Ostrich Cyber-Risk

Birdseye™ Level 2

Calculating The Effect of NIST CSF Maturity Levels on Risk Reduction

Version 1.0.4

11/6/2022

Birdseye™ Level 2

Calculating The Effect Of NIST CSF Maturity Levels On Risk Reduction

Table of Contents

1. Introduction

This paper presents a simple method of Monte Carlo simulation for calculating the risk reduction effect of Maturity Level on preventive cybersecurity controls. The simulated ALE (Annual Loss Expectancy) is used as a quantitative measure of risk for a hypothetical scenario. A conversion table is first used to map NIST CSF Maturity Levels to Resistance probability distributions. The expected risk for the scenario is then calculated at the Vulnerability level (for existing inherent risk) and at each of the five NIST CSF Maturity Levels (for residual risk after adding a new control). The average ALE is tabulated at each level of residual risk for the hypothetical scenario to illustrate the risk reduction effect of Maturity Level. Table 1 (below) shows a summary of the simulation results.

		O-FAIR	FAIR-U	B-FAIR	B-Binary	B-Prop
Primary LM = ($100K, $200K, $1M) Loss Event Frequency = (1, 2, 3)		Ave ALE	Ave ALE	Ave ALE	Ave ALE	Ave ALE
		$858,000	$633,100	$633,250	$633,250	$633,250
Threat Event Frequency = (1, 2, 3) Vulnerability = (50%, 75%, 100%)		Inherent Risk Shown Below				
		$645,200	$473,425	$475,000	$475,000	$475,000
Threat Event Frequency = (1, 2, 3) Threat Capability = (50%, 75%, 100%)		Residual Risk Shown Below				
NIST CSF PRISMA	Resistance [Strength]	Ave ALE	Ave ALE	Ave ALE	Ave ALE	Ave ALE
Maturity Level 0	(0%,0%,1%)		$632,750	$633,000	$475,000	$474,000
Maturity Level 1	(18%,20%,22%)	$859,900	$632,775	$633,000	$475,000	$380,000
Maturity Level 2	(45%,50%,55%)	$837,700	$632,775	$633,000	$475,000	$237,500
Maturity Level 3	(72%,80%,88%)	$291,200	$203,575	$204,500	$174,000	$95,000
Maturity Level 4	(81%,90%,99%)	$99,600	$42,250	$44,750	$41,000	$47,500
Maturity Level 5	(88%,98%,100%)	$48,200	$3,625	$3,500	$3,000	$16,000

Table 1 - Comparison of Effect of NIST Maturity Level on Annual Loss Expectancy for Hypothetical Scenario

The column labeled "O-FAIR" was generated using the free Open FAIR™ spreadsheet tool from The Open Group™. The Column labeled "FAIR-U" was generated using the free FAIR-U™ website from RiskLens™. The columns labeled "B-*" were generated using the Birdseye™ risk simulator. The cells highlighted red in Table 1 indicate Maturity Levels where the calculated residual risk exceeded inherent risk in FAIR-type calculations for this scenario.

If you are interested in the internal details of the PRA (Probabilistic Risk Analysis) calculations used to create Table 1, continue reading at Section 2 (Page 4).

If you are interested to learn how to perform this same type of risk-reduction analysis using your own preferred Monte Carlo simulation tool, continue reading at Section 9 (Page 18). A summary of the risk analysis results starts at Section 13 (Page 27).

2. Probabilistic Risk Analysis

The Birdseye risk simulator uses a standard method of risk calculation knowns as PRA (Probabilistic Risk Analysis). Depending upon the context, PRA can also stand for Probabilistic Risk Assessment. Engineers are taught to use Assessment to answer "What is?" and to use Analysis to answer "What if?". But a more practical distinction is that Process Engineers take off Event Trees from Process Flow Diagrams for Risk Assessment but take off Fault Trees for Risk Analysis. Another distinction is that the sort of PRA used for CRQ (Cyber Risk Quantification) has nothing at all to do with the sort of Risk Assessment/Analysis done by professional Process Engineers (in the author's opinion). This paper will only discuss the sort of PRA used in CRQ.

The type of PRA often used in CRQ (including Birdseye, FAIR, and Open FAIR) is sometimes referred to as the "simplest-possible" expression of PRA because this was the presentation used in the nationally televised debate in 1976 between Prof. Norman Rasmussen and Ralph Nader regarding the safety of nuclear power. For the television audience, ***Rasmussen was constrained to keep the mathematics in his presentation within reach of an 8th grade education***. Rasmussen used Probabilistic Risk Analysis in his papers and debate, and this is also used by the Birdseye risk simulator. Despite the Three Mile Island (TMI) nuclear meltdown that followed shortly after the televised debate, this simple method of PRA later found wide acceptance, particularly in Investment Analysis and related fields. It is acceptable to use this type of PRA when only money is at hazard because a court of law can always resolve any disputes that may occur. However, no court of law can satisfactorily resolve claims of unnecessary loss of human life, so this simple type of PRA may not be used in engineering when human safety is at hazard.

This section will attempt to show:

- How parts of basic PRA are used in common by the Birdseye, FAIR, Open FAIR and other CRQ simulators,
- How an improper addition ("Resistance Strength") to this PRA was popularized by the FAIR and Open FAIR frameworks, and
- How instead to add "Resistance" to PRA calculations using the standard definition of the term from Physics.

Calculating the effect of the NIST CSF Maturity Levels on risk is an excellent demonstration of the similarities and differences between of preventive cybersecurity controls. For example, this Maturity Level analysis shows how the "Resistance Strength" calculation from FAIR and Open FAIR can allow calculated residual risk to exceed inherent risk. Analysts outside the field of CRQ would generally consider this to be an unacceptable feature, so Birdseye also offers the "Resistance" calculation option. These same calculations for Maturity Level also show that inherent risk is the natural upper limit of residual risk when the "Resistance" calculation is used.

The rest of this section deals with PRA history, theory, and equations. *Readers who do not immediately wish to devote time to these topics can now skip to Section 6 (Page 16).*

The essential PRA equation that Prof. Rasmussen presented on television is shown below:

$$\text{A Risk} = (\text{Probability of Failure}) \text{ X } (\text{Consequence}) \quad 2.01$$

This equation expresses the fundamental relationship of EVT (Expected Value Theory) which has been widely known and used for centuries. In this formula, Consequence is the impact of Failure. Types of Consequence can include Economic, Environmental, Safety, Security, etc. An appropriate method is used to evaluate the impact of each different type of Consequence.

In the field of CRQ, Consequence is generally restricted to Economic Impact, and so "Loss Event Magnitude" (in currency units, such as dollars) is substituted for "Consequence" in Equation (2.01). Because the scope of Consequence is thus limited to Economic Impact, the calculated "Risk" should always be normalized for a specific time frame. Thus, economic "Risk" can be replaced here by "Annual Loss Expectancy" (defined in dollars per year).

Also, CRQ analysts commonly do not attempt to compute probability density functions for their system states. Consequently, a substitution must be made in Equation (2.01) to replace Probability with Frequency using the Frequentist Theory of Probability. Thus, "Probability of Failure" is replaced by "Loss Event Frequency". This latter substitution is justified by the fact that in this theory Probability is defined to be the limit of Frequency as the Number of Trials goes to infinity. This substitution is usually a very good approximation as the Number of Trials approaches infinity, but it is often a poor approximation as the Number of Trials approaches zero. These three substitutions yield the equation below:

$$\text{Annual Loss Expectancy} = (\text{Loss Event Frequency}) \text{ X } (\text{Loss Event Magnitude}) \quad 2.02$$

We will now begin adding these equations to a simple Conceptual diagram to show how common PRA is related to the risk ontology used in FAIR and Open FAIR. The "bubbles" on this type of Conceptual Design artifact can include pictures, language, and pre-calculus mathematics. The "lines" on the diagram show the Relationships between Concepts. The bubbles on the left side of the diagram show a Conceptual expression of the common PRA equation, while the bubbles on the right side of the diagram show how this same equation is expressed in the ontology of FAIR and Open FAIR. For consistency with FAIR and Open FAIR diagrams, mathematical relationships between Concepts are not explicitly depicted on the diagram.

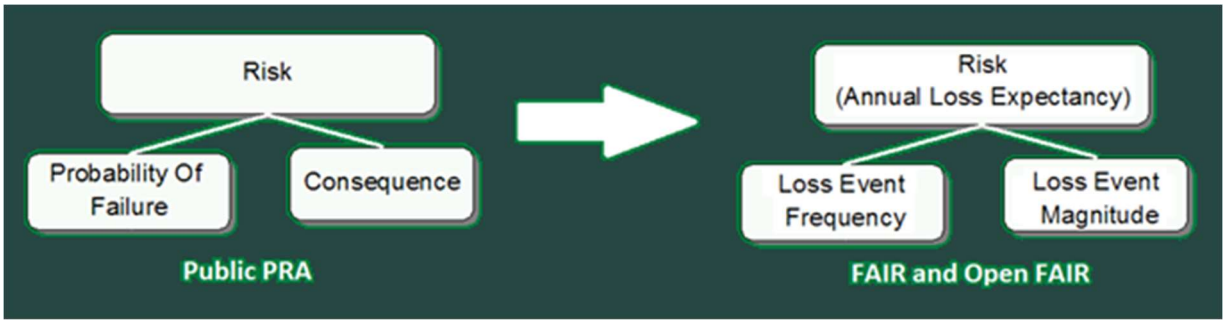

Figure 1 - Risk = (Probability of Failure) X (Consequence)

The next fundamental PRA equation tries to account for human influence in risk. This equation is shown below:

$$A\ Risk = (Threat)\ X\ (Vulnerability)\ X\ (Consequence)^{2.03}$$

This influence is added by assuming that the "**Probability of Failure**" can result from the product of two independent probabilities:

1. the probability of attempted failure ("**Threat**"), and
2. the conditional probability of success for an attempt ("**Vulnerability**").

