ψ

Alan Bell, PhD
Hamid Shirazi
July 2026

1 ABSTRACT

In mathematics, there are many constants that emerge as useful tools for simplifying complex relationships and associated calculations. Arguably the most well-known constant is referred to by the Greek letter π, an indispensable number used in geometry. The Greek letter ψ (psi, pronounced like "sigh") represents a new constant presented in this paper for use in aviation safety risk management. This paper briefly reviews several background items that are important to the derivation of the constant, then demonstrates its applicability to risk analysis. The objective of this paper is to present ψ as a simple tool that can be applied by any risk analyst, regardless of background, to support risk quantification in safety analyses where hazards exist but have not yet caused an incident or aircraft accident. For completeness, the paper also includes a brief description of the mathematical underpinnings and limitations for those who desire a more formal mathematical explanation of the concept.

2 BACKGROUND

The concept of ψ emerges out of the interaction between several other concepts. Each relevant concept is briefly reviewed in this section to build a foundation for explaining the derivation of ψ.

2.1 THE BINOMIAL DISTRIBUTION

The Binomial distribution is a useful tool for understanding the nature of probability associated with discrete events such as a series of coin flips. If the probability of "heads" on a coin flip is 50%, the binomial distribution allows us to ask a question such as, "if the coin is tossed 20 times, what is the chance of the result being *heads* on seven of the twenty flips?" The answer provided by the binomial distribution is 7.4%. It will also tell us the probability of any possible outcome, from 0 to 20.

2.2 THE RISK FRAMEWORK

In Decision Science, decision trees are frequently used to quantify the risk associated with two or more alternatives, allowing a risk-based decision that governs consistent selection of the alternative that is expected to deliver the highest value (or utility). Each section of the tree associated with an alternative is referred to as a risk framework, allowing the risk associated with that alternative to be quantified, then compared to the risk of other alternatives. In the context of aviation safety risk management, a risk framework is a tool that organizes data and enables basic calculations to determine the *likelihood* of a hazard causing or contributing to an incident or aircraft accident, and that *occurrence* reaching some level of severity. A generalized version of the risk framework for hazard analysis is shown in Figure 1.

