Factor Analysis of Information Risk (FAIR) Model

Standard Artifact | Version 3.0

January 15, 2025

Table of Contents

1 | Introduction

1.1 | Overview of the FAIR Model

The FAIR Model™ is a tool for analyzing the factors that drive risk, aiding in understanding, measuring, and communicating risk. Based on a first-principles approach, the FAIR Model decomposes risk into its fundamental components, enabling better analytic focus, data application, and quantitative risk measurement and management. This offers a clearer pathway to understanding and cost-effectively managing exposure to loss.

In addition to being an analytical tool for measuring and managing risk, the FAIR Model can be used as a framework for critically thinking about risk or simply to provide a set of clearly defined terms when communicating about risk. Although the FAIR Model is mostly used to quantify risk in financial terms, it can also be used to measure risk in non-financial terms—e.g., mission failure probabilities and effects, or even to measure risk qualitatively.

1.2 | This Document's Purpose

This document is designed to be a concise and authoritative reference of the FAIR Model structure for anyone wishing to use it to manage risk. It includes definitions and descriptions of the factors within the model and their relationships to one another. It also briefly addresses some common misperceptions about the FAIR Model. Note that this document does not cover extensions to the FAIR Model, such as FAIR-CAM or FAIR-MAM. Those standards are described in other documents.

Note that this document is not a "how-to" guide for using the FAIR Model, as that information is available in other FAIR Institute publications and the book Measuring and Managing Information Risk, A FAIR Approach#. The FAIR Institute’s annual conference (FAIR Conference) provides many best practices and case studies for using FAIR to manage enterprise cyber risk.

2 | The FAIR Model

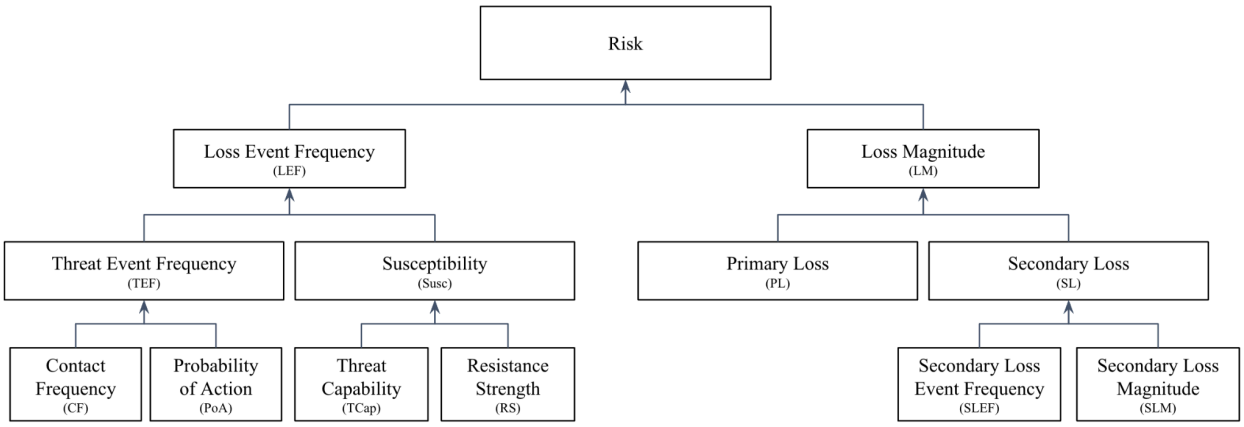

2.1 | Risk

"The probable frequency and probable magnitude of future loss"

This definition emphasizes the fact that measuring risk is inherently a forward-looking problem, which means that it must be measured probabilistically. This, in turn, means that there will always be uncertainty in the measurements, which should be accounted for in both the inputs to the model and its outputs[1].