We can therefore substitute "**Threat** X **Vulnerability**" for "**Probability of Failure**" in our Conceptual diagram, as shown in bubbles on the left side of Figure 2 (below). The bubbles on the right side of the diagram show how this same equation is expressed in the ontology of FAIR and Open FAIR. One potential problem with this substitution is that the assumption of statistical independence between "Threat" and "Vulnerability" may not be justified.

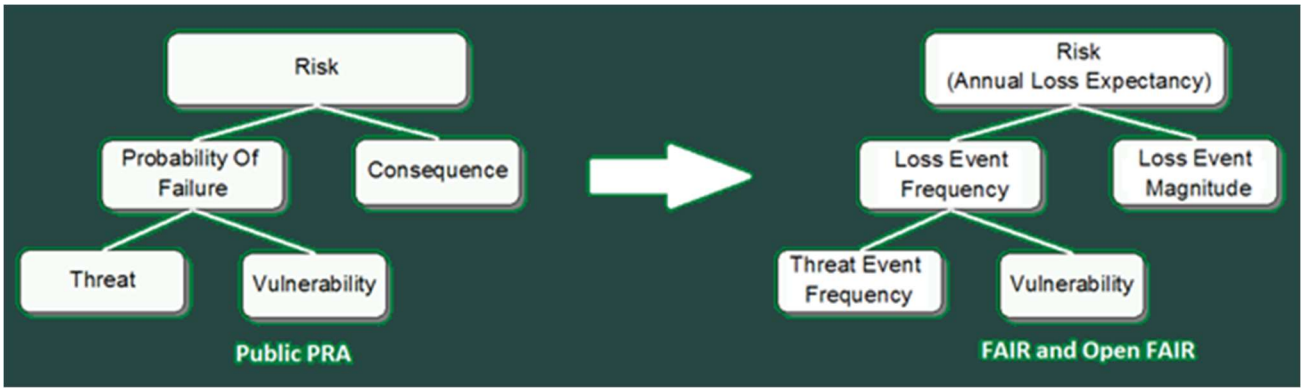

Figure 2 - Probability of Failure = (Threat) X (Vulnerability)

After the 911 terrorist attack in 2001, the USCG (United States Coast Guard) began working to add the effects of terrorism into PRA. The USCG started releasing results of their work in 2004, with a final report in 2005 that defined their **MSRAM (Maritime Security Risk Assessment Method)**. Probably the best-known equation from MSRAM is:

Threat = (Intention) X (Capability) [2.04]

Here "Intention" is the probability that an attack will be made, and "Capability" is the conditional probability that the attack will succeed. In the FAIR and Open FAIR ontology, "Intention" is replaced by "Probability of Action" and "Capability" is replaced by "Contact Frequency". There is one potential problem with this substitution of "Threat" for analysts using the NIST CSF (Cybersecurity Framework). The current version of the NIST CSF does not seem to include any cybersecurity controls that map directly to Factors below "Threat". By contrast, NIST CSF controls in the "Protect" function can be mapped to Factors below "Vulnerability", while NIST CSF controls in the "Detect", "Respond" and "Recover" functions can be mapped to Factors below "Consequence".

We can therefore substitute "Intention X Capability" for "Threat" in our Conceptual diagram, as shown in bubbles on the left side of Figure 3 (below). The bubbles on the right side of the diagram show how this same equation is expressed in the ontology of FAIR and Open FAIR. At this point, the algebraic formulae used on both sides of Figure 3 are mathematically equivalent (just the variable names are changed) EXCEPT for the bubbles labeled "Consequence" and "Loss Event Magnitude".

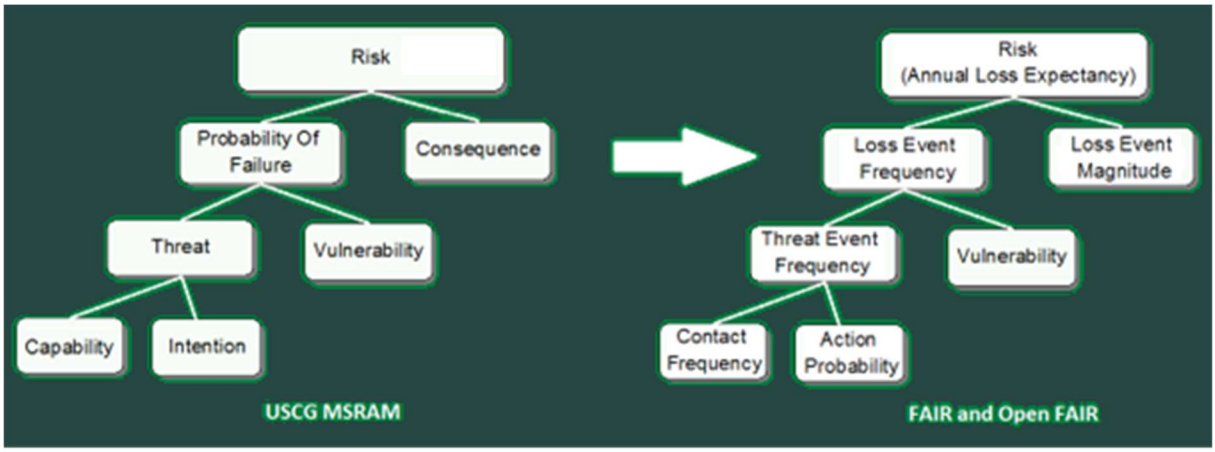

Figure 3 - Threat = (Intention) X (Capability)

Economic Impact analysis in engineering is almost universally done using standard CFA (Cash Flow Analysis). CFA obviously requires a modicum of calculus (e.g., to compute Marginal Rates of Return), but this can usually be made fully transparent to users of modern computerized systems. In the same way, cash flow indexing can usually also be made transparent for users. Cash flow re-indexing for Probability can be easily accomplished by Expected Value calculations, and re-indexing for Time can be done using TVM (Time Value of Money) calculations. Re-indexing cash flows for Space (geographical location and currency) can also be made transparent for users, but this requires statistical correlation from actuarial data (i.e., a database). One reason for the popularity of CFA in engineering is that it properly handles economic contingencies (e.g., Opportunity Cost).

A simplified approach for Economic Impact analysis in CRQ has been popularized by the FAIR and Open FAIR frameworks, becoming a sort of de facto standard in the field. This approach consists of using a

required cash flow with 100% probability and an optional cash flow with conditional probability. The Total for each cash flow can be decomposed into six pre-defined Sub-Total categories. TVM for cash flows is not considered. An obvious advantage to this approach is that it eliminates any potential need for calculus by reducing the calculations to simple arithmetic. But this simplification may also introduce a sort of weakness into the economic analysis. Let us consider negative cash flows to be a type of "bill" that must eventually be paid. In engineering CFA, in addition to optional properties (like Probability) the analyst must typically provide at least the following for each potential "bill":

- Amount due to be paid,
- Time when payment is due, and
- Party responsible to pay (e.g., standard organizational accounting code).

The simplified approach to Economic Impact analysis does not require any specification of Time or Party for bills. This lifts a significant burden from the analyst, but it also makes it very easy for analysts to "pad" their economic analysis with "bills" that may never actually become due to be paid or that cannot be addressed to any responsible party for payment. This weakness should logically reduce the relative Utility of the simplified economic analysis, but that seems to be a reasonable trade-off considering typical quality of input data and intended use for CRQ.

We therefore replace Cash Flow Analysis in our Conceptual diagram with the simplified equations for Economic Impact analysis used in the FAIR and Open FAIR frameworks, as shown in the bubbles on Figure 4 (below). After this substitution, the algebraic formulae used on both sides of Figure 3 (Page 7) are mathematically equivalent.

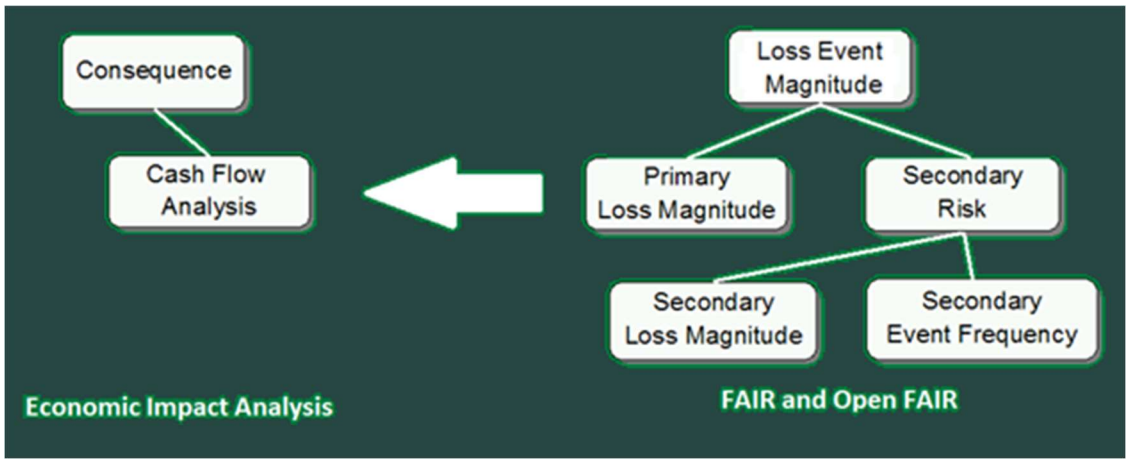

Figure 4 – Replacement of Cash Flow Analysis for Economic Impact Analysis

The final substitution to be made on the Conceptual diagram will create a difference between the calculations used by the Birdseye simulator versus the calculations used by the FAIR and Open FAIR frameworks. In FAIR and Open FAIR, this next equation adds the Factors "Threat Capability" and "Resistance Strength" below the "Vulnerability" Factor. The author of this paper was unable to determine the exact source of the mathematical equation that was used. This was in part complicated by the

"wandering ontology" of the "Resistance Strength" Factor over time. As usual, no mathematical relationships are expressed in the FAIR and Open FAIR Conceptual diagrams, so the only source of such relationships available to the author was narrative text in the proprietary (copyrighted) literature for these frameworks. Using the obvious mathematical interpretation from the narrative text, results from the author's first version of the Birdseye risk simulator were a surprisingly good match to results from the certified "FAIR-compliant" simulator available on the FAIR-U website (e.g., compare the FAIR-U versus the B-FAIR columns in Table 1, Page 3). However, calculating the effect of NIST CSF Maturity Levels on risk reduction immediately demonstrated a serious problem with this "Resistance Strength" calculation implemented from the narrative text description. Based upon the name "Resistance Strength", the author's expectation was that the inherent risk would be the upper limit for residual risk as "Resistance Strength" goes to zero. However, the tabulated results clearly showed that residual risk could rise past inherent risk (based on "Threat Event Frequency" and "Vulnerability") and up to the level of uncontrolled risk (i.e., based on "Loss Event Frequency").

