MARY KAY O'CONNOR PROCESS SAFETY CENTER
TEXAS A&M ENGINEERING EXPERIMENT STATION

22nd Annual International Symposium
October 22-24, 2019 | College Station, Texas

# Reverend Bayes, Meet Process Safety: Use Bayes' Theorem to establish site specific confidence in your LOPA calculation

Keith Brumbaugh, Dave Grattan*
10375 Richmond Ave. Suite 800
Houston, TX 77042
*Presenter E-mail: dave.grattan@aesolns.com

## Abstract

The Process Industry has an established practice of crediting IPLs (Independent Protection Layers) to meet risk reduction targets as part of LOPA (Layer of Protection Analysis) studies. Often the risk targets are calculated to be on the order of 1E-4 per year or lower. Achieving the risk target on paper is one thing, but what is missing from the LOPA calculation is a statement of the confidence in the result. LOPA is an order-of-magnitude method, however, this only reflects the tolerance of error, not the tolerance of uncertainty. It is often stated that LOPA uses generic credits that are conservative, thereby implying the LOPA result should be conservative. By itself this statement is dubious because the generic data used in LOPA did not originate from the facility for which the statistical inferences are being made (which for frequentist-based statistics makes the inference invalid). Worse, when conservative credits are multiplied together to produce a rare-event number, does the conservative property emerge from the combination?

There is no way to answer this question without performing IPL Validation (i.e., ensuring the IPL will function when needed). However, IPL Validation and related Safety Life-cycle methods (e.g. functional safety assessments and cyber-security audits related to barrier integrity) are purely qualitative and have no apparent relation to the quantitative risk target. There is a need therefore, to bridge the qualitative results of IPL validation with the quantitative result of the associated LOPA calculation, as a way to establish a site-specific confidence level in the risk target we are trying to achieve.

This is where Bayes' Theorem comes in. Bayes' Theorem is an epistemological statement of knowledge, versus a statement of proportions and relative frequencies. It is therefore a method that can bridge qualitative knowledge with the rare-event numbers that are intended to represent that knowledge.

Bayes' Theorem is sorely missing from the toolbox of Process Safety practitioners. This paper will introduce Bayes' Theorem to the reader and discuss the reasons and applications for using

Bayes in Process Safety related to IPLs and LOPA. While intended to be introductory (to not discourage potential users), this paper will describe simple Excel™ based Bayesian calculations that the practitioner can begin to use immediately to address issues such as uncertainty, establishing confidence intervals, properly evaluating LOPA gaps, and incorporating site specific data, all related to IPLs and barriers used to meet LOPA targets.

**Keywords:** Bayes' Theorem, IPL Validation, Uncertainty, Risk Analysis, Tolerance

## Disclaimer



## 1 Introduction

If probability is a measure of uncertainty, then inference is used to make a statement of how accurate that probability is. Statistics is the tool by which we make inferences. Of course, an inference itself is subject to the same kind of analysis concerning its accuracy. And this is how we arrive at different kinds of statistics, namely, Frequentist and Bayesian.

That there exist different kinds of statistics may come as a surprise to many Process Safety practitioners, as most of us have never given much thought to the inferences implied in the numbers we use. We excel as Probability Calculators. This paper is asking you to become a Probability Thinker, which is much more important.

We use many numbers in Process Safety associated with predicting the likelihood of catastrophic events (e.g., failure rates, demand rates, incident rates, probability of failure, probability of ignition, etc.). Very rarely do we think about how good (i.e., trustworthy) the numbers are.

The LOPA calculation presents a unique case in statistical inference. It is neither practical nor ethical to determine rare event frequencies of catastrophic accidents by experiment [1]. Instead, the rare event frequency must be inferred. And it is actually worse than this, because several inferences must be made to arrive at the calculated LOPA number (i.e., each of the individual probabilities of failure are themselves an inference). And it is actually still worse than this, because the data I am using for the inferences are not from my plant. In this paper, I am interested in my own situation, not everyone else's.

This paper takes the position that Bayesian inference is the correct statistical tool to use for making process safety decisions regarding catastrophic rare events in my plant. It is the most consistent and rational method to update current beliefs about safe operation, as new data and evidence trickles in. It is the best we can do in a complex changing operating environment, where we can't afford to wait even 10 years to gather enough data to use Frequentist based methods.

There is one more question to answer before we get started. Does any of this matter? In other words, why do we need a different type of statistics to describe the rare event frequency? There is a simple answer. An inference on the LOPA number for a given rare event is one of the most (arguably *the* most) important Leading Indicator in your plant. On paper, calculated with generic data, the LOPA calculation is "*only the starting point"* (as someone once brilliantly but unwittingly said). Paired with Bayes' methods, the LOPA inference is *the best* plant-specific statement of how safely you are operating with respect to that potential hazard.

## 2 Frequentist vs. Bayesians. Why Process Safety Practitioners should be Bayesians.

Most of us reading this paper have been educated in Frequentist based statistics, learning concepts such as the Law of Large Numbers, Maximum Likelihood Estimate, hypothesis testing, etc. Frequentist based statistics assumes that the relative frequency of an event (i.e., how many times an event occurs over the sample space) is the same as the probability of the event occurring. Unfortunately, this assumption only applies to situations where many identical trials can be repeated (mathematically defined as infinity, but colloquially known as "in the long run"), with well-defined outcomes (think games of chance where the odds are known *in advance*), and where a "true" fixed parameter value can be assumed to exist in the population (e.g., average height of males versus females taking Calculus 101 for Fall 2019, at Texas A&M).

