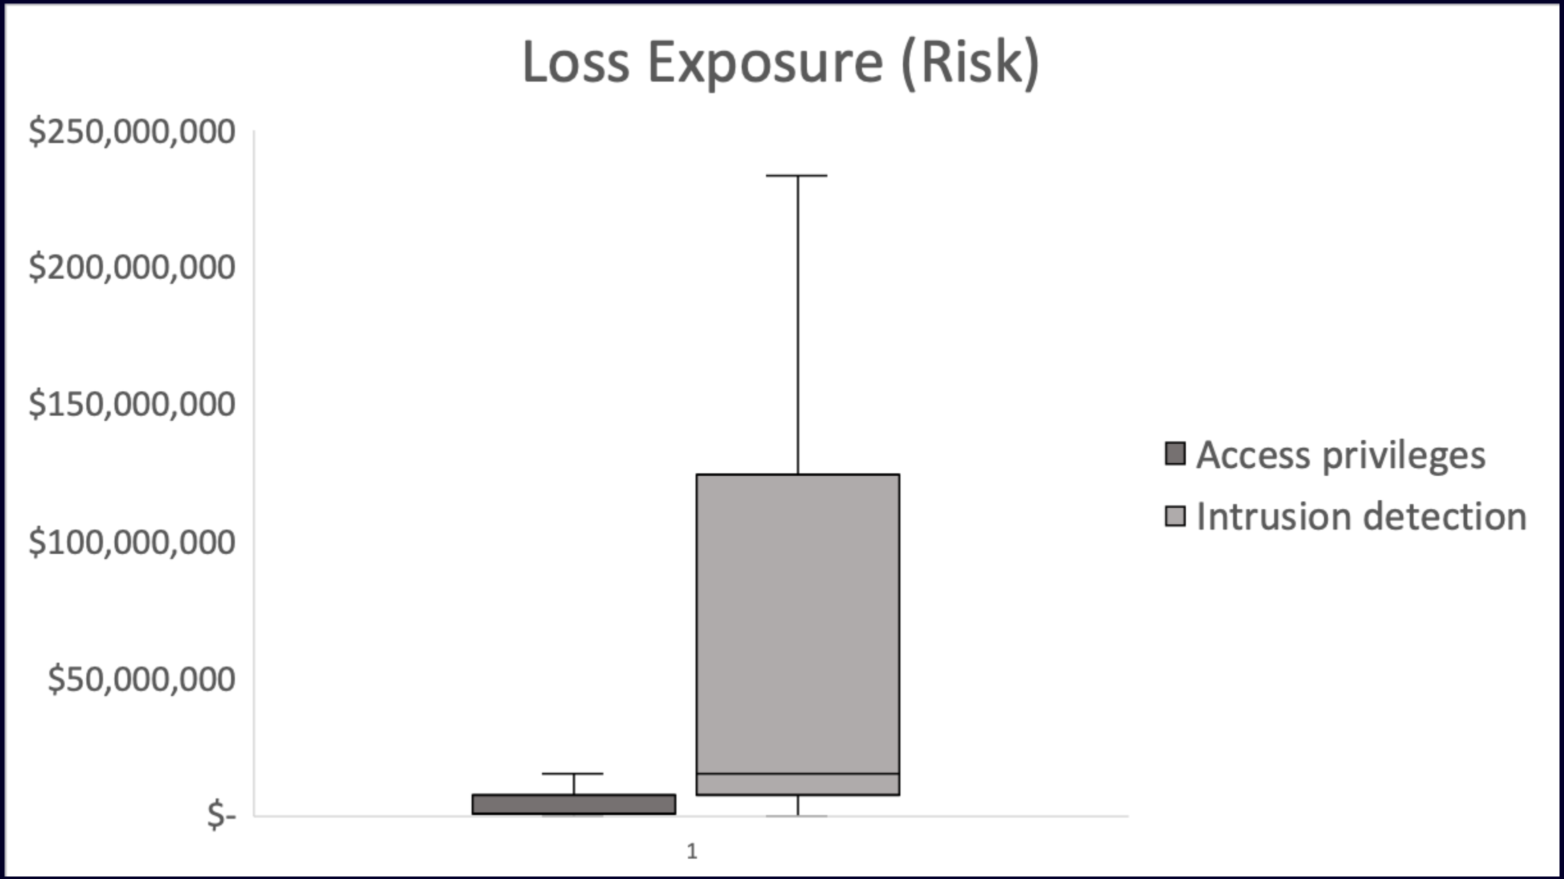
Loss Exposure (Risk)
$250,000,000
$200,000,000
$150,000,000
$100,000,000
$50,000,000
$-
1
Access privileges
Intrusion detection

Common concerns

The most common concerns (besides data)

- How hard is it to measure risk this way?
  - Requires training (to break bad habits)
  - Requires decent critical thinking skills
  - Can be done in Excel, R, etc., or commercial tools
    - Free online learning application: https://app.fairu.net/

- Challenges with adopting these practices?
  - It represents a paradigm shift
  - It increases the cost of risk measurement
  - You need to know what problem(s) you’re trying to solve with it
  - You need to avoid the pursuit of “perfect” data

The siren song of automation…

- The good news…
  - Better loss magnitude data is becoming available
- The not so good news…
  - Threat data is very easy to misinterpret/misapply
  - Controls data is a mess

Controls-related data challenges

- Context sensitivity (e.g., applying weighted values)
- Control relationships and dependencies
- Framework assessment scores (e.g., NIST CSF)
    - Ambiguous definitions
    - Insufficient granularity
    - Undefined/poorly defined scoring scales

Key take-aways…

Automated cyber risk measurement is incredibly easy to screw up.

When it's screwed up, all you've done is automate poor decision-making.

Wrapping up

The value of risk measurement best practices

- Enables accurate risk measurement in economic terms
  - Significantly improves prioritization
  - Supports cost-benefit analysis of security efforts
  - Economic expression of risk is familiar to many executives
  - Enables comparing risk vs. other economically measured organization imperatives (revenue, cost, etc.)
- Also…
  - Surfaces assumptions so they can be recognized and challenged
  - Improves risk-related conversations and collaboration

Summary

- We measure risk in order to make decisions and take actions that affect the frequency and magnitude of loss event scenarios.

- Common risk measurement practices today do not enable us to measure risk reliably.

- Risk measurement best practices require:
  - Clarity: You can't reliably measure what you haven't clearly defined
  - An accurate model: Note that all models require assumptions
  - Explicit consideration of data: Data will always have uncertainty. The key is to faithfully account for and communicate uncertainty.

- Those requirements are true for qualitative or quantitative risk measurements.

FAIR Resources

- The FAIR Institute (www.fairinstitute.org)
- The Open Group (www.opengroup.org/certifications/openfair)
- RiskLens (www.risklens.com/resources)
- Measuring and Managing Information Risk: A FAIR Approach (www.amazon.com)