The author reviewed the calculations in the simulator code and immediately noticed that the "Resistance Strength" calculation for derived Vulnerability was ***almost but not quite*** the calculation required for the standard "Resistance" relationship used in physics. Likewise, the narrative text seemed to be just ***eight little words*** short of matching the physical definition of "Resistance". Without any regard for CRQ tradition, the author arbitrarily decided to replace the "Resistance Strength" calculation used by FAIR and Open FAIR with the "Resistance" relationship from physics. The author decided to also add the ability to select between *two of the three standard forms of first-order Resistance*, i.e., binary or proportional.

In physics, a single mathematical equation using differential calculus is used to define Resistance in many types of physical systems (i.e., electrical, thermal, hydraulic, chemical, etc.). In all these systems, Resistance is the first derivative of Potential (or Driving Force) divided by the first derivative of Flow. In Control Systems texts, first derivatives are customarily expressed as Deviation Variables, i.e., ΔX is the first derivative of X. Thus, the standard definition of Resistance would be written as follows:

$$\text{Resistance} = (\Delta\text{Potential})/(\Delta\text{Flow})^{2.05}$$

For the Birdseye risk simulator, the following Information Risk analogy was made for Resistance:

$$\text{Resistance} = (\Delta\text{Threat Capability})/(\Delta\text{Vulnerability})^{2.06}$$

This analogy will be explained further in Section 6 (Page 16). The obvious question at this point is: If "Threat Capability" is the Driving Force and "Vulnerability" is the Flow, then what exactly is being driven or flowing here? A simplistic answer would be "Probability of Loss". By this analogy, we replace "Resistance Strength" in our Conceptual diagram with "Resistance" as shown on Figure 5 (below). Note that Figure 5 would be technically "illegal" in the Engineering Process because it attempts to incorporate calculus into a Conceptual Design artifact. There is an alternative "legal" way to accomplish this same feat, but that is outside the scope of this paper.

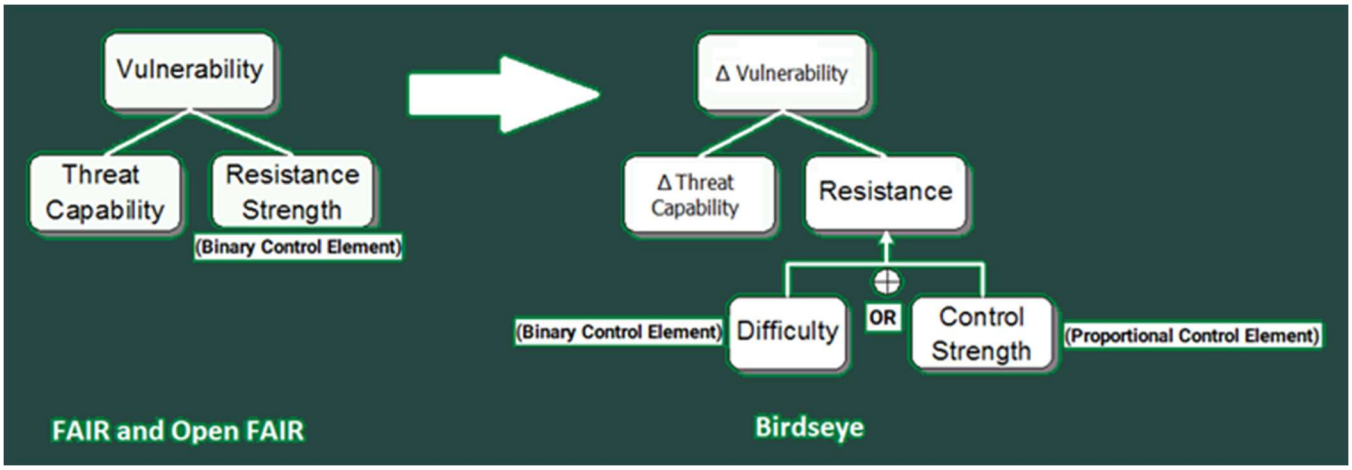

Figure 5 - Replacement of "Resistance Strength" with "Resistance"

The difference in results made by changing the form of the equation becomes obvious if the "Resistance strength" relationship is used in a hydraulic simulation. ***When used for hydraulic simulation, it can readily be shown that the "Resistance Strength" equation can allow transferring more fluid out of a system than was initially present in the system, which would be physically impossible.*** On the other hand, the nature of the "Resistance" equation always guarantees appropriate "balance" for transfer of fluids (or electricity, heat, etc.). For the hydraulic simulation this distinction is vitally important because experiments can be run in a laboratory to determine exactly what is "Reality" in terms of physical hydraulics, thus requiring use of the "Resistance" equation. In the CRQ field, this distinction seems to matter little, possibly because there does not seem to be much laboratory experimentation to determine exactly what is "Reality" in Information Risk. ***When used for Information Risk simulation, the "Resistance strength" equation can allow transferring an "impossible" quantity of "probability", thus making it possible for calculated residual risk to exceed inherent risk.***

This concludes the discussion of PRA needed to calculate the effect of NIST CSF Maturity Levels on risk reduction. Note that the final Conceptual diagrams (Figure 6, Page 14 and Figure 7, Page 15) remain entirely within the limits of 8th grade mathematics. Understanding of Statistics is not required to use these diagrams. For example, anyone can enter the (Minimum, Most Likely, Maximum, Confidence) parameters for a BetaPERT probability distribution without any real understanding of Statistics. By contrast, understanding the formulas required to calculate the first four standard statistical moments of the BetaPERT distribution would require some knowledge of Statistics.

Feedback from early demonstrations of the Birdseye simulator suggests that the Proportional Control Element may be suspicious to some analysts. This new control model conflicts with that of the FAIR/Open FAIR frameworks where all preventive cybersecurity controls can have only Binary Control Effect. In other words, in these popular frameworks every preventive cybersecurity control has a *thresholded return on investment*, so that *investment in improving control strength can possibly have no risk-reduction effect until a certain "threshold" is crossed*. Thus, a control model having Proportional Control Effect (*which consistently provides incremental return on investment*) would be considered unnecessary (at best) and misleading (at worst).

As Table 1 demonstrates, making an incorrect choice of control model could lead to very bad investment decisions. For example, compare the difference between inherent risk and residual risk for Maturity Level 2 in the "FAIR-U" and "B-Prop" columns. An analyst might find that the change in risk due to upgrading from the current level of risk prevention (inherent risk) to a control at Maturity Level 2 would definitely not be worthwhile based on residual risk in the "FAIR-U" column. If we assume that Resistance effect can only be due to binary control effect, then an analyst might be misled into thinking this same investment could be worthwhile based on the "B-Prop" column.

Note that the author of this paper makes no assertion that any cybersecurity controls (e.g., from sections 3.11 or 3.18 in NIST SP 800-53 R5) consistently provide *incremental return on incremental investment*.

However, the author does assert that if any cybersecurity controls providing incremental return on incremental investment *ever should be discovered*, then using a Proportional Control Element in a First-Order Systems model would be an obvious way to analytically model the risk reduction effect of those cybersecurity controls.

3. Return on Investment

Some economic Measure of Merit is required to compare investment alternatives. Common metrics used as Measures of Merit include: Present Worth (e.g., NPV), Benefit-Cost Ratio, Return on Investment, and Rate of Return (e.g., ERR, IRR, DCFRR, etc.). Methods for defining these metrics fall into two basic categories: Private Sector (based on Profit) or Public Sector (based on Benefit). Because there is no profit involved in reducing security risk, the measures defined here are based on Benefit, i.e., using BCA (Benefit-Cost Analysis) aka CBA (Cost-Benefit Analysis). Benefit can encompass a vast multitude of possibilities, but common examples include reduction of cost and reduction of risk. For this example using NIST CSF Maturity Levels, the expected risk reduction is simply the difference between the inherent risk and residual risk.

In the following equations, Inherent Risk is the current risk without improved security, while Residual Risk is the risk after making the proposed security improvement. Investment Cost is the cost of making the proposed security improvement. For simplicity, average Annual Loss Expectancy is used for Inherent Risk and Residual Risk values. Investment Cost is therefore typically amortized to an annual basis over the life of the investment.

Using this approach, the simplest Measure of Merit is used to compare alternatives where no investment is made:

$$\text{Expected Benefit} = \text{Reduction of Risk} = (\text{Inherent Risk}) - (\text{Residual Risk}) \quad 3.01$$

If an investment is made, then the Expected Benefit must be reduced by the invested amount:

$$\text{Expected Net Benefit} = (\text{Expected Benefit}) - (\text{Investment Cost}) \quad 3.02$$

The basic metric for Measure of Merit with investment is the BCR (Benefit-Cost Ratio):

$$\text{Benefit-Cost Ratio} = (\text{Expected Net Benefit}) / (\text{Investment Cost}) \quad 3.03$$

A more common metric for Measure of Merit is the ROI (Return On Investment):

$$\text{Return On Investment in \%} = (\text{Benefit-Cost Ratio}) \times 100 \quad 3.04$$

ROI values expressed as percentages can be used as a convenient way to compare risk reduction effects between scenarios that involve improvements to different preventive cybersecurity controls.