For reasons to be discussed, these concepts are not useful to the Process Safety practitioner attempting to quantify rare event frequencies. Ultimately this affects our ability to make good decisions regarding risk reduction allocation, as well as having confidence that we are truly operating safely. At a more basic level, Frequentist based statistics is the wrong math to use when making inferences about rare events that haven't happened yet in my plant.

First, a rare event is a one-off (i.e., a single-case probability). There is no meaningful interpretation of a rare event occurring many times in my plant "in the long run" (one event is too many). Second, in Process Safety, the odds of the rare event outcome are not known in advance. A common gimmick sometimes seen at trade shows is to roll several 10-sided dice ("LOPA dice") to make an analogy to the probability of all barriers failing (or being failed) at the same time. This is misleading, because in Process Safety we do not know *in advance* what dice we are gambling with! Third, the Frequentist concept of "identical" trials is not valid to Process Safety, because of the complexity of our systems (e.g., human and organizational) that change with time. Related to this is the Frequentist notion of a "true" or fixed parameter value (e.g., probability of failure, initiating event frequency, rare event frequency, etc.) describing the population. These parameters are not fixed (because over time the population changes). For example, there is not a "true" average probability of failure on demand for a safety interlock (because the systematic interactions of these safety interlocks are constantly in flux, resulting in the total PFD of the system never remaining constant, meaning you can never have a true fixed average). Process Safety parameters are not point values, they are random variables that change with time. Bayesian statistics provides the correct interpretation for this.

Contrast the Frequentist interpretation of probability to that of the Bayesians. Bayes' Theorem (also known as Bayes' Rule) provides the likelihood of occurrence for one-off events (e.g., the first roll of the dice with unknown bias, the next task, the next operation, the next demand). Bayes' probability is not defined by long-run averages. In Bayes' Rule, qualitative Knowledge (e.g., validation) of a process can be used to quantify the uncertainty of our assumptions about said process (e.g., the reliability of a barrier). It helps us quantify the odds of the safety dice *before* we throw them. The mechanism to do this is the Bayesian Prior. Bayes' Rule works with sparse data, treats parameters as random variables (not fixed point values), and provides a way to update a parameter as new evidence (data) is gathered (as opposed to waiting ad infinitum to pool enough data to make a valid Frequentist inference). Bayes' rule is also able to account for information that may not be showing up in your data.

I don't think there is much controversy is stating that Process Safety practitioners should be Bayesians. That said, a traditional Hardware Reliability person (e.g., Safety Instrumented System purists) may be resistant to accepting Bayesian thinking. I suppose if you could strip the human element from and focus purely on the hardware widget, if you could doggedly collect data on said widget from your plant over several decades, ignoring the fact that the samples separated in time may not be "identical" (a requirement of Frequentist methods), then eventually you could use Frequentist methods to calculate whatever parameter you need. However, this is *not* what you need. You need the parameter *now*. And this is just one parameter of many that you need to reliably infer the rare event frequency for your plant. Bayesian methods are the only logical and rational way to get there.

Another way to describe the difference between Frequentists and Bayesians is to look at what they mean when they use the word "Uncertainty" (a qualitative English word) [2].

To the Frequentist, uncertainty is due to the underlying randomness (i.e., aleatory) of a known but indeterminate outcome. Think of a gaming die, precision manufactured, with equal probability of landing on each of 6 sides (i.e., unbiased). When rolled with sufficient force in open space (a necessary qualifier because there are tricksters that can deterministically flip coins and roll die to a desired outcome), you don't know which side it will land on, but you know the odds in advance, or you could determine the odds by rolling the die unlimited times (assuming the edges of the die do not wear which would change the outcome). This interpretation has very limited application in process safety for the reasons discussed above.

To a Bayesian, uncertainty is associated with our degree of ignorance. The more we know about a process, the less uncertainty there is. Take the die example above. Instead of being handed an unopened package labelled "gaming die," say you found a die on the sidewalk leading to a Bingo gaming room, and you had no knowledge of its inherent bias. Further suppose you forgot to bring your lucky die with you, and so you needed to quickly and reliably estimate the bias of the die by rolling it a few times (think periodic proof-test) on the sidewalk, and that if you used a noticeably biased die in the game room, you would be thrown out (or worse). This, is Process Safety.

**Figure 1** attempts to compare Frequentist and Bayesian interpretations of Uncertainty on a scale representing typical Process Safety systems and processes related to quantification of Barrier performance.

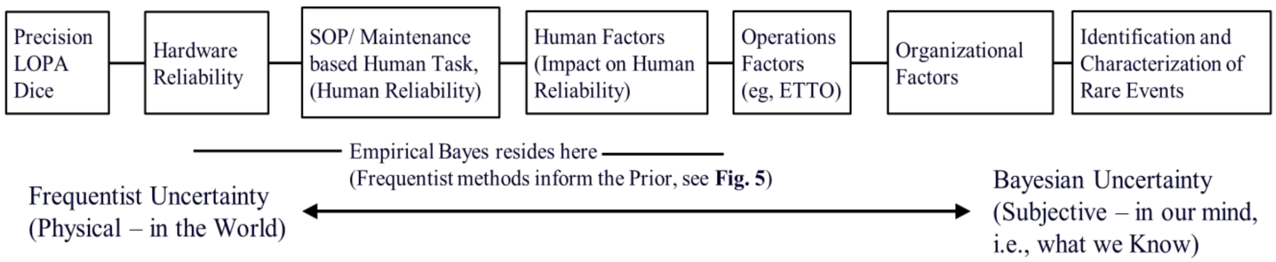


**Figure 1**. Frequentist vs. Bayesian Interpretation of Uncertainty related to systems and processes that affect parameters used to calculate Barrier performance. A significant part of Barrier performance resides in the Subjective Uncertainty portion of the scale. ETTO = Efficiency Thoroughness Trade-offs [3] (e.g., "unknown knows", i.e., things we should know but don't, or things we know but refuse to acknowledge may be a problem).