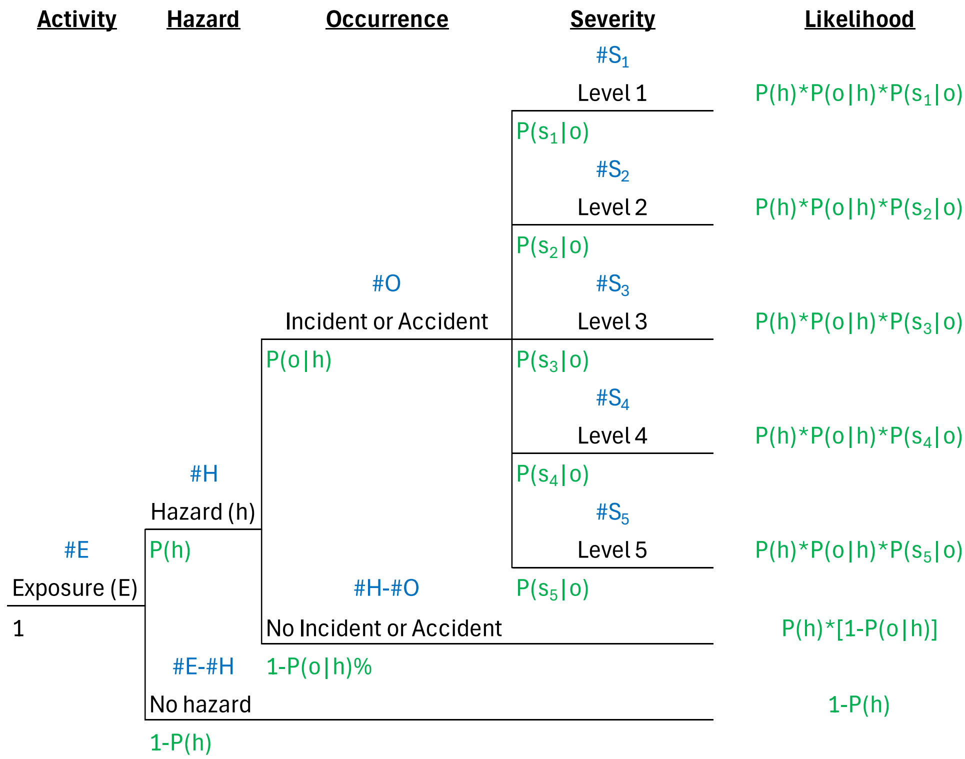

Figure 1: Generalized Risk Framework for Hazard Analysis

2.3 Expected Value

Risk is comprised of two elements: an uncertainty term such as *probability* or *likelihood*, and an outcome term such as *severity* or *objectives*. A risk framework relies on probabilities shown in green beneath each branch and at the end of the tree. The numbers above each branch in blue text are expected values. Within the risk framework, an expected value is found by multiplying the probability by the number of associated events. For example, if 15 million instances of a hazard are multiplied by the probability of an *occurrence*, such as 0.03%, the result is about 5,000 occurrences. Because of this simple relationship, if the number of occurrences is known, the probability of the occurrence can be estimated by simply dividing the number of occurrences (5,000) by the exposure (15 million). In general, for the discrete case described in this paper, the expected value (EV) equals the number of "events" (n) multiplied by the probability of the event (p), or in equation format, EV = n * p.

2.4 Confidence

The coin flip example is useful for explaining the concept of confidence in statistics. Using the binomial distribution, the probability of each possible outcome can be summarized in a bar graph. Then, the bar graph can be used to illustrate the concept of statistical confidence. For example, if the probability of observing 9, 10, or 11 *heads* are summed together, the result is approximately 50%. Thus, if a coin is tossed 20 times, there's about a 50% confidence level that the resulting number of heads will be either 9, 10, or 11. These potential outcomes are contained by green bars in Figure 2.

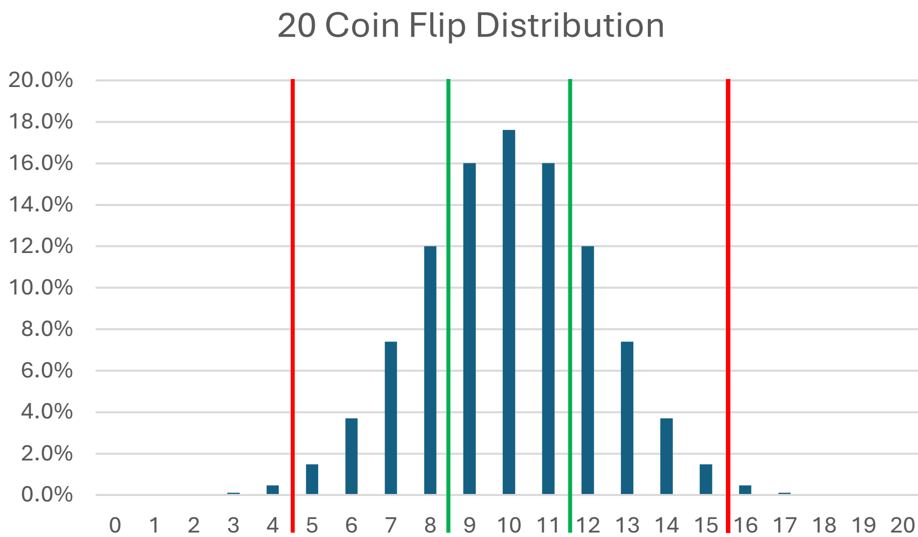

Figure 2: Confidence Intervals for 20 Coin Flips

In general, the wider the range between the colored bars shown in Figure 2, the more confidence there is that if a coin is thrown 20 times, the result will fall between the bars. When stating a confidence interval (or a margin of error), the amount of confidence associated with the result is discretionary. That is, the analyst decides what level of confidence is needed or desired, and then reports the results associated with that decision. In the case of twenty coin flips, if a 99% confidence is desired, then the confidence interval is between 5 and 15 as shown by the red bars in Figure 2.

3 Discussion

The coin flip examples thus far rely on the knowledge that the probability of heads is 50%. What if that probability isn't known? Suppose a coin is flipped twenty times and the result is 14 heads. What can be said about the probability of heads based solely on this observation? An answer to that question is provided in a video presentation, Probability of Probabilities Part 3, and while this topic builds a foundation for the development of ψ, a complete answer is beyond the scope of this paper.