A special name for this Measure of Merit sometimes used in the CRQ field is ROSI (Return On Security Investment), reserved for the special case where "Expected Benefit" in Equation (3.01) can consist ONLY of reduction of risk due to security improvement. The defining equation for ROSI is often given as something like the following:

$$\text{ROSI in \%} = ((((\text{ALE} \times (\text{Mitigation Ratio}))) - (\text{Investment Cost})) / (\text{Investment Cost})) \times 100 \quad 3.05$$

Table 2 (below) tabulates the annual Expected Benefit from Equation (3.01) calculated by five different simulators for adding a new control to the hypothetical scenario at each of five NIST CSF Maturity Levels. If Investment Cost for each Maturity Level were known, then Equations (3.02-3.04) could be used to calculate ROI at each Maturity Level. Maturity Level 0 in Table 2 approximates inherent risk (i.e., very little control effect other than the inherent level), so the Expected Benefit (i.e., risk reduction) for Maturity Level 0 should be approximately zero (by definition).

	O-FAIR	FAIR-U	B-FAIR	B-Binary	B-Prop
NIST CSF PRISMA	Benefit	Benefit	Benefit	Benefit	Benefit
Maturity Level 0	-$214,700	-$159,325	-$158,000	$0	$1,000
Maturity Level 1	-$214,700	-$159,350	-$158,000	$0	$95,000
Maturity Level 2	-$192,500	-$159,350	-$158,000	$0	$237,500
Maturity Level 3	$354,000	$269,850	$270,500	$301,000	$380,000
Maturity Level 4	$545,600	$431,175	$430,250	$434,000	$427,500
Maturity Level 5	$597,000	$469,800	$471,500	$472,000	$459,000

Table 2 - Comparison of Effect of NIST Maturity Level on Annual Expected Benefit for Hypothetical Scenario

One advantage of using the "Resistance" equation (i.e., the B-Binary and B-Prop options in Table 2) is that the Expected Benefit values always work out sensibly, which might be an advantage in presenting results. Using the "Resistance Strength" equation (i.e., the O-FAIR, FAIR-U, and B-FAIR columns in Table 2), having negative values for Expected Benefit may not be noticeable if ROI is computed only for one or two operational alternatives, say for a particular Denial of Service or Data Breach scenario. This is especially true if a relatively high "Difficulty" probability distribution is used. But systematically varying the Levels of Risk (as required to calculate the effect of NIST CSF Maturity Levels on risk reduction) can identify these problems with residual risk. It might be awkward to explain how Expected Benefit can be negative (rather than zero) after investing in additional strength for preventive cybersecurity controls. For example, some people might wonder if trying to add a new cybersecurity control at Maturity Level 1 would increase their risk exposure by automatically compromising their existing level of security.

Feedback from early demonstrations of the Birdseye simulator suggests that negative Expected Benefit values may not be considered a significant issue in the CRQ field. It has been pointed out that when a positive Investment Cost is used, cells in Table 2 with either a zero or negative value will result in negative ROI. The investment alternatives corresponding to these cells will be automatically rejected, and thus they will never be presented for review. So, there is no need for concern regarding any confusion that might be caused by negative Expected Benefit values.

Nonetheless, the author of this paper prefers having zero risk reduction as the minimum possible Expected Benefit for improving cybersecurity by adding Resistance. Consequently, for the rest of this paper, all presented results will be generated using the Birdseye simulator with the "Resistance" calculation option and either the "Difficulty" or "Control Strength" Factors. These options produce the results in the "B-Binary" and "B-Prop" columns in Table 2.

4. Risk Scenario Factors

When creating a scenario in Birdseye, *the available risk Factors include all those from the well-known FAIR™ (Factor Analysis of Information Risk) and Open FAIR™ frameworks*. For more information, see the links below:

https://www.fairinstitute.org/

https://blog.opengroup.org/tag/open-fair/

However, the standard FAIR/Open FAIR frameworks provide only a single Control Element model for computing derived Vulnerability from the factors for Threat Capability and Resistance Strength. In industrial Process Engineering, this model would be considered a typical binary (on/off) Control Element. The Birdseye simulator provides an additional Factor to support a supplementary Control Element model corresponding to a typical proportional Control Element from Process Engineering.

The diagram in Figure 6 (below) shows the typical dependency relationship between the Factors in the FAIR or Open FAIR frameworks. At various times, these frameworks have used the terms "Control Strength", "Resistance Strength", and "Difficulty" to all refer to the same highlighted box in this Factor dependency diagram, i.e., the Factor that is used in conjunction with the "Threat Capability" Factor to calculate the derived "Vulnerability" Factor. The current FAIR standard seems to have dropped "Control Strength" from its ontology altogether, and "Resistance Strength" now appears to be synonymous with

"Difficulty". Thus, FAIR and Open FAIR have always provided only one control model for "Resistance Strength", although the Factor name on the diagram has changed over time.

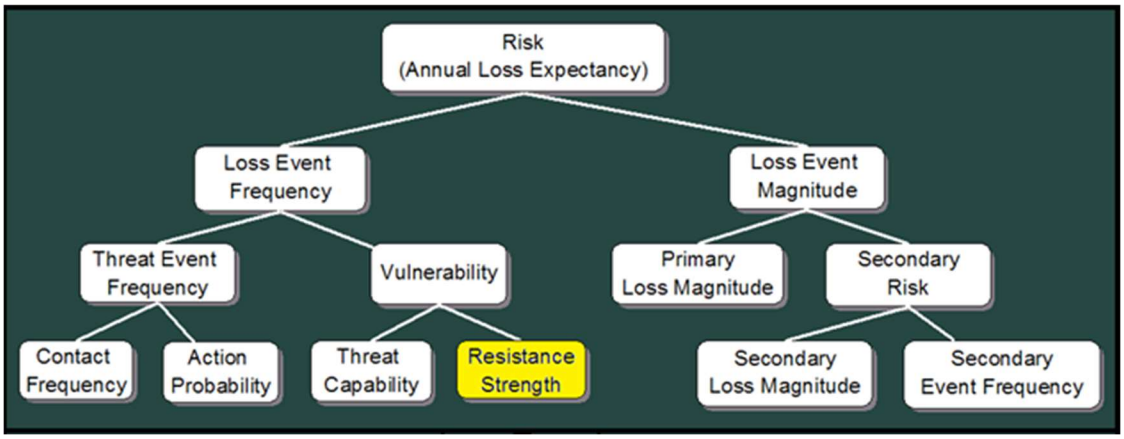

Figure 6 - Risk Factor Dependency Diagram for FAIR and Open FAIR

The Birdseye Level 2 simulator has a "FAIR-like" option that uses the "Resistance Strength" calculation. This option is enabled by default. However, the simulator also has an option to use the "Resistance" calculation depicted in Figure 5 (Page 10). Therefore, the popular ontology shown in Figure 6 (Page 14) needed to be modified to conform with the Resistance model from physics.

The Factor dependency diagram for the Birdseye simulator in Figure 7 (below) has been altered to change the highlighted Factor name from "Resistance Strength" to simply "Resistance" to indicate that a First-Order Systems approach is used for calculation. In the typical FAIR/Open FAIR ontology, "Resistance Strength" means the same thing as "Difficulty". But in the Birdseye simulator "Resistance" can be due to the presence of either a binary Control Element ("Difficulty") or a proportional Control Element ("Control Strength"). The Birdseye simulator thus allows choosing either the "Difficulty" Factor or the "Control Strength" Factor to use as the "Resistance" Factor.

Birdseye uses the Factor names in Figure 7 (below) to provide cyber risk analysts performing economic-only PRA with a better degree of interoperability with dominant cyber risk simulation products.

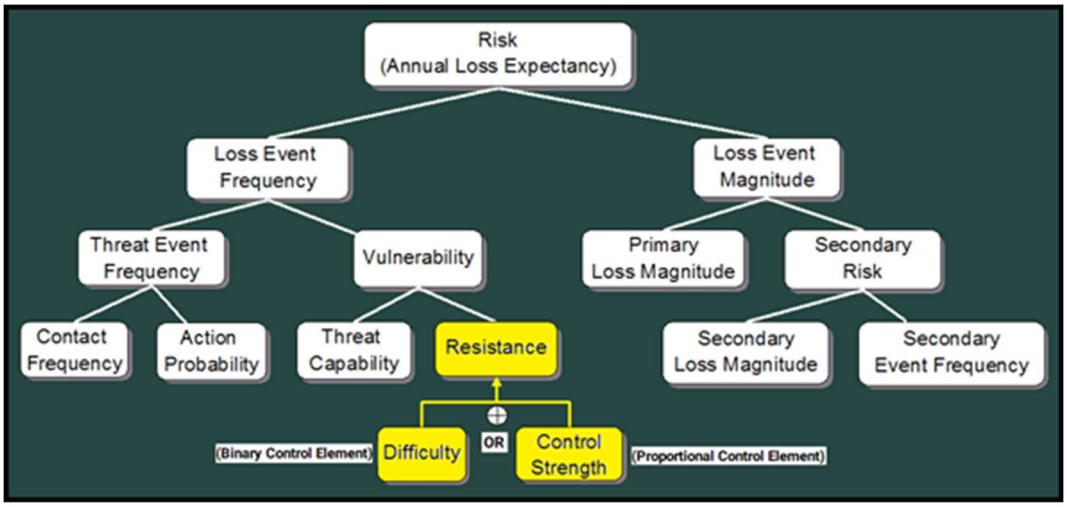

Figure 7 - Risk Factor Dependency Diagram for Birdseye

5. Updated Risk Ontology

The change from "Resistance Strength" to use the physical definition of "Resistance" from First-Order Systems theory alters the calculations related to the highlighted Factors in Figure 7 (Page 15). Consequently, it seems necessary to alter the risk ontology to allow analysts to readily distinguish between the original and modified methods of risk calculation. For the modified calculations, the Factor "Resistance Strength" is changed to "Resistance" to agree with the mathematical definition from science and engineering. The Factor "Difficulty" is used exclusively to refer to Resistance due to a binary Control Element. The Factor "Control Strength" is used exclusively to refer to Resistance due to a proportional Control Element. The risk analyst using the Birdseye Level 2 simulator should be aware of these specific differences in ontology between "Resistance Strength" and "Resistance".