## 3 How Bayes' Rule works

Bayes' Theorem (also known as Bayes' Rule) is a simple formula for updating current beliefs based on new evidence (data, quantitative and qualitative) as it trickles in. Contrast this to Frequentist methods, which require a large pool of data all at once, to make a valid inference. Depending on the parameter of interest (e.g., initiating event frequency) this could take decades to collect enough data to be "statistically significant" (i.e., 95% confidence level), using Frequentist methods. (A corollary issue is how Frequentist methods update the "true" parameter value when new data/ evidence is collected. This is the Frequentist "magic" alluded to in **Fig. 2**.)

Knowledge is "belief justified." So when we talk about "beliefs" in this context, we are describing a degree of knowledge.

What are the "beliefs" we are trying to validate. Those related to LOPA include for example,

> "I believe the initiating event frequency is x."
>
> "I believe the probability of failure on demand is x."
>
> "I believe the frequency of this rare event occurring is x."

In each case the belief is typically the parameter of interest we are making an inference on. You could also call the belief the hypothesis.

The genius of Bayes' Rule comes in two ideas.

1. Use of a Prior probability (or probability distribution) representing our initial belief "prior to" the collection of evidence (data). Frequentists do not use a Prior.

2. Inference in the correct direction, that is, of the parameter given (that we know) the data. Frequentists make the opposite inference, that is, of the data given (that we know) the parameter. But we don't know the parameter, that is what we are trying to find!

**Section 5** will develop these two ideas more fully.

Bayesian inference infers probability directly from data (via the Prior). Frequentist inference assumes the (long-run) relative frequency *is* probability. Firstly, a specific plant/ facility often doesn't have long-run frequency data related to process safety events (or, it would take decades to gather). Second, to make the assumption valid (i.e., that relative frequency *is* probability), complex techniques are needed, that often remain hidden to the user. Which leads to the potential for abuse of the methods. **Figure 2** shows this schematically.

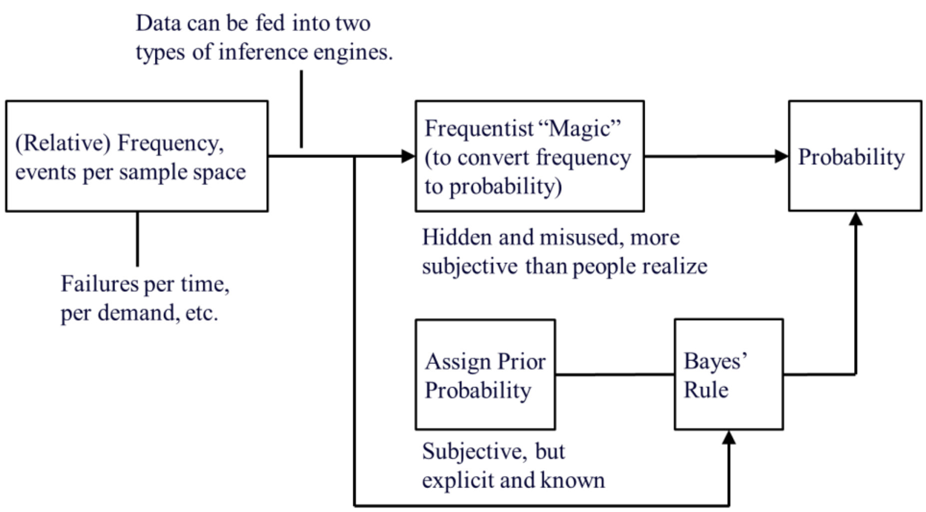


**Figure 2**. Bayesian inference is superior to Frequentist inference because it infers probability directly, versus Frequentist inference which requires tests and comparisons to hypothetical samples that may not be real. Because of the hidden complexity, Frequentist methods are open to abuse, as recently cited by the American Statistical Association [4].

A theorem is a mathematical statement that has been proven to be true. The mechanics of computing probabilities using Bayes' Theorem are described in **Figures 3 and 4** [5].

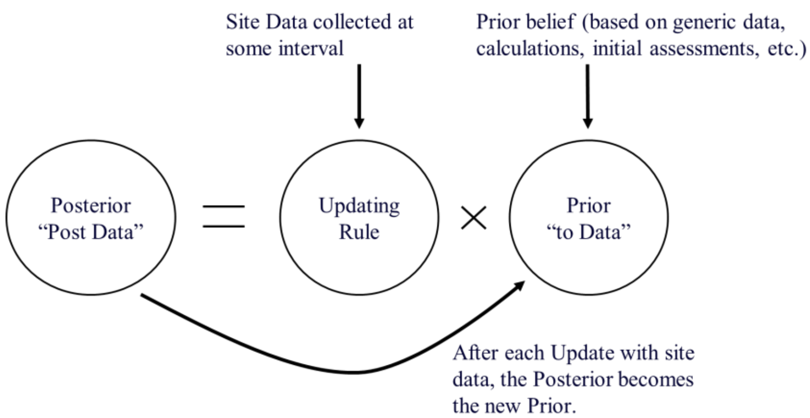


**Figure 3**. A Bayesian engine in simplified form. Prior belief is the starting point typically based on calculations using generic data and initial audits and assessments. As the Bayes engine runs, the confidence in the Posterior result increases. The Posterior is an inference on the future, i.e., between the Prior and the next update.