A more specific question is, what if the coin is believed to have a 50% chance of heads, but when it is tossed 20 times, the result is 0 occurrences of heads? If the coin had a 50% chance of heads, this outcome is possible, but the chance of it happening is so small (less than one in a million), it would be very hard to believe. Observing 20 consecutive outcomes of *tails* would lead to an updated belief regarding the true probability of heads. The more times the coin tosses are repeated with a result of 0 heads, the greater the magnitude of the change in belief regarding the true probability of heads. The coin toss example is a simplified demonstration of Bayesian probability, and how a prior belief in a probability should be updated based on the results of an experiment or observation.

3.1 Credibility

The concept of *credibility* is closely related to the concept of *confidence*. One of the challenges in explaining the concept of credibility is that the term *credible* has been widely used as a qualitative metric in aviation risk management. This paper offers a quantitative explanation of *credibility* by deriving a definition from probability theory, and more specifically, from the binomial distribution.

Like confidence, the threshold of credibility is subjective, and it is established by the analyst. To illustrate the concept, if a coin flipped 20 times and the result is zero heads, what would the chance of heads have to be for it to be considered a credible outcome? Intuitively, it's relatively easy to see that the chance of heads must be lower than the 50-50 chance expected with a typical coin, but how much lower does it need to be for a result of 0 heads to be credible?

The binomial distribution provides an answer to an important question that is directly related: "if a coin with some chance of heads is tossed 20 times, what is the probability of observing zero heads?" For instance, if the true chance of heads is 10%, then the answer is there's about a 12% chance of seeing that outcome. If the true chance of heads is 5%, then the answer increases to about 36%. This means that if a coin is flipped 20 times, the result of zero heads is more credible for a coin having a 5% chance of heads than one with a 10% chance. The more the chance of heads is lowered, the more credible a result of zero heads becomes. In the context of aviation risk management, *credibility* is a metric that quantifies the chance of seeing zero outcomes over some period of exposure given the probability of that outcome.

3.2 Zero Occurrence Observations

A frequent challenge in aviation safety risk management is trying to understand what the probability of an *occurrence* could be if a hazard of interest has not been cited as causal or contributing to an incident or accident. A variation of this problem is that as the airspace system is improved, the risk associated with hazards that caused serious accidents in the past may have been reduced to a level where, in the current system, there have been no additional incidents or accidents.

To illustrate this challenge, consider a notional example. Over a 15-year period, there is an exposure of 1M flight hours per year, or 15M total flight hours. During this period, a hazard, "Hazard X," is present 5,000 times but causes 0 incidents or aircraft accidents. If 0 is entered into a risk framework for the number of occurrences, the result is that the likelihood of effects at all levels of severity becomes zero as well. This relationship is illustrated in the risk framework shown in Figure 3.

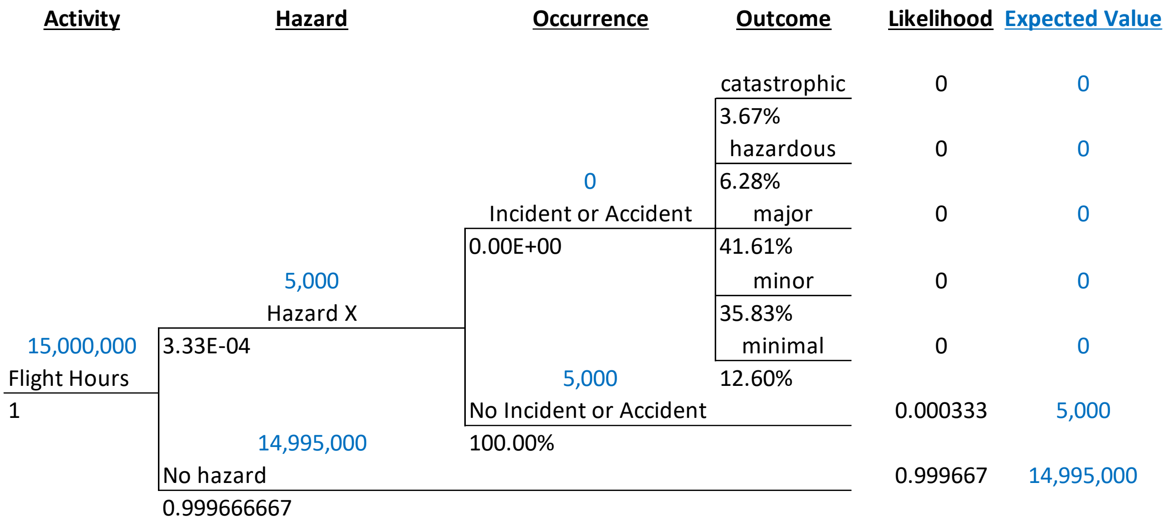

Figure 3: Risk Framework with 0 Occurrences[1]