Based on ontology, the risk analyst should expect the following:

- If "Resistance Strength" is used, then calculations will follow the legacy FAIR/Open FAIR approach that allows residual risk to exceed inherent risk. Only the binary model of control effect will be available. This is an "on" or "off" effect, where the derived Vulnerability depends upon whether the Threat Capability exceeds the Resistance Strength (which is a synonym for Difficulty).

- If "Resistance" is used, then calculations will follow the analogy of Resistance in physical First-Order Systems, ensuring that residual risk can never exceed inherent risk. In addition to the binary model of control effect (where the Difficulty Factor represents the control's break point), a proportional model of control effect will also be available. The proportional effect reduces derived Vulnerability in proportion to the Control Strength.

6. Physical Process Analogy for Vulnerability

In industrial Process Engineering, many physical systems are commonly modeled as First-Order systems because their behavior results from similar systems of Linear Differential Equations. These systems are also referred to as first-order lag systems or single-stage exponential systems. The *response* of these systems to any specified *forcing function* (e.g., step, ramp, or sinusoidal) can be non-linear in time and complicated.

Industrial systems are typically organized into *Control Systems*, which can consist only of Process Elements, Control Elements, Controller Elements, and Measuring Elements. A Comparator Element is also required for closed-loop, feed-back control, but in practice this element is now often combined with the Controller Element. Note that this system definition precludes the use of any Human Element in a formal Control System, because such an element (with free will) would prevent Automatic Control, which is the purpose of a Control System.

Mathematical solutions for control of common First-Order systems have been cataloged in engineering handbooks for easy reference by Control System configuration. By using *physical analogies*, common solutions for system response and control can be developed. Simply changing the variable names (by analogy) allows the same mathematical solutions (e.g., for system lag time, capacitance, Controller Element *gain*, etc.) to be readily shared by various scientific and engineering disciplines.
The most common physical analogies for First-Order systems are:

Electrical Thermal Hydraulic Chemical

For all these systems, Resistance is defined as the first derivative of Potential (or Driving Force) divided by the first derivative of Flow. For this discussion, we will be using the Hydraulic Analogy. Thus, we will consider *Vulnerability to be an analog for volumetric flow rate*. All system variables will be considered stochastic, so values will be specified by probability distributions, e.g., (10, 25, 50, λ) for a BetaPERT probability distribution.
In our First-Order hydraulic system, the Process Element will be a vertical pipe filled with toxic effluent, while the Control Element will be an obstruction that prevents flow from the pipe. For a binary Control Element, we will choose a membrane to cover the mouth of the pipe. For a proportional Control Element, we will choose a valve to close the pipe. Toxic effluent leaking from the pipe must be collected and disposed of safely (at cost).

In our Information Risk analogy to the hydraulic system, *the Threat Capability is the fluid pressure (driving force) of the toxic effluent*. *The Vulnerability is the flow rate of toxic effluent from the pipe*. *The Loss Amount is the quantity of toxic effluent that leaks from the pipe over a defined time interval. Resistance is a value that defines either:*

1) *the breaking strength ("break point") of the membrane for a binary Control Element, or*
2) *the valve setting ("set point") on a linear unit scale for a proportional Control Element.*

7. Control Element Model for "Difficulty" Factor

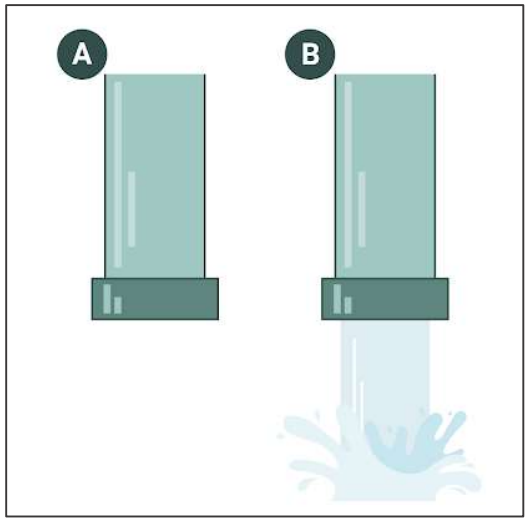

Figure 2 - Hydraulic Analogy for Difficulty

Figure 8 (at left) shows the hydraulic analogy for the Difficulty Factor used in Birdseye simulations. In this analogy, the fluid pressure of the toxic effluent in the standing pipe is the Threat Capability, and the flow rate of toxic effluent out of the pipe is the Vulnerability. **Difficulty** is the breaking strength of the membrane (which is the binary Control Element) that covers the pipe opening. **Threat Capability and Difficulty must both be defined using the same scale.** For example, in the hydraulic analogy, this would be like using psi (pounds per square-inch) to define both fluid pressure and breaking strength. Case A in the diagram occurs when the Difficulty exceeds the Threat Capability. In this case, the membrane does not rupture, and the derived Vulnerability is 0% of the uncontrolled flow rate. Case B occurs when the Threat Capability exceeds the Difficulty. In this case, the membrane ruptures, and the derived Vulnerability is 100% of the flow rate for the fluid pressure.

8. Control Element Model for "Control Strength" Factor

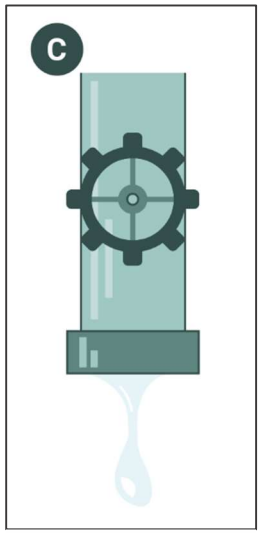

Figure 3 - Hydraulic Analogy for Control Strength

Figure 9 (at left) shows the hydraulic analogy for the Control Strength Factor used in Birdseye simulations. In this analogy, the fluid pressure of the toxic effluent in the standing pipe is the Threat Capability, and the flow rate of toxic effluent out of the pipe is the Vulnerability. **Control Strength** is the setting of the knob control for the valve (which is the proportional Control Element). Control Strength here is a dimensionless value supported over a range from zero to one. This value can easily be scaled into the range from 0% to 100%. Vulnerability is attenuated in proportion to the Control Strength by an ***inverse forcing function*** which is dependent upon the valve knob setting (i.e., the ***set point*** for the Control Element). For example, when the Control Strength is 100%, the derived Vulnerability (flow rate) is 0% of the uncontrolled flow rate. When the Control Strength is 50%, the derived Vulnerability is 50%. When the Control Strength is 0%, the derived Vulnerability is 100%. Control Strength is increased by moving the set point of the proportional Control Element toward a higher value. Moving the set point value toward a ***higher*** value moves the derived Vulnerability toward a ***lower*** value, thus reducing the Loss Event Frequency and the related Annual Loss Expectancy (see Figure 7, Page 15).

9. Converting NIST CSF Maturity Scores to Resistance

To compare the two Control Element models available in Birdseye, we will propose a hypothetical risk scenario and then examine the effect that systematically changing the Resistance (i.e., Difficulty or Control Strength) has on the ALE (Annual Loss Expectancy) for the risk scenario. We will vary the Resistance to correspond to the five levels of maturity on the Program Review for Information Security Management Assistance (PRISMA) scale for the NIST CSF. These five ***maturity levels*** translate to a ***score range*** from one to five for each control in the NIST CSF subcategories.

- Maturity Level 1: Policies
- Maturity Level 2: Procedures
- Maturity Level 3: Implementation
- Maturity Level 4: Testing
- Maturity Level 5: Integration

Some sort of conversion table is required to map the NIST Maturity Level scores for controls to Resistance probability distributions for use in the cyber risk simulation. This table could be created using expert opinion (i.e., from ***Subject Matter Experts***) to specify resistance of the scenario controls to cyber-attack.

Purely for illustrative purposes, we will use the conversion table from Adeyinka Bakare's thesis[1]. Bakare used two computer programs, ESM (Enterprise Security Management) and TPP (Technology, Process & People), to assign properties to NIST controls and rollup these scores to create a "maturity score". The maturity scores of multiple controls in NIST categories and subcategories were then further rolled up to generate a new score called "total maturity" value. This total maturity value was used to correspond to the level of compliance for aggregated multiple controls representing the strength of an organization in mitigating and resisting cyber threats. For demonstration, we will use Table 3 (below), adapted from Table 3 (page 22) in Bakare's thesis.

Resistance is (Min, Most Likely, Max).

NIST Maturity Level	Resistance
0	(0%, 0%, 0%)
1	(18%, 20%, 22%)
2	(45%, 50%, 55%)
3	(72%, 80%, 88%)
4	(81%, 90%, 99%)
5	(88%, 98%, 100%)

Table 3 - Conversion of NIST Maturity Scores to Resistance

[1] Adeyinka A. Bakare, *A Methodology for Cyberthreat ranking: Incorporating the NIST Cybersecurity Framework into FAIR Model*, (Thesis submitted to Graduate School of the University of Cincinnati, 2020), p.27-29.