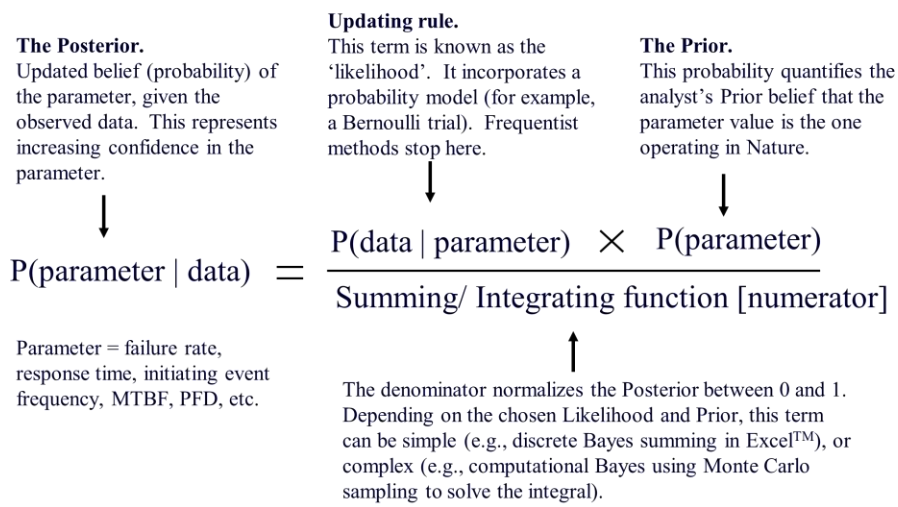


**Figure 4**. Bayes' Rule. Bayes' Rule works with both probabilities or probability distributions. Frequentist methods stop at the Likelihood term and then subject it to a "p-value" test *on the data* assuming the parameter is true. This is the opposite inference we need. Calculating the value of the denominator determines the complexity of implementing Bayes, ranging from Excel™ based discrete methods, to Markov Chain Monte Carlo (MCMC) computational methods.

## 4 Process Safety, Bayes, and the Problem of Induction

Little is known about Thomas Bayes. He was a Presbyterian minister trained in theology and mathematics at the University of Edinburgh. As a Presbyterian, he was an English Dissenter (from the Church of England) and as such faced persecution. His intertwining of theology and mathematics to answer big theological questions based in mathematics such as, "what is the conditional probability Jesus rose from the dead given eyewitness testimony?" helped lead him to his famous rule, which was published a few years after his death (c. 1761). [6]

Although Bayes never could definitively prove that "God exists," or that "Jesus rose from the dead," Bayes' rule was seen at the time (and today) as a credible challenge to answer David Hume's famous Problem of Induction (i.e., the Future is only knowable when it occurs, at which time it has already become the Past). The Problem of Induction even today is still unsolved. Bayes' Posterior is an attempt to glimpse the Future. Frequentist methods have no ability to do this, relying solely on looking at the Past (data).

A lot of Process Safety practice involves trying to glimpse the Future. Think about *Leading* Indicators. They are not intended to be backward looking (we already know the Past), rather what we need is an indication of where we are today and where we are headed tomorrow. In this way too, Leading Indicators are an answer to the Problem of Induction.

Similarly, the probabilistic calculations we make to infer a rare event frequency are meant to be forward looking. What good would it be to make a statement only on the Past (data)? We care about tomorrow. Don't confuse inductive inference as being a prediction. It is a statement of how trustworthy a prediction is.

All that said, the Bayesian Prior is a key component to making the inference.

Diaconis and Skyrms say it best [7],

> "…if you were going to risk a lot on the next few trials, it would be prudent for you to devote some thought to putting whatever you Know into your Prior."

For process safety, the next "few trials" is referring to the next "day's" Operation, task, or demand on a safety critical function. And this is the basis of Empirical Bayes. Using Frequentist methods to inform the Bayesian Prior. **Figure 5** shows this graphically.

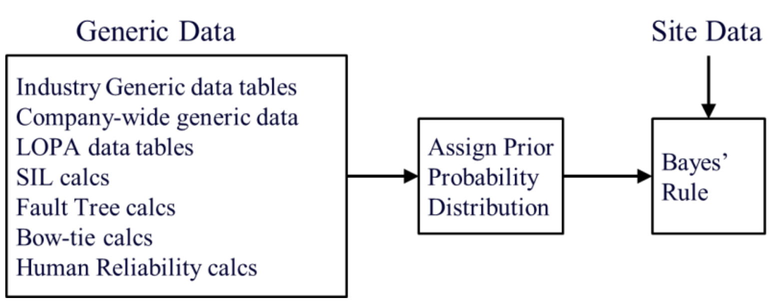


**Figure 5**. Empirical Bayes. The front-end engineering work becomes an important part of selecting the Prior Uncertainty distribution. Depending on the site data update rate for the function, task or operation (i.e., of the inferred parameter), the Prior can dominate for years, or be quickly over-taken by the updates.