Note that this is a measurement of loss event scenario frequency and magnitude, which must be clearly defined. For example, *"The risk associated with compromise of PII by cybercriminals via an email phishing attack."*

Although FAIR analysis results often express risk as an annualized value, it is perfectly acceptable and sometimes preferable to express analysis results in different ways. For example, it may be more useful to express analysis results in terms of the probability that losses of a particular magnitude are likely to occur. For example, an analysis might conclude that in the next 12 months there is a 50% probability that a loss greater than $100,000 will occur.

Risk is derived from:

- Loss Event Frequency (LEF)
- Loss Magnitude (LM)

[1] Well-established methods for representing measurement uncertainty include the use of ranges or distributions.

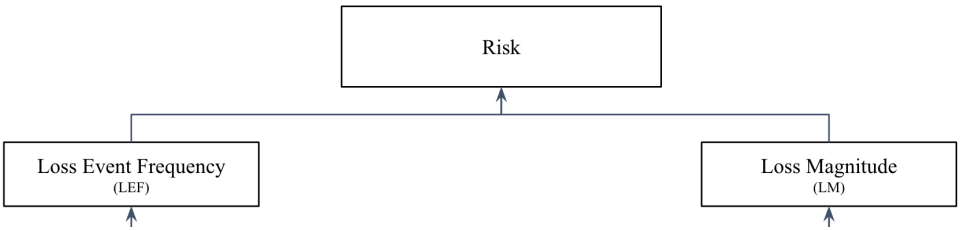

2.2 | Loss Event Frequency (LEF)

"The probable frequency, within a given timeframe,
that a threat agent's actions will inflict harm upon an asset."

Note that the specific type of harm must be clearly defined in the scenario scope (e.g., data theft, outage, loss of data integrity, destruction of facilities, etc.).

Although this factor is labeled as a "frequency," it can also be expressed as a probability of loss occurring within a given timeframe. For example, rather than state that a loss event scenario has an LEF of 1 in 20 years, it's perfectly acceptable to describe this as a 5% probability of occurring in the next 12 months.

LEF is derived from:

- Threat Event Frequency (TEF)
- Susceptibility (Susc)

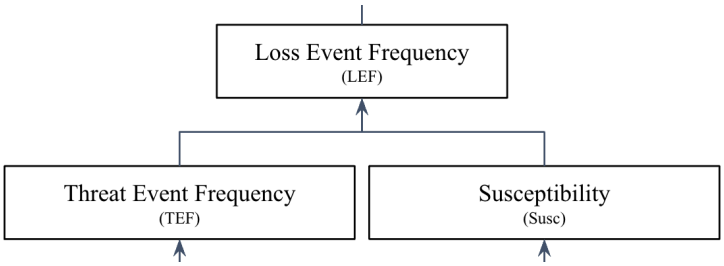

2.2.1 | Threat Event Frequency (TEF)

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

Note that "threat agents" may be human, animal, technological, or natural, and their actions may arise intentionally (either maliciously or non-maliciously), accidentally, or through natural

processes. Another way to think about it is that threat agents are any actor whose action may result in adverse effects.

TEF is derived from:

- Contact Frequency (CF)
- Probability of Action (PoA)

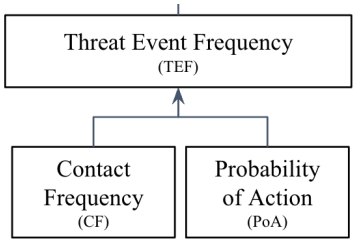

2.2.1.1 | Contact Frequency (CF)

"The frequency within a given timeframe in which a threat agent will come into contact with an asset."

Being in "contact" with an asset means that a threat agent is in a position to act against the asset (physically or logically), given their capabilities and resources. For example, a cybercriminal is in "contact" with a web application on the Internet if they are aware of the web application's existence and have unimpeded network access to it.

2.2.1.2 | Probability of Action (PoA)

"The probability that a threat agent will act against an asset once contact has occurred."

Just because contact occurs between a threat agent and an asset does not guarantee that a harmful action will occur. The probability will vary by the type of threat agent. For example, when Mother Nature is the threat agent, the action is almost certain. However, when contact between a cognitive and rational[2] threat agent (e.g., a human or other animal) and an asset occurs, the probability that a threat event will occur is uncertain.