10. Lambda Parameter for Probability Distributions

The hypothetical scenario used as an example here can be created easily by using the Birdseye Prototype simulator. On the "Define Scenario" page, select the "From Default" option in the "New Scenario" group box, and then click the "Create New Scenario" button, and the "Save" button. The default Birdseye scenario uses the BetaPERT probability distribution, so you must specify the four BetaPERT distribution parameters for each of four risk Factors (including "Difficulty"). We will write these parameters for a Factor in the following form:

(Minimum, Most Likely, Maximum, Lambda)

Here "Lambda" is the kurtosis (or shape) parameter for the distribution. Birdseye allows you to enter any desired value of Lambda in case you wish to use Maximum Likelihood Estimation (MLE) to compute the distribution parameters from your own historical Loss Event data (e.g., from risk registers) or if you wish to obtain parameter values from commercial reports. The Birdseye data entry page provides a handy conversion table for converting qualitative "Confidence" categories into suggested numerical values for Lambda. ***But Birdseye does not directly use*** qualitative ***Confidence settings, as that would prevent using numerical Lambda values that are statistically determined from historical data.***

For simplicity, all probability distributions in this hypothetical scenario will use a Lambda value of 4.

This value *may* correspond to "Low Confidence" on the 5-point scale for Open FAIR or to "Medium Confidence" on the 3-point scale for FAIR. Neither FAIR nor Open FAIR disclose the Lambda values corresponding to Confidence settings in their free simulators, but the value of 4 was found to give the best agreement here.

11. Hypothetical Scenario – Levels of Risk

When using the Birdseye default scenario, parameters for four risk Factors must be entered on the "Factor Data" page. The parameter values used in the hypothetical scenario are as follows:

- Primary Loss Magnitude: ($100K, $200K, $1M, 4)
- Threat Event Frequency: (1, 2, 3, 4) in Events Per Year
- Threat Capability: (50%, 75%, 100%, 4)
- Difficulty: (18%, 20%, 22%, 4) for Maturity Level 1

These parameters will determine the residual risk after adding a new control (i.e., Difficulty) to the scenario. But let us first determine the inherent risk for the current situation before adding any new control. The inherent risk can be considered as the risk for the existing security controls that affect the scenario, which requires using the "Vulnerability" Factor rather than the "Threat Capability" and "Difficulty" Factors.

But even before determining the inherent risk (at the current level of control), let us determine the risk without any loss preventive measures in place. In this case, there is no "Vulnerability" Factor, and "Loss Event Frequency" is used instead of "Threat Event Frequency" because every Threat Event will become a Loss Event. The probability distribution parameters for this version of the scenario are shown below:

- Primary Loss Magnitude: ($100K, $200K, $1M, 4)

- Loss Event Frequency: (1, 2, 3, 4) in Events Per Year

This is the level of **risk** (i.e., **Annual Loss Expectancy**) that is shown at the top of the results in Table 4 (Page 27). Because there is no control effect whatsoever to reduce the risk, this should be the highest level of risk.

The next level of risk down from this top level in Table 4 is the level of **inherent risk** with the existing control effect in place. Here are the parameters for **inherent Vulnerability** of the hypothetical scenario:

- **Primary Loss Magnitude**: ($100K, $200K, $1M, 4)
- **Threat Event Frequency**: (1, 2, 3, 4) in Events Per Year
- **Vulnerability**: **(50%, 75%, 100%, 4)**

Note that "Threat Event Frequency" is used instead of "Loss Event Frequency" here because the "Vulnerability" distribution determines how many Threat Events become Loss Events per year. The effect of Vulnerability should always reduce the uncontrolled Loss Event Frequency, so the inherent risk at this next lower level of Table 4 should never be greater than the risk at the top level of the table. ***Using the Hydraulic Analogy, this corresponds to using the flow rate from the pipe without any additional Resistance (i.e., no new membrane or valve present).***

The next step is to compute residual risk by replacing the "Vulnerability" Factor with the "**Threat Capability**" and "**Difficulty**" (or "**Control Strength**") Factors. The **Threat Capability** and **Difficulty** are used together to compute the **derived Vulnerability** for the scenario. The derived Vulnerability is used to calculate the residual risk (i.e., ***based on the newly controlled flow rate***) due to ***adding Resistance from the new preventive control***.

- **Primary Loss Magnitude**: ($100K, $200K, $1M, 4)
- **Threat Event Frequency**: (1, 2, 3, 4) in Events Per Year
- **Threat Capability**: (50%, 75%, 100%, 4)
- **Difficulty**: **(18%, 20%, 22%, 4)** for Maturity Level 1

The residual risks will be shown below the inherent risk in Table 4, with one row of residual risk for each new alternative being considered. In the hypothetical scenario, we are evaluating the effect of adding the new control at each of the NIST CSF Maturity Levels. Therefore, we will have one row of residual risk in the table for each of the alternative Maturity Levels. **The Resistance probability distribution in each row of the table will vary according to the different alternative control measures to be considered.** Inherent risk should be the upper limit for residual risk because risk should never increase with the addition of positive Resistance.

12. Hypothetical Scenario – Example Risk Analysis

This section presents example risk analysis for the hypothetical scenario at the inherent level versus the level of NIST Maturity Level 3. A separate technical paper is available from Ostrich Cyber-Risk that presents a detailed comparison of risk reduction effects of binary versus proportional Control Elements at all NIST Maturity Levels.

All histograms and charts presented in this section were generated using the Birdseye Level 2 risk scenario simulator running 1,000 simulations with the "Resistance" option selected. Data labels are displayed on the histogram bars at this low number of observations. The simulator is fully stochastic, so every run produces somewhat different results. Increasing the number of simulations per run increases both accuracy and reproducibility of results. The smallest available number of simulations per run is 1,000. Reasonably good results can be obtained at the level of 10,000 simulations per run. The tabulated results in Table 4 were generated using 1,000,000 simulations per run, and all results are the average of four independent simulation runs.

Note that all results reported in Table 1 (Page 3) are also the average of four independent simulation runs. The Open FAIR™ spreadsheet tool uses the Triangle probability distribution and so the results in the "O-FAIR" column in Table 1 differ from those in the "FAIR-U" and "B-FAIR" columns (which use the BetaPERT distribution).

12.1 Inherent Risk Analysis

This sub-section analyzes the risk at the inherent level, where there is already some existing control effect, but the proposed new Control Element is not present. Here are the probability distribution parameters used for Vulnerability in the hypothetical scenario:

- Vulnerability: (50%, 75%, 100%, 4)

Here are example Birdseye results for the inherent Vulnerability and Risk (ALE) of the hypothetical scenario:

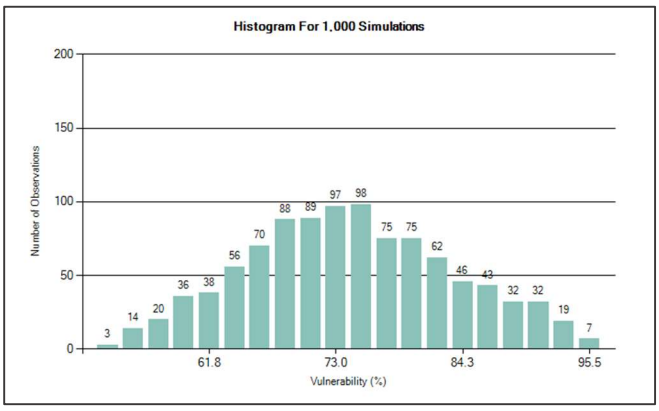

Figure 10 - Histogram for Inherent Vulnerability

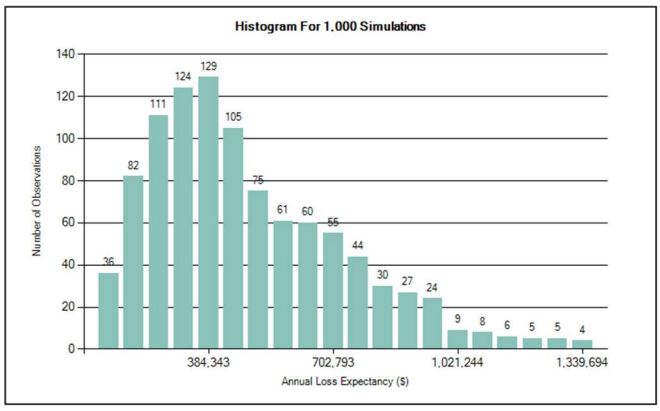

Figure 11 - Histogram for ALE from Inherent Vulnerability

The Loss Exceedance curve computed from the histogram in Figure 11 (Page 21) is shown in Figure 12 (below). Note that Risk Appetite or Risk Tolerance curves were not defined for this scenario, so these curves are not shown.

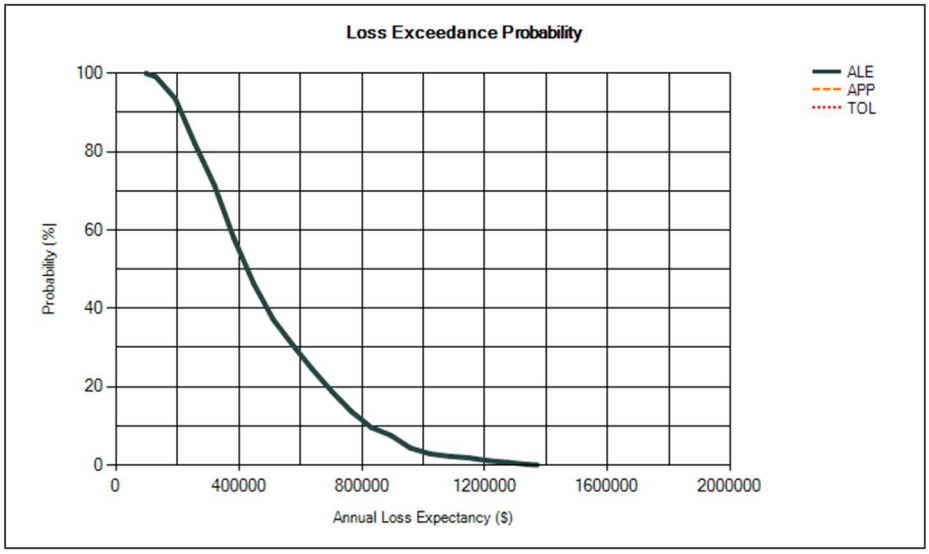

Figure 4 - Inherent Loss Exceedance Probability for Scenario

The percentile table for inherent risk (i.e., ALE) computed from Figure 12 is shown in Table 3 (below):

Percentile	ALE
0.01%	$1,371,538
0.10%	$1,369,568
1%	$1,226,518
5%	$948,768
10%	$821,086
25%	$632,729
50%	$426,719
75%	$304,140
90%	$212,341
95%	$179,754
99%	$131,029
99.90%	$98,362
99.99%	$97,737

Table 3 - Inherent Risk Percentiles for Hypothetical Scenario

Referring to the table percentile values shows there is a 10% chance of Annual Loss Expectancy being equal to or greater than $821K, while there is a 90% chance of Annual Loss Expectancy being equal to or greater than $212K.