## 5 How to Think like a Bayesian

Human beings are not good probability calculators. We have not evolved a probability sense organ. And our use of heuristics (mental short-cuts and other rules-of-thumb applied incorrectly) results in systematic bias when making probabilistic estimates. The Literature is replete with examples [8, 9]. The judgements we make in **Figure 7** are susceptible to the same.

Humans are good at collecting relative frequency data on the past (e.g., this event occurred 5 times out of 20). Converting this to a probability of occurrence for the next trial (i.e., the $21^{st}$) is what Bayes' rule does [10].

Bayes' rule uses two concepts that can help Process Safety practitioners become better Probability Thinkers, apart from using the rule itself. They are:

1. Making sure you're using conditional probability, *in the correct direction*.
2. Don't neglect the Prior, also known as the Base Rate.

**Figure 4** shows that the concept of conditional probability is a key part of Bayes' rule. Further, **Figure 4** shows two conditional probabilities, in opposite directions. First thing to note is, the probabilities are not the same. It depends on what direction you are looking. Four examples will be given.

**Example 1** is from Bayes' Rule itself.

P(data | parameter). "The probability of getting the data given that I know the parameter". This is the conditional probability used by Frequentists. It's in the wrong direction! You don't know the parameter, that is what you're trying to find! Also, Frequentists don't use a Prior, they neglect it.

**Example 2** is from Pop Culture.

In Season 1 Episode 17 of the Cosby Show (1984), titled "Theo and the Joint," Cliff and Claire find a marijuana cigarette in one of Theo's text books from school. They get in a discussion of whether or not it is *his cigarette*, and being in *his* text book, Claire (the lawyer) is ready to convict. But she is getting her conditional probability wrong.

P(Joint being in Theo's text book | It's his Joint). This is the conditional probability Claire was calculating, notice, *assuming it's his joint*. This conditional probability is near 100%.

P(It's his Joint | Joint being in Theo's textbook). This is the conditional probability Claire should have been calculating. Which is much lower. As it turns out, it was not Theo's joint. A classmate had hidden it there when the teacher walked in the classroom. What Claire was also missing was the Prior (her belief before seeing the joint in Theo's text book, that Theo was a pot smoker).

This example is typical of many problems seen in the Justice System (answering the wrong conditional probability). Ref [11] gives several good examples.

**Example 3** is from Process Safety. This one is notorious. Suppose I have a data trend that shows the number of incident or accidents, is trending down-ward. We are operating safer, right! Not so fast.

P(Data trend downward | We are operating safer). This is the conditional probability most people answer in this type of situation, which is in the wrong direction.

P(We are operating safer | Data trend downward ). This is the correct conditional probability, and depends on many factors (e.g. the Prior being one), that may not be reflected in the data.

**Example 4** is whimsical, but one of my favorite. Study **Figure 6**.

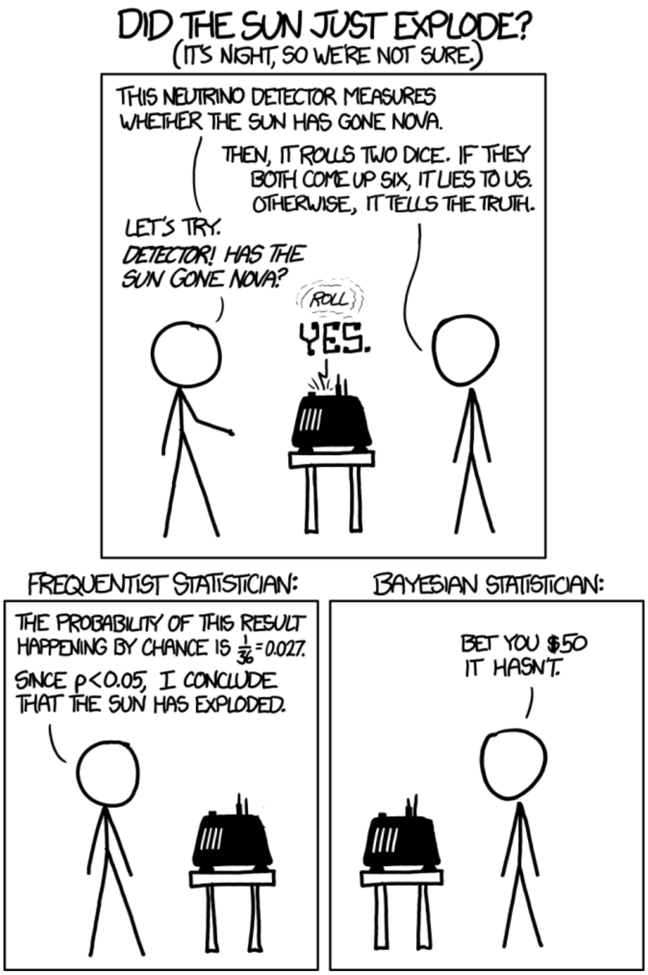


**Figure 6**. A Twist on the Sunrise Problem from Philosophy. Source: xkcd: A webcomic of romance,
sarcasm, math, and language. https://xkcd.com/1132/

P('getting' the data | the sun has exploded). Similar to **Example 1**.