4 Estimating Probability Through Iteration

In aviation, and especially in commercial aviation, the probability of a hazard causing an incident or an aircraft accident is typically very low. As a result, it is common for safety data to reflect that there are no incidents or accidents caused by the hazard, even in relatively large safety databases that cover many years. However, the lack of an accident in the database does not imply that the probability of an accident is zero. The capability that ψ unlocks is to be able to transform raw data that includes many flight hours, flights, or operations without incident or accident into a *credible* probability. That capability is described and illustrated in the following sections.

4.1 Standardizing Credibility

Before iterative techniques can be discussed, a standard "rule of thumb" must be established for credibility. While there is no rule that requires its use, a 95% confidence level is frequently applied in statistical applications. Similarly, there is no rule that would govern credibility levels, and analysts are free to increase or decrease credibility on a case-by-case basis. However, a credibility level of 50% provides a solid starting point for any analysis.

Returning to the coin flip example, if the true probability of *heads* is 3.4%, then the binomial distribution can be used to show the chance of observing zero heads in 20 flips is about 50%. Stated another way, if a coin with a 3.4% chance of heads is flipped 20 times, there is about a 50-50 chance of the result being zero heads, so that result would be credible. Table 1 provides a more complete summary of the binomial distribution results and provides a graphic depiction of the term *credibility*.

[1] The risk frameworks presented in this paper include a distribution for potential outcomes at each severity level if the hazard does cause an occurrence. These distributions are derived from historic data, but a complete explanation of applicable techniques for developing a severity distribution is not relevant to the concept of ψ and is therefore beyond the scope of this paper.

Table 1: Binomial Distribution for a 3.4% chance of Heads

# of Heads	Probability
0	50%
1	35%
2	12%
3	2%

Credibility

Another way of thinking about this is to consider the Bayesian implications. In the previous example, where the prior belief that the coin was "fair," 20 consecutive coin flips with a result of tails would suggest a significant update is needed in the prior probability, leading to a posterior probability that is significantly lower. However, if the prior probability of heads is 3.4%, a result of 0 heads in twenty coin flips would be expected, on average, about every other time the experiment is conducted. Therefore, there would be virtually no update to the prior probability as a 0-heads outcome is clearly a credible result. The constant ψ directly depends on the level of credibility established by the analyst. Therefore, for the purpose of aviation safety risk analysis, the level of credibility proposed as a standard is 50%.

4.2 An Iterative Approach to Probability Estimation

The risk framework presented in Figure 3 includes a 0% conditional probability of an occurrence given the hazard is present. To eliminate the implication that there is zero risk, the zero probability and its associated expected value must be replaced with a positive number. One way this can be done is by iteratively adjusting the probability until there is a 50% chance of observing 0 incidents or accidents over an exposure period of 15M flight hours, and in which the hazard is present 5,000 times.

A good starting point for this process is to start with a hypothesis of "1 in X," where X is the number of hazards that are present within the exposure period. So, in this case, an initial hypothesis would be 1 in 5K. If this is done, the binomial distribution reveals that the chance of observing 0 occurrences is only 37%. Therefore, this hypothesis would be rejected as it does not achieve the 50% credibility standard. The probability of an occurrence must be lower than 1 in 5K in order for a zero occurrence outcome to be considered credible.

A new hypothesis of a lower probability could be established at 1 in 10K. Once again, the binomial distribution can be used to determine the probability of 0 occurrences, with the result being about 61%. This number is well above the threshold for credibility. As a result of these two hypothesis tests, it's clear that the probability resulting in a 50% chance of 0 occurrences over an exposure period with 5K hazards lies somewhere between 1 in 5K and 1 in 10K. By iteratively changing the hypothesis, an analyst could eventually determine the probability of occurrence that creates a credible chance of observing 0 occurrences in response to the 5K hazards. By selecting a 1 in 7K chance, the binomial distribution returns a 49% chance of observing zero occurrences, and for the purpose of illustration, this result is close enough.