12.2 Residual Risk Analysis for NIST Maturity Level 3 Using Difficulty

This sub-section analyzes the risk at NIST Maturity Level 3 when using a binary Control Element (i.e., when using Difficulty as Resistance). Here are the probability distribution parameters used for Threat Capability and Difficulty at NIST Maturity Level 3:

- Threat Capability: (50, 75, 100, 4)
- Difficulty: (72, 80, 88, 4) for Maturity Level 3

Here are example Birdseye results for the residual Vulnerability and ALE (Annual Loss Expectancy):

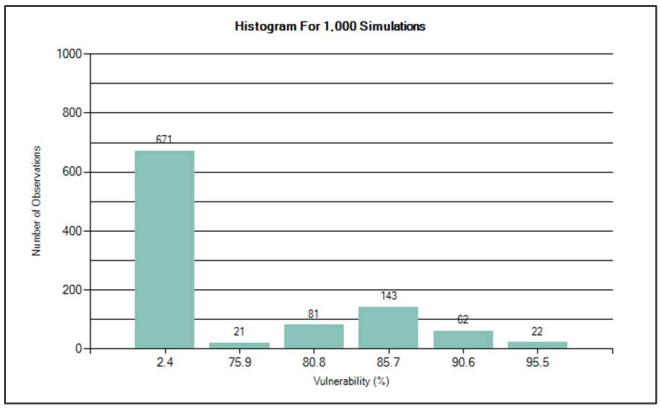

Figure 13 - Histogram for Derived Vulnerability, ML = 3

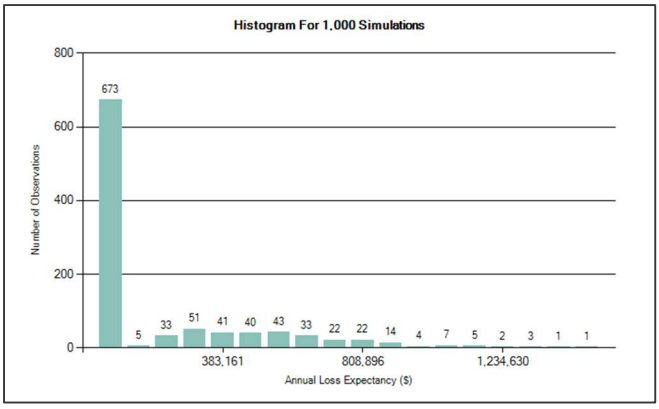

Figure 14 - Histogram for ALE from Derived Vulnerability, ML = 3

Note that the minimum Difficulty (72%) is higher than the minimum Threat Capability (50%), and the maximum Difficulty (88%) is lower than the maximum Threat Capability (100%). Thus, we expect that some fraction of the threat attempts will fail (producing observations with 0% Vulnerability for the histogram), and some fraction of the threat attempts will succeed (producing observations with non-zero Vulnerability). In Figure 13, you can see that about 67% of the simulations had 0% Vulnerability, and thus about 33% had non-zero Vulnerability.

Figure 14 shows the ALE distribution resulting from the derived Vulnerability distribution. The fraction of the simulations with 0% Vulnerability are all in the "spike" at the first column in the ALE histogram because there was no Loss for these simulations. The remaining columns in the ALE histogram chart are due to the observations that had non-zero Vulnerability, resulting in some distributed non-zero Loss amounts.

The degenerate frequency distribution for ALE shown in Figure 14 can be considered to consist of two separate distributions: one distribution for the "spike" at the left-most edge of the histogram (from 0% Vulnerability), and one distribution for the "hump" that follows after the spike (from non-zero Vulnerability).

Considering the ALE probability distribution as a two-piece distribution helps to explain the Loss Exceedance curve shown in Figure 15 (below). The sharply dropping line at the left-most edge of the chart is due to the fraction of simulations with 0% Vulnerability. In fact, you can read the ratio of Zero% to non-Zero% Vulnerability in simulations directly off Figure 15. Observe that the left-most "spike" of the curve

drops from 100% Probability to about 33% Probability. Therefore, the amount of "probability" in this part of the curve is about 67% (i.e., 100% minus 33%), so the fraction of simulations found in the 0% column of the histogram should be about 67%. This is confirmed by the histogram in Figure 13 (see Page 23).

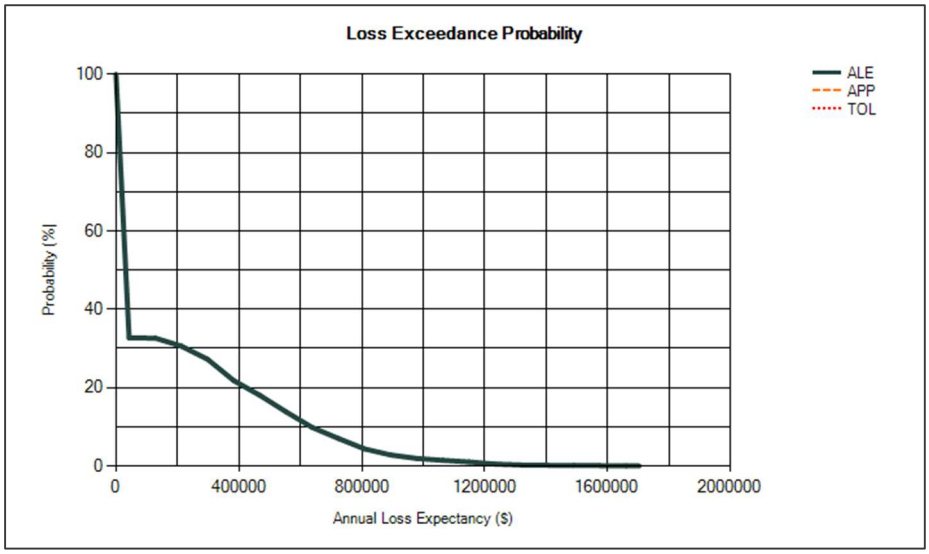

Figure 15 - Residual Loss Exceedance Probability, Maturity Level = 3

The "plateau" part of the curve that follows the initial "spike" is due to the degenerate probability distribution. The smallest Loss value in the fraction of simulations with non-zero Vulnerability determines the right-most edge of the "plateau". Increasing or decreasing the Minimum value in the probability distribution for the Loss Magnitude will make this plateau wider or narrower. The curved slope that follows to the right of the "plateau" is, of course, the typical Loss Exceedance curve for the fraction of simulations with non-zero Vulnerability (due to failure of Difficulty to contain Threat Capability).

12.3 Residual Risk Analysis for NIST Maturity Level 3 Using Control Strength

This sub-section analyzes the risk at NIST Maturity Level 3 when using a proportional Control Element (i.e., when using Control Strength as Resistance). Here are the probability distribution parameters used for Threat Capability and Control Strength at NIST Maturity Level 3:

- Threat Capability: (50, 75, 100, 4)
- Control Strength: (72, 80, 88, 4) for Maturity Level 3

Example histograms for residual Vulnerability and ALE for this scenario are presented below:

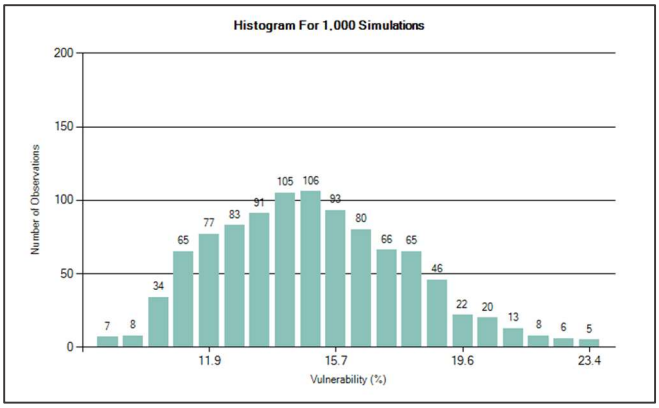

Figure 16 - Histogram for Derived Vulnerability, ML = 3

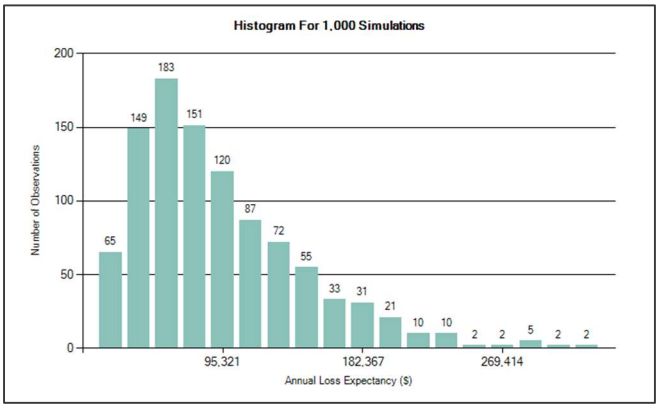

Figure 17 - Histogram for ALE from Derived Vulnerability, ML = 3

Because the Most Likely value for Control Strength is 80% at this Maturity Level, we expect the residual Vulnerability to be about 20% (i.e., 100% minus 80%) of the inherent Vulnerability. Comparing Figure 10 (see Page 21) to Figure 16 (above), you can see that the right-most bin in the histogram has dropped from 95.5% to 23.4%. Also, the residual distribution of Vulnerability appears to be compressed somewhat more in comparison to the inherent distribution, but both distributions have approximately the same shape.