P(the sun has exploded | the data). This is the conditional probability you want, and the Bayesian knows it! And again, accounting for the Prior (i.e., a billion+ previous sunrises) should not be neglected!

## 6 Converting Qualitative Knowledge (Subjective Belief) to a Prior Confidence Level

**Figure 5** shows how it is possible to use quantitative Frequentist methods to inform the Prior distribution. This is possible, because the Frequentist methods themselves incorporate a statement of confidence. For example, SIL Calc standards (i.e., S84/ 61511) require a 70% single-sided confidence on the failure rate data used (i.e., 70% confident that the True value falls within the upper bound, i.e., equal to or lower than the value used). Of course, the problem with this is, the data is *generic*, it did not come from my plant. So with no further information, the actual confidence in the number (in my application) is *unknown*. Think about that! Enter the qualitative life-cycle methods (i.e., Validation, Assessments, audits, etc.).

But how do we convert purely qualitative Knowledge (i.e., from Validation, Assessments, audits, etc.) into a number that can inform the Prior distribution? That is the subject of this section.

Of all the content in this paper, this is the most subjective. From **Figure 1**, we are on the far right side of the scale, into Operational and Organizational factors for which there is no generic data tables to calculate a parameter (as a Prior). That said, there are factors (or multipliers) from Human Reliability science that can offer guidance to scale a parameter off of some base rate. See for example Refs [12, 13, 14, 15, 16]. Still, many operational and organizational factors don't have even a scaling rate, and therefore are purely qualitative.

The good news is the Bayesian Prior is completely transparent (see **Figure 2**), such that if you don't agree with my numbers, one, you know it, and two, you can come up with your own that you think are better. Conceptually what we are trying to do is shown in **Figure 7**.

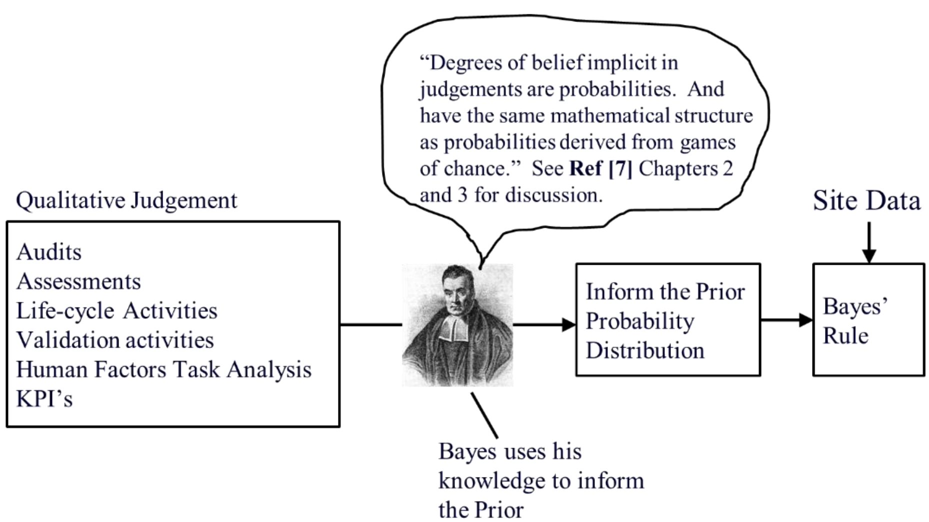


**Figure 7**. Converting Knowledge into a number is an exercise in Subjective Probability. Site data which is quantitative will soon dominate the Prior if the parameter you are making the inference on can be easily measured.

Frequentist Statistics says my confidence increases as the square root of 'n' times increase in the sample size, e.g., after 4 samples, my confidence level increases by a factor of 2 [9]. If each sample is a Bernoulli trial (success/ fail), by the 4th sample my confidence has gone from (for example) 50% to 75% (see **Figure 8**).

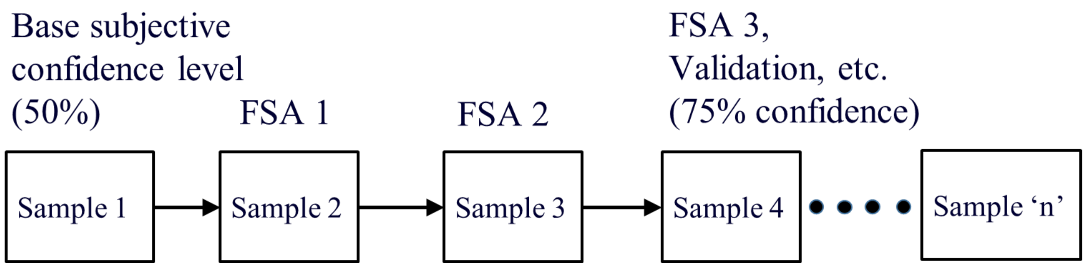


**Figure 8**. An attempt to quantify qualitative Knowledge. A 50% base confidence represents a site's attempt to meet industry best practices related to PSM including Barrier management, derived from lessons learned of high-profile catastrophic accidents. The number 50% is a purely subjective assignment of probability. It is meant to approximate the confidence in an average. If a site were missing a particular segment of best practice, the initial subjective confidence would be lower. Each sample represents an assessment or audit, increasing knowledge about the parameter of interest. After 4 samples I've reached 75% confidence in the parameter, this assuming each sample represents a Bernoulli trial (success/ fail). FSA = Functional Safety Assessment.