4.3 An Iterative Approach to Expected Value Estimation

In addition to a conditional probability of occurrence, the risk framework in Figure 3 also includes the corresponding expected value. If the probability of occurrence has not yet been determined, an alternative to iteratively updating a probability hypothesis is to update an expected value hypothesis. For the example described in the previous section, the original hypothesis was 1 in 5K. The "1" in this case is the expected value. The second hypothesis could be restated as 0.5 in 5K, leading to a 61% credibility. Through iteration, 0.6931 in 5M is the expected value that results in a 50% credibility.

5 IDENTIFYING Ψ AS A CONSTANT

In the previous sections, an iterative process led to the identification of either a *probability* or an *expected value*. The same process could be used regardless of the exposure or the number of hazards. For instance, if over 100M flight hours and the hazard was present 44,000 times, a probability of 1.58×10^{-5} would lead to a 50% chance of experiencing 0 occurrences during that period. If we then apply the expected value equation, $EV = n * p$, the expected value is equal to 0.6931. Table 2 provides several additional examples.

Table 2: Expected Number of Occurrences Needed to Achieve 50% Credibility

# of Hazards	P(Occurrence | Hazard)	Expected # of Occurrences
5,000	0.0139%	0.6931
25,000	0.0028%	0.6931
44,000	0.00158%	0.6931
100,000	0.0007%	0.6931
1,000,000	0.00007%	0.6931

This table could be expanded to include any number of hazards. Therefore, $\psi = 0.6931$.

6 APPLICATION

To demonstrate the application of ψ, a real-world example is useful. Air carriers in the United States operating under 14 CFR § 121 logged over 265M flight hours between 2008 and 2022. During this time, *Oil System Malfunctions* were reported to the FAA as hazards 1,440 times by operators. However, this hazard was not cited as causal or contributing to any incidents or aircraft accidents during the same period. A risk framework for this example is presented in Figure 4.

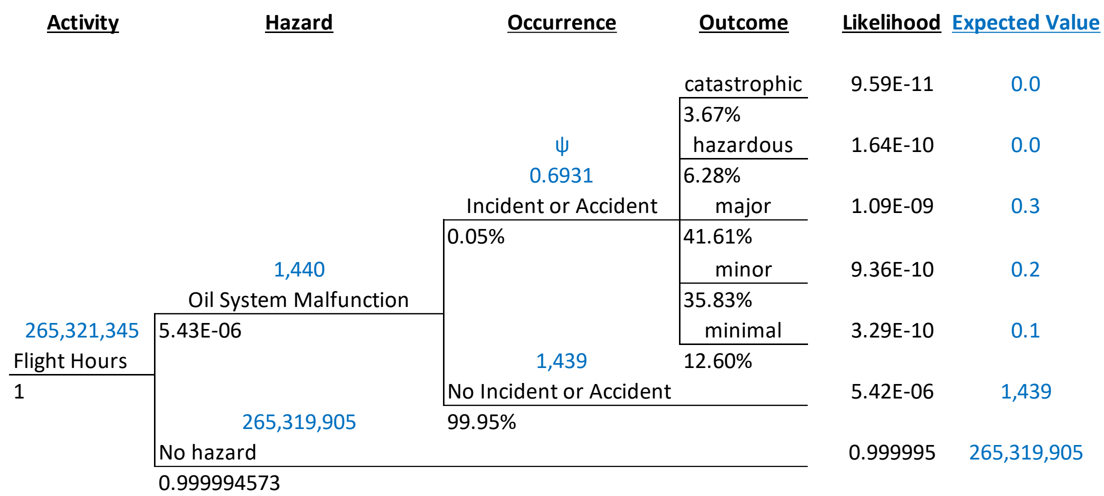

Figure 4: Risk Framework for Part 121 Oil System Malfunctions, 2008-2022

Since no incidents or accidents have been observed, replacing the expected value for the number of occurrences with ψ produces data driven likelihood estimates that are further explained in the following section.

7 INTERPRETING Ψ ANALYSIS RESULTS

The results of a typical risk framework may be used to quantify the risk associated with a hazard such as an oil system malfunction. However, the results of a ψ analysis provide a different perspective on risk. The likelihoods calculated by the risk framework cannot be viewed as a point estimate of risk, and the same holds true for the expected values. Instead, the results of a ψ analysis produce a ceiling on credible risk levels. For example, Figure 4 shows a likelihood of 1.09×10^{-9} for major severity outcomes. This should be interpreted as the highest likelihood that could still result in zero outcomes at that severity level with credibility. Clearly, the actual likelihood of this outcome could be much lower and still produce zero effects at this severity level. Similarly, the expected value should be interpreted as an expectation of 0.3 major severity outcomes *or less* per 265M flight hours. Finally, statistically valid performance targets typically include an error margin on both sides of the expected value, but in the case of a ψ analysis, they would include a one-sided error tolerance on the high end and no lower bound on the low end.