Comparing Figure 11 (Page 21) to Figure 17 (above), you can see that the residual distribution for ALE also has approximately the same distribution as the inherent distribution, and the residual distribution has also been shifted to the left (toward zero). Thus, the frequency distributions for both residual Vulnerability and ALE simply appear to be attenuated versions of the inherent distributions.

The Loss Exceedance curve for a proportional Control Element is shown in Figure 18 (Page 26). We would likewise expect this curve to be simply an attenuated version of Figure 12 (Page 22). Examining Figure 12, you can find the Median (not the Mean!) value for the ALE distribution by intersecting the ALE curve at the 50% Probability line. This gives an apparent Median ALE of about $427,000 for inherent risk. Examining Figure 18 (below) to find where the 50% Probability line intersects the ALE curve gives a much-reduced Median Annual Loss Expectancy of about $80,000 for residual risk at NIST Maturity Level 3.

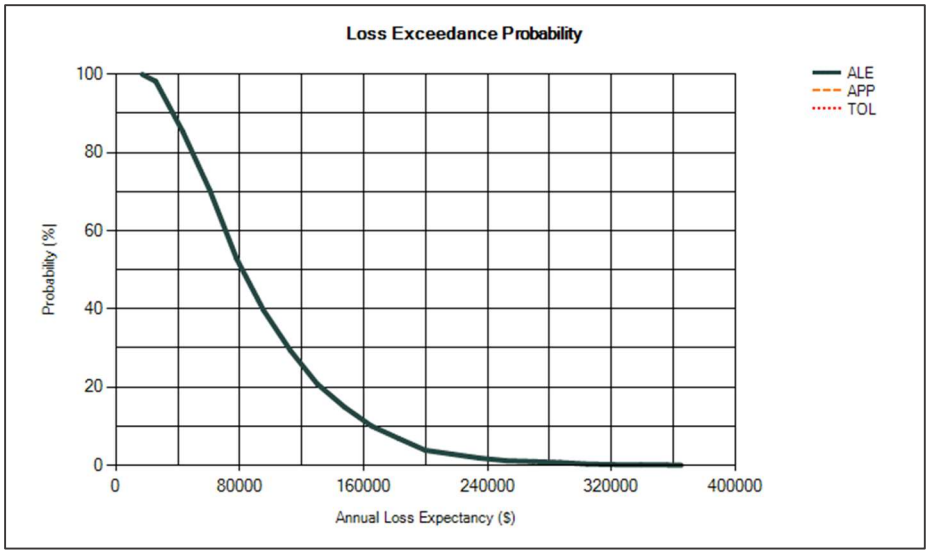

Figure 5 - Residual Loss Exceedance Probability, Maturity Level = 3

This comparison shows that increasing Control Strength (i.e., increasing the *set point* for a proportional Control Element) appears to attenuate the inherent ALE curve by shifting it to the left (toward zero) and compressing it (thus making the curve slightly steeper) to produce the residual ALE curve.

13. Summary – Risk Comparison for All NIST CSF Maturity Levels

The simplest way to compare risk between any two scenarios is by comparing the average (mean) and maximum values of the ALE (Annual Loss Expectancy) probability distribution. Using the Birdseye scenario simulator, it is easy to construct a table to compare the inherent and residual risks for the two available Control Element models. Table 4 (below) shows the various NIST Maturity Levels and their associated Resistance probability distributions (from Table 3, Page 22). The average and maximum ALE values are compared for both the binary Control Element model (Difficulty) and the proportional Control Element model (Control Strength).

In Table 4, values in the Vulnerability row are the *inherent risk* found by using Vulnerability instead of derived Vulnerability (which is computed from Threat Capability and Resistance). The *residual risk* (based on derived Vulnerability) is shown for each of the five NIST Maturity Levels. The values in the row for NIST Level 0 are approximated by using the probability distribution (0, 0, 1, 4) instead of (0, 0, 0, 4).

		Binary CE		Proportional CE	
Primary LM = ($100K, $200K, $1M) Loss Event Frequency = (1, 2, 3)		Ave ALE	Max ALE	Ave ALE	Max ALE
		$633,250	$2,575,500	$633,250	$2,618,750
Threat Event Frequency = (1, 2, 3) Vulnerability = (50%, 75%, 100%)		Inherent Risk		Inherent Risk	
		$475,000	$2,260,000	$475,000	$2,310,250
Threat Event Frequency = (1, 2, 3) Threat Capability = (50%, 75%, 100%)		Residual Risk Shown Below		Residual Risk Shown Below	
Control Maturity	Resistance	Ave ALE	Max ALE	Ave ALE	Max ALE
NIST Level 0	(0%,0%,1%)	$475,000	$2,325,250	$474,000	$2,244,000
NIST Level 1	(18%,20%,22%)	$475,000	$2,278,500	$380,000	$1,815,500
NIST Level 2	(45%,50%,55%)	$475,000	$2,305,000	$237,500	$1,144,000
NIST Level 3	(72%,80%,88%)	$174,000	$2,246,000	$95,000	$528,750
NIST Level 4	(81%,90%,99%)	$41,000	$2,265,750	$47,500	$347,250
NIST Level 5	(88%,98%,100%)	$3,000	$2,162,250	$16,000	$182,000

Table 4 - Comparison of Effect of NIST Maturity Level on Annual Loss Expectancy for Hypothetical Scenario

The simulated Average ALE values in the table converge much more quickly than the Maximum ALE values, so there is still some variability in the reported Maximum ALE values between runs. We would theoretically expect the residual risk at the NIST Level 0 to correspond to the inherent risk, because Level 0 effectively applies no additional Resistance. Note that the Average ALE values in Table 4 agree with this theoretical expectation. The Average ALE value for the Proportional Control Element at NIST Level 0 is slightly lower because the probability distribution for Resistance at NIST Level 0 was approximated as described above.

13.1 Summary for Difficulty (Binary Control Element)

An interesting feature of the progression of average residual ALE values is that they remain constant at the value of $475K (i.e., the average inherent ALE) for a wide range of Maturity values. The average residual ALE values only begin to decrease when the probability distribution for Difficulty begins to overtake the distribution for Threat Capability. For example, there is no apparent effect on average ALE when going from NIST Maturity Level 1 up to Level 2 for this scenario because the minimum Threat Capability is greater than the maximum Difficulty for both levels. When the distributions for Difficulty and Threat Capability begin to overlap significantly, then average ALE begins to decrease dramatically, as in going from Maturity Level 2 to Level 3. After the Difficulty distribution has passed over the Threat Capability distribution, then the average ALE value will drop to almost nothing (i.e., the average value for 0% Vulnerability most of the time). Thus, there is relatively little decrease in average ALE when going from Maturity Level 4 to Level 5 for this scenario. Even at Maturity Levels where the average risk is negligible, the maximum risk can still be very high due to the very small (but possible) chance of control failure allowing at least one simulation run with a high Threat Capability when both the Threat Event Frequency and Primary Loss Magnitude are also high.

Average ALE may or may not decrease as Difficulty increases, depending upon the relative characteristics of the Threat Capability and the Difficulty probability distributions.

13.2 Summary for Control Strength (Proportional Control Element)

The progression of ALE values when using the proportional Control Element is entirely intuitive. Conceptually, when the Control Strength is 0%, then the derived Vulnerability (volumetric flow rate) is at its maximum for the given Threat Capability (fluid pressure). When the Control Strength is 100%, then the derived Vulnerability would be zero.

As Control Strength increases incrementally, the derived Vulnerability decreases incrementally. The rate of decrease in Vulnerability is initially very low, but then that rate (of decrease) increases as the Control Strength is increased. At some value of Control Strength (here around Maturity Level 3) a maximum rate (of ALE decrease) is attained, and thereafter the rate of ALE decrease begins to decrease as Control Strength continues to increase.

Average ALE always decreases as Control Strength increases.

14. Conclusion

The Birdseye Level 2 scenario simulator can easily be used to calculate the effect of NIST CSF Maturity Levels on risk reduction. The default simulator setting is to use the only preventive security control model available in the popular FAIR/Open FAIR approach to risk simulation, i.e., the "Resistance Strength" calculation with Difficulty modeled as a binary Control Element. An optional simulator setting provides two alternative ways of modeling the Resistance of preventive security controls based on analogy to physical First-Order Systems.

In addition to comparing the risk reduction effect of NIST CSF Maturity Level, a comparison is also made here between the risk reduction effects of these two alternative Control Element models that are based on physical Resistance. This same comparison of risk reduction using the demonstrated Levels of Risk can be performed with any risk scenario simulator that supports the typical PRA "Factors" in the FAIR and Open FAIR frameworks.

The selection of a Control Element model should never be arbitrary, but rather should always be based upon consideration of the control effect of the security control(s) affecting the cyber risk simulation. If the control effect for certain preventive security controls is unknown, then the binary Control Element (Difficulty Factor) can be chosen for compatibility with risk analysts using the FAIR/Open FAIR frameworks.

The Difficulty Factor (i.e., a binary Control Element) should be preferred if the security control(s) that affect the simulation are of the type where investment in increasing Resistance (i.e., Maturity Level) provides little or no return until a certain threshold value is achieved. A drastic reduction in risk follows when Resistance increases above this threshold. While the Threat Capability remains below the "break point" for these controls, risk will be negligible.

The Control Strength Factor (i.e., a proportional Control Element) should be preferred if the security control(s) that affect the simulation are of the type where every significant incremental increase in Resistance (Maturity Level) produces some reduction in risk. The risk mitigation effect of these controls will follow a logistic curve, with increasing returns on investment in improving Resistance up until a certain value is reached, followed by diminishing returns on investment thereafter. The "set point" for these controls determines the proportional reduction in derived Vulnerability, which produces a related reduction in risk.