The caveat of **Figure 8** is the confidence is increasing for the "knowns" i.e., what you find during an assessment. The same square root of 'n' rule does not apply to "unknowns" i.e., what you're not sampling. Here, confidence increases much more slowly for the absence of findings. This is contributory to the Black Swan problem. For this reason, the Prior uncertainty distribution should always allow for unknowns. This is a qualitative judgement. There is no failure data that will contain information on a Black Swan event (until it happens, then it's too late!).

## 7 Bayes' Rule and LOPA Inference, an Overview.

**Figure 9 and 10** show how we are and are not using the Bayes' engine to infer the population parameter. Recall from **Figure 4** the types of parameters we are interested in for LOPA. In each case, we are concerned with the "right tail" of the distribution, which corresponds to high initiating frequency, high probability of failure on demand, etc. We use a cumulative distribution to make a statement that we are equal to or less than a certain value (that corresponds to the confidence we want). In this way, we move decision making from "Pass/ Fail" (as is practiced today), to a degree of confidence. Think of the implications. The following list represents a potential FAIL of the risk target in each case. Decision making would be improved if we instead looked at how the confidence in the desired parameter value has been affected.

1. I'm off my risk target by a factor of 3.
   Bayes says: "We don't need to do anything because my confidence level at that increased factor hasn't changed."
2. My risk targets changed an order of magnitude more conservative.

Bayes says: "Before we go trigger a massive capital spend to close gaps, let's evaluate how the confidence in meeting the rare event freq. has changed for each scenario."

3. My generic SIL calc shows I have a residual gap.
   Bayes says: "Before I go spend money, I want site specific quantification of said gap. Assess and collect data over the year and run a Bayes' engine. Then make a decision"
4. My initiating frequency is higher than assumed in my calculation, but Operations is making changes to fix it.
   Bayes says: "Lower the confidence level in the (assumed) initiating event frequency and monitor next year's operation closely. Update Bayes' engine with new data and see where we are."
5. My Stage 4 FSA (functional safety assessment) found I can't test this safety function as often as I initially assumed.
   Bayes says: "Redo the SIL calc and input to the Prior. How has the confidence level in the target changed?"

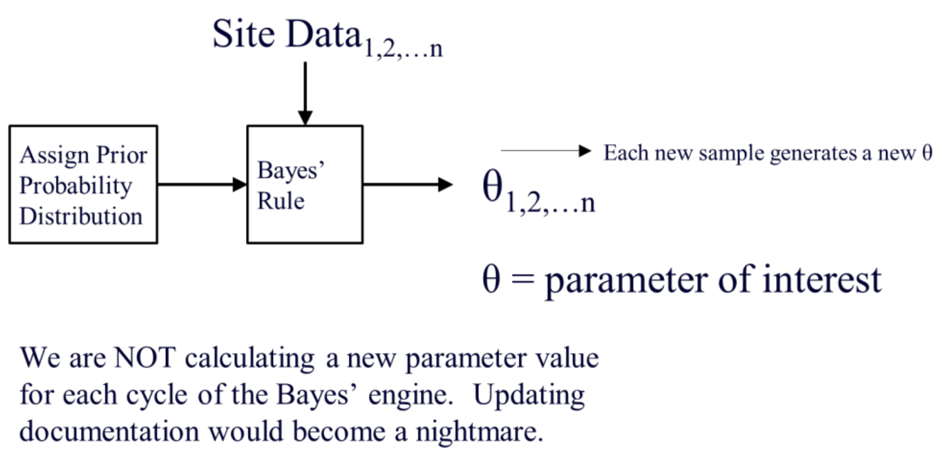


**Figure 9**. What we are NOT doing with Bayes (but you could do).

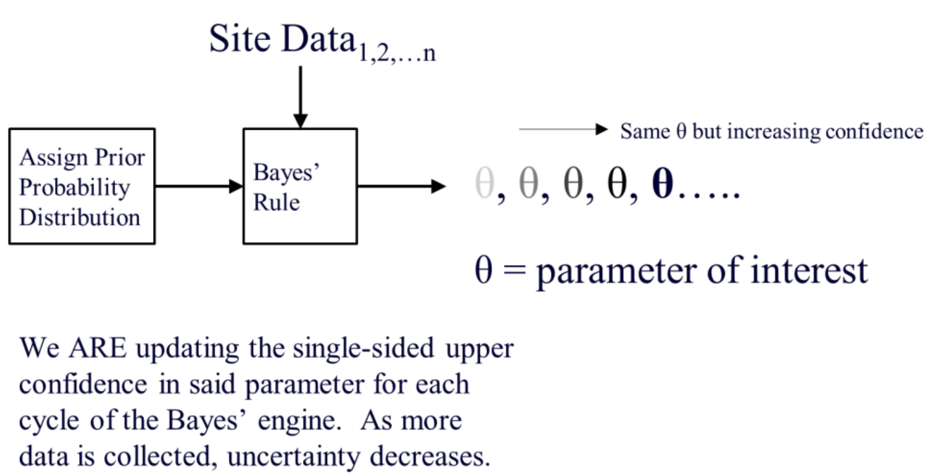


**Figure 10**. What we ARE doing with Bayes.

**Figure 11** is intended to show the big picture of how all this would look. For each IPL (independent protection layer) as well as the initiating event frequency, a Bayes engine can be built to perform what we show in **Figure 10**. An obvious and easy place to begin is with inferring the initiating event frequency. Generic initiating event frequency data typically used in LOPA is wildly inaccurate in many cases. Often times initiating event rates are much greater than the generic 1/10 years. In other cases, certain demands have never been seen in the history of the plant. You may object by saying, "what's wrong with using a generic Frequentist based average?" Answer: because you can't demonstrate you are operating safely based on generic averages. Also, in each case, there are important qualitative insights to be derived from a closer quantitative look at the initiating event frequency. For example, for events happening more frequent than assumed, the risk is increased. Instead of say, once per 10 years, think of rolling the LOPA dice once per year instead (add in the element of unknown bias and the situation is worse). For a more rare initiating event, the qualitative take-away is operation's lack of actual practice with the event, decreased situation awareness, and decreased ability to respond when things don't go as planned (e.g., the automated trip doesn't activate).