Probability of Action (for rational threat actors) is derived from the following:

- The threat agent's perceived value from acting
- The threat agent's perceived level of effort in acting

[2] Not all cognitive threat agents will be rational.

- The threat agent's perceived level of risk from acting

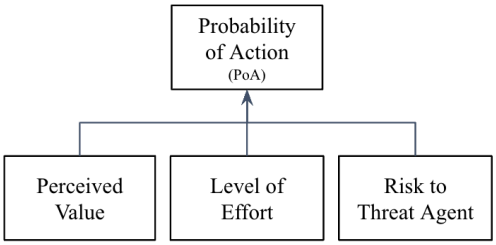

2.2.2 | Susceptibility (Susc)

"The probability that a threat event will become a loss event"

When a threat event occurs, there is some probability that harm will result. This is referred to as "Susceptibility". The degree of Susceptibility (i.e., the probability of a harmful outcome) depends upon whether the threat agent's capabilities are greater than the resistive controls that are in place.

Susceptibility is derived from:

- Threat Capability (TCap)
- Resistance Strength (Susc)

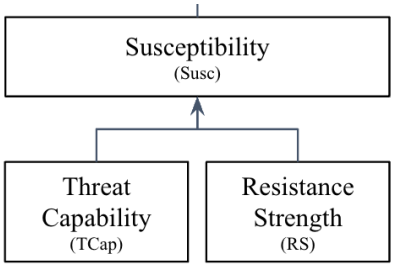

2.2.2.1 | Threat Capability (TCap)

Threat Capability is a measurement of a threat agent's ability to defeat resistive controls. It is an abstract measurement (rather than an empirical one) that takes into account the skills and resources a threat agent is able to bring to bear. It is estimated as a percentile relative to the overall threat community. For example, the overall threat community capability ranges from the 1st to the 100th percentile – i.e., the capability of all threat agents fall somewhere within that range. The capability of a specific threat community will fall somewhere within that range (e.g., 50th to 75th percentile).

2.2.2.2 | Resistance Strength (RS)

Similar to Threat Capability, Resistance Strength is an abstract measurement that represents the efficacy of resistive controls against the overall threat population. For example, a resistive control whose efficacy is estimated to be between 70% and 90% is expected to successfully repel attackers who fall below this range. Any attack whose capability is above 90% is expected to always be successful.

2.3 | Loss Magnitude (LM)

LM is the probable magnitude of loss resulting from a loss event. These losses may result directly or indirectly from an event. This LM is divided into two primary forms: Primary Loss (PL) and Secondary Loss (SL).

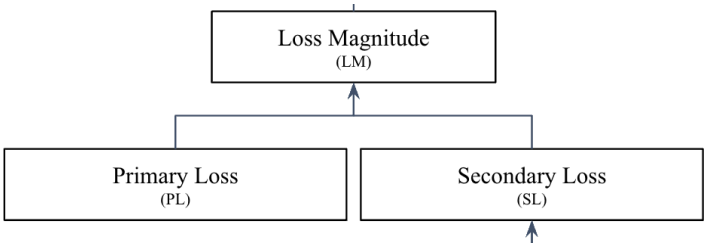

2.3.1 | Primary Loss (PL)

Primary Loss occurs directly as a result of the loss event. These losses are more certain to materialize when a loss event occurs, but their magnitude is often lower than Secondary losses.

2.3.2 | Secondary Loss (SL)

Secondary Loss occurs indirectly from a loss event, typically due to actions or reactions by secondary stakeholders who have been harmed from or are motivated by an event. Examples include, but aren't limited to:

- Customers
- Investors
- Employees
- The community

- Business partners
- Regulators

The indirect nature of secondary losses is often referred to as “fallout” from an event. An important distinction regarding secondary losses is that the probability of their occurrence is less than 100% *because they are indirect*.