To reflect the results of a ψ analysis, in addition to the traditional point estimate of likelihood in a risk matrix, a tail can be added to each point to reflect that the likelihood could be anything less than the plotted point. However, anything greater than this would not be a credible estimate of risk. Figure 5 provides an example of an FAA Order 8040.4C risk matrix for commercial aviation with the results of the ψ analysis from the previous section.

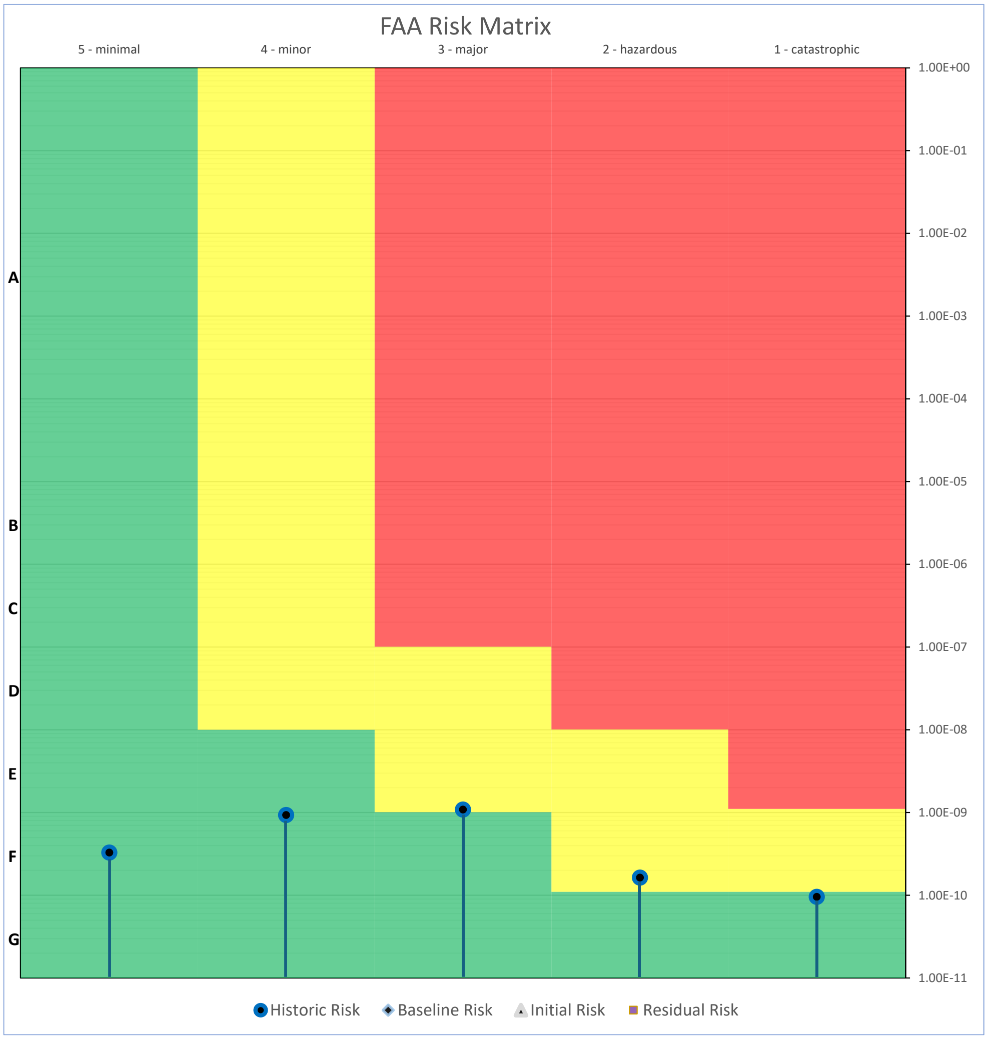

Figure 5: Historic Risk Matrix for Part 121 Oil System Malfunction Hazard, 2008-2022

8 Underlying Mathematics and Limitations

The purpose of developing ψ is to estimate the maximum probability that may have existed even when there have been no observations of the outcome of interest. Equally important is the opportunity to use this tool in risk analysis without the need for deep knowledge of probability or statistics theory. However, this constant has limitations. If an operation is only performed a small number of times, and no undesired outcome occurs, there would be no value to applying ψ because all that could be said is the risk could be extremely high.