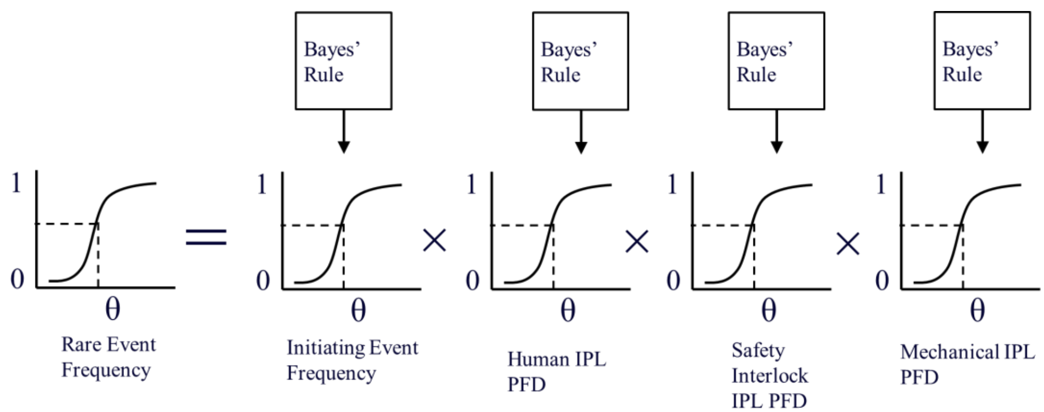


**Figure 11**. Putting it all together. Each parameter of the LOPA equation can be inferred using a Bayes' engine. The result is a cumulative distribution of the rare event frequency, from which the single-sided upper confidence limit in the LOPA target can be determined. This provides a better decision making tool than using a single point value with unknown confidence for the site application. As data is gathered and the Bayes' engines run, confidence in meeting the target will increase. Using Bayesian inference is the most consistent and rational way to incorporate site data to update our belief in meeting rare event targets. This method can be easily implemented in Excel™ (see **NUREG-6823** Chapter 6 for details).

## 8 Conclusion

Process Safety practitioners have the choice of using either Frequentist or Bayesian based inference methods. This paper has laid out the case for Bayes. Ultimately it is about making better decisions related to managing rare event scenarios. The Bayesian interpretation is where we start with a Prior informed by engineering design, assessments, and audits, and then update to a Posterior when periodic proof-tests, incident reports, continuing audits, etc. are evaluated, and is the only rational and consistent way to build confidence as a degree of belief that we are operating safely.

## 9 References


1. Hans Pasman and William Rogers, "How trustworthy are risk assessment results, and what can be done about the uncertainties they are plagued with?" Journal of Loss Prevention in the Process Industries, 55 (2018) 162-177.
2. Kristin Lennox, "All About that Bayes: Probability, Statistics, and the Quest to Quantify Uncertainty" https://www.youtube.com/watch?v=eDMGDhyDxuY
3. Hollnagel, E., Safety – I and Safety II The Past and Future of Safety Management, Ashgate, 2014.
4. Ronald L.Wasserstein, "ASA Statement on Statistical Significance and P-Values", The American Statistician, 70:2, 129-133
5. Ben Lambert, "A student's Guide to Bayesian Statistics", Sage Publishing, 2018.
6. Judea Pearl, "The Book of Why – The New Science of Cause and Effect", Basic Books, New York, 2018.
7. Persi Diaconis and Brian Skyrms, "Ten Great Ideas about Chance", Princeton University Press, 2018.
8. Daniel Kahneman, "Thinking Fast and Slow", Farrar, Straus & Giroux, New York, 2011.
9. Nassim Nicholas Taleb, "Fooled by Randomness – The Hidden Role of Chance in Life and in the Markets", Random House, New York, 2004.
10. James R. Henderson, "Bayesian Meets Frequentist: A Marriage Made in California", published in "Probability is the Very Guide of Life", Ed. Henry Kyburg. 2003.
11. Leonard Mlodinow, "The Drunkard's Walk – How Randomness Rules our Lives", Pantheon Books, New York, 2008.
12. Robert Taylor, "Human Error in Process Plant Design and Operations – A Practitioner's Guide", CRC Press, Taylor & Francis Group, Boca Raton, 2016.
13. Barry Kirwan, "A Guide to Practical Human Reliability Assessment", CRC Press, Taylor & Francis Group, Boca Raton, 1994.
14. Swain and Guttmann. NUREG/ CR – 1278. "Handbook of Human Reliability Analysis." 1983.
15. James Reason, "Organizational Accidents Revisited", CRC Press, 2016.
16. Institute for Energy Technology, "The Petro-HRA Guideline." Report IFE/HR/E-2017/001, 2017.