Secondary Loss is derived from:

- Secondary Loss Event Frequency
- Secondary Loss Magnitude

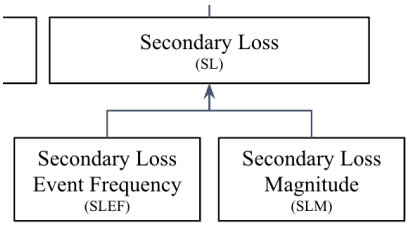

2.3.2.1 | Secondary Loss Event Frequency (SLEF)

Secondary Loss Event Frequency represents the number of Loss Events that are expected to have secondary effects (fallout). Note that although labeled as a “frequency,” it is measured as a percentage. In other words, a loss event scenario might have a SLEF of 80% – i.e., eighty percent of loss events are expected to experience secondary loss.

2.3.2.1 | Secondary Loss Magnitude (SLM)

Secondary Loss Magnitude represents how much loss is expected to materialize from secondary stakeholder reactions. The manner in which these losses materialize is captured in the forms of loss that are described in the next section.

3 | Forms of Loss in FAIR

FAIR categorizes loss into six main forms, providing a reasonably comprehensive understanding of the potential impacts of risk events[3]. Some of these forms of loss are more likely to occur as Primary or Secondary loss, but they can occur in either depending on the specifics of a scenario.

3.1 | Productivity Loss (ProdL)

Productivity loss represents the reduction in organizational efficiency and effectiveness due to the loss event. This includes reduced revenue due to system downtime, decreased employee performance, and interruption of business processes.

3.2 | Response Costs (RespC)

Response costs encompass the expenses incurred while addressing and mitigating the effects of a loss event. These costs cover activities such as incident response, forensic investigations, communication with stakeholders, and legal consultations.

3.3 | Replacement Costs (ReplC)

Replacement costs are the expenditures related to restoring or replacing lost or damaged assets. This includes purchasing new equipment, software, or data recovery services.

3.4 | Fines and Judgments (FinJu)

Fines and judgments refer to the financial penalties and legal settlements resulting from regulatory violations, criminal penalties, or lawsuits related to the loss event. These can include fines imposed by regulatory bodies or damages awarded in civil litigation.

3.5 Reputation Damage (RepuD)

Reputation damage represents the potential long-term impact on the organization's brand and customer trust. This form of loss can lead to decreased customer loyalty, reduced market share, and lower revenue.

[3] The FAIR-MAM (Materiality Assessment Model) provides increased granularity and clarity in loss magnitude analysis and measurement.

3.6 | Competitive Advantage Loss (CAdvL)

Competitive advantage loss occurs when the loss event undermines the organization's market position or strategic initiatives. This can result from the loss of intellectual property, diminished innovation capacity, or weakened market perception.

4 | Misperceptions

This section addresses a few common misperceptions about FAIR that have developed over the years.

4.1 | Model Depth

When using FAIR, it is not necessary to operate at the deepest layer in the model. It's perfectly acceptable to make estimates at, for example, the TEF layer rather than the deeper Contact Frequency and Probability of Action layer in the model. Time constraints and availability of data will typically influence which layer of the model an analyst will choose to work from.

4.2 | Data Sources

There is a misperception that using subject matter expert (SME) estimates is required when using FAIR. In fact, nothing in the FAIR standard stipulates the use of specific data sources. If an analyst has at their disposal security telemetry, industry data, etc., they are free to use that data. Many businesses employ cyber risk management software solutions that ingest data to automate significant aspects of FAIR analysis, and there are third parties.

4.3 | Distributions

There is a misperception that FAIR mandates using PERT[4] (or Beta PERT) distributions when performing risk assessments. In fact, there are no stipulations regarding the use of any specific statistical methods or distributions. Analysts and technologies that use FAIR are expected to leverage methods and distributions that best fit their needs and constraints.

[4] Developed in 1958, program evaluation and review technique (PERT) is a statistical method of analyzing the tasks involved in completing a project.