To see the limitation of ψ, it's useful to start with the extreme case of a single observation without a negative outcome. If 50% is the threshold for credibility, then the value of ψ is 0.50, well below the 0.6931 value described previously. While an objective of this paper is to ensure readability, regardless of the reader's background in probability theory, the validity of the underlying mathematics is important. To address the mathematics, the derivation of ψ can be expressed via the following equation:

$$f(n) = n * (1 - 0.5^{1/n})$$

For any n between 1 and ∞. With this function, as n approaches infinity, the value of ψ approaches 0.6931.

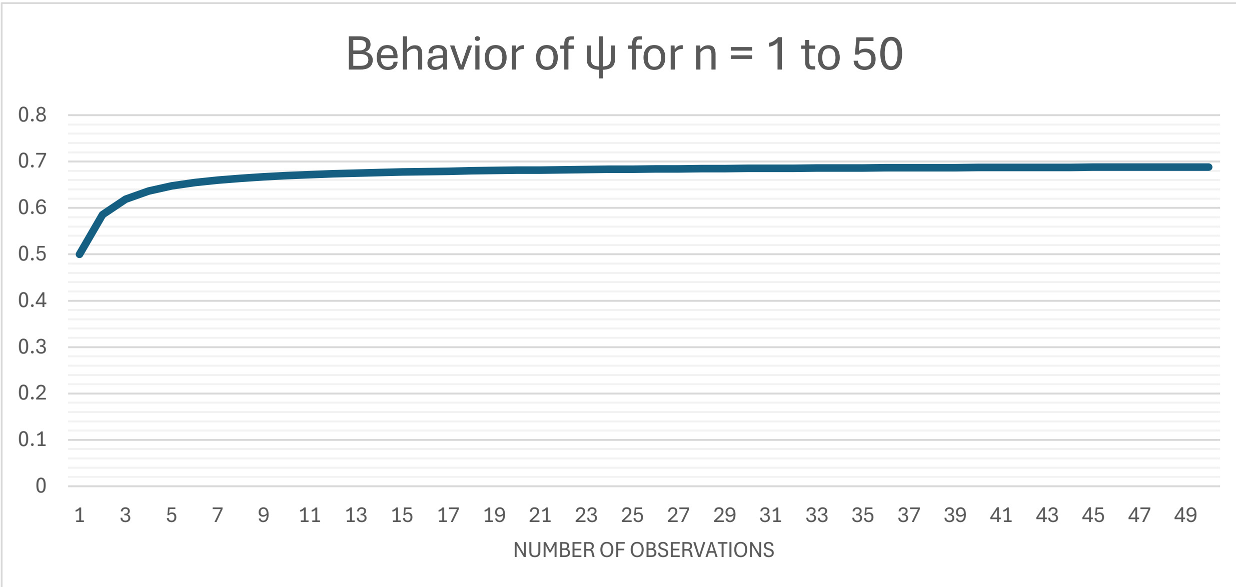

Figure 6: Limit-like Behavior of Psi

Thus, ψ is not really a constant as suggested at the outset, but instead, a limit. However, when used with any large number of observations, the number 0.6931 can be thought of as a constant for the purpose estimating the maximum risk that could have credibly existed based on the number of observations with no outcomes of interest.

9 CONCLUSIONS

Aviation safety risk management frequently involves estimating risk for rare events. The absence of incidents and aircraft accidents over some period of exposure cannot be interpreted as a zero probability of occurrence. However, there is a valuable technique available for estimating the highest risk that could exist and still produce zero occurrences as a credible outcome.

This paper presents much of the theory leading to establishment of ψ. However, there is no reason for risk managers to have a complete understanding of the underlying theory behind its derivation. Instead, all that is needed is a standardized risk framework with ψ entered as the expected value for the number of occurrences. Combined with data or estimates of the exposure, number of hazards present, and a severity distribution, the maximum likelihood of effects that are credible at all severity levels